



Intro to Image Generation

Turn text into pictures — and short videos — with modern AI tools.

What are image generation AI tools?

Computer programmes that use AI to create images from text descriptions (or other inputs).

They learn from massive collections of images, recognising patterns and styles. When you describe what you want, the model generates an image that matches.

The result: anyone — designer, marketer, or curious beginner — can sketch a visual idea in seconds without picking up a pencil or opening Photoshop.




Image generated by AI

SECTION 1 OF 6

Where it's used

Real-world applications



Every image here was AI-generated.



Concept development

Character art, UI sketches, mood boards.



Marketing

Visuals, ad concepts, tailored social content.



Game development

2D assets ready to drop into game engines.



Fashion & retail

Generate styles and 'try-on' looks.



Industrial design

Product mock-ups in realistic settings.



Graphic design

Logos, posters, illustrations, packaging.

The 2026 tool landscape

There are dozens of image-generation tools. Most fall into one of four buckets:

All-rounders

Midjourney v7, Leonardo.ai, ChatGPT (GPT Image 2).
Type a prompt, get an image.

Pro editors

Adobe Firefly Image 4 inside Photoshop 2026 — generative fill, expand, remove. Now ships with FLUX.2 Pro and Nano Banana Pro alongside Firefly.

Video generators

Sora 2 (OpenAI), Veo 3.1 (Google), and aggregators like Higgsfield AI bundle 15+ video models in one subscription.

Specialist & open

FLUX.2 Pro for photorealism, Stable Diffusion 3.x for self-hosting, custom-trained models for brand-specific looks.

Today we'll use Leonardo.ai — a friendly all-rounder that demonstrates the core ideas you'll see across every modern tool.

What we'll use today: Leonardo.ai


Leonardo.ai

A popular web-based image generator with a full suite of generative-AI features.

Why Leonardo for this stream?

- Free tier is generous enough for learning.
- The interface mirrors what you'll see in most other tools (prompt bar, model picker, style settings).
- It bundles still-image generation and short-video generation in one place.

We'll explore only the basics today — Leonardo has many more advanced features that build on these foundations.



Leonardo.ai homepage (screenshot may differ — UI evolves).

Activity overview



Generate an image

Use a simple text prompt — and see what comes back.



Add motion

Turn one of your still images into a short video clip.



Notice the limits

Spot common errors, biases, and safety blocks in action.

These cover the basics. Most image generators have more powerful options — but these three steps give you a good sense of what AI can and can't do.

SECTION 2 OF 6

Generate an image

Getting set up

1. Open Leonardo.ai


On the NAIS laptops, use the Chrome 'favourites' bar — the Leonardo tab is already pinned with the NAIS login filled in.

2. Click 'Image Creation'

It sits on the home page next to Realtime Canvas, Motion, Upscaler and Canvas Editor.

3. You're ready to prompt

If you're using your own laptop, you can sign up for a free Leonardo account and follow along — most features in this demo work on the free tier.



Leonardo's main navigation: Realtime Canvas · Realtime Gen · Motion · Image Creation · Upscaler · Canvas Editor.

Writing a simple prompt

Most image generators have a single prompt bar at the top. Type a description, hit Generate, and the model returns 1–4 images.

Try this:

“A blue cat sleeping in a library.”

Simple, vivid, specific. Good prompts read like camera directions.

Three quick prompting tips

- Be specific: subject + setting + style.
- Use sensory words: “soft light”, “mist”, “golden hour”.
- Iterate: change one thing at a time and re-generate.



Leonardo's output for the prompt above.

Generation options

On the left of the prompt screen you'll see a sidebar of controls. Most of these aren't critical for a first try — but knowing what they do helps later.

Model

Which AI does the drawing. Different models = different styles.

Prompt Enhance

Auto-adds detail to your prompt. Big impact on the result.

Style

Photoreal, anime, illustration, cinematic — preset moods.

Dimensions

Square, portrait, landscape, widescreen. Match it to where the image will live.

Number of images

Generate 1, 2, 3 or 4 variations from one prompt.

Private mode

Hide your generations from the public feed (paid tiers).

SECTION 3 OF 6

Generate motion


Generating a short video

Leonardo can turn a still image into a short video using its **Motion** feature.

From the home page, click the Motion button (next to Image Creation).

