

I've created thousands of AI images and these are the best AI image generators of 2024

[Buying Guides](#)

By [Ryan Morrison](#)

last updated 2 days ago

Create stunning art from text or images



(Image credit: Adobe/Midjourney/Ideogram)

I have created a lot of content using generative AI platforms. One of the first things I do every morning is open Ideogram and create something random, [animate it in Runway](#) and create a [voiceover with Hedra](#). I have tested multiple AI models, particularly image models and they each have something unique to offer.

AI image generation