Heads-up: Motion isn't free. Use the NAIS laptops (we have a paid plan) or sign up for a Leonardo subscription on your own account.

*Around the wider 2026 landscape, **Sora 2** (OpenAI), **Veo 3.1** (Google), and aggregators like **Higgsfield AI** are the headline video models — Leonardo's Motion feature is a friendly entry point.*



Motion sits between Realtime Gen and Image Creation in Leonardo's nav bar.

Video models in 2026



Sora 2

OpenAI. Object permanence, fabric/liquid physics. 4–12 s clips at 720p.



Veo 3.1

Google. 1080p video from text or reference images, up to 8 s per clip.



Higgsfield AI

Aggregator: Sora 2, Veo 3.1, Kling 3.0, Seedance 2.0 — one subscription.

Motion settings

Three clicks to a short clip:

1

Pick a source image

Choose one of your Leonardo generations — or upload your own picture.

2


Adjust motion strength

Slider on the right: drag right for more movement, left for subtler animation.

3

Press Generate

Wait a few seconds — your clip appears in your library alongside your images.



Leonardo's Motion dialog (source-cited screenshot).

SECTION 4 OF 6

Understanding limitations

Generation errors

Image generators aren't perfect. Knowing what's hard to draw saves you generations (and tokens).

Common failure modes

- **Hands & fingers** — too many, too few, or fused together.
- **Text inside images** — gibberish letters, even with simple words.
- **Perspective & straight lines** — bricks bend, windows tilt, horizons drift.
- **Pattern bleeding** — a shirt's stripes leak onto the wall behind.
- **Repeated subjects** — ask for two people, you might get four limbs.



Hands: fingers fused or duplicated.



Text: looks like writing, isn't.

Bias in training data



A generic prompt yielded this couple — note the narrow beauty standards.

What is the model leaving out?

Image models learn from enormous training datasets — and any imbalance in that data shows up in their outputs.

Common patterns:

- Vague prompts default to stereotypes (occupation, gender, ethnicity).
- ‘Average’ people are rare — outputs skew toward conventionally attractive, Instagram-style looks.
- Cultural blind spots show up in clothing, architecture, landscapes.

Your defence: be explicit. Prompts that name age, body type, ethnicity, profession, and setting return more varied — and often more honest — results.

Safety filters and blocked imagery

What gets blocked?

Every commercial generator screens prompts and outputs for unsafe, inappropriate, or copyrighted material.

Why it matters:

- Protects users, the provider, and the public.
- In many jurisdictions, this screening is now a legal requirement — including under the EU AI Act, which is fully applicable from 2 August 2026.
- A single trigger word can block an otherwise reasonable prompt.

Imperfect by design

These systems are hypervigilant but they make mistakes. Expect occasional false blocks on innocent prompts — rephrase and try again.



Some prompts hit a wall — by design.

SECTION 5 OF 6

Other features

Realtime Canvas


Draw, watch it become art.

Realtime Canvas generates a finished image, live, based on the rough sketch you doodle in the left panel.

Useful for:

- Locking compositions to specific shapes or colours.
- Exploring variations without typing a thousand-word prompt.
- Free, unlimited use — no token costs.

Give it a go. It's the most tactile way to experience what these models actually do.



Sketch on the left → AI render on the right (Leonardo Realtime Canvas).

Realtime Generator: Elements


Watch the image update as you type.

Realtime Generator re-renders your image every time you change the prompt — or move an **Elements** slider.

Why this is a great learning tool:

- Each Element is essentially an extra prompt with a weight.
- Try positive, negative, and zero values — see the model react in real time.
- The same idea (weighted prompts) is hidden inside every other image generator.

Play with the Elements panel. Notice how a small change in one slider can completely transform the result.



Elements panel: weak / strong / negative sliders for each style.

Summary

Today you turned text into pictures, pictures into video, and saw where the seams still show. Modern image-generation AI is creative, fast, and imperfect — exactly the combination that makes it powerful for sketching ideas.

Prompts are camera directions

Subject + setting + style. Iterate one change at a time.

The 2026 toolkit is broad

Leonardo for learning, then Midjourney v7, Firefly Image 4, FLUX.2 Pro, Sora 2 and Veo 3.1 as you grow.

Know the limits

Hands, text, perspective, bias, safety blocks — expect them, plan around them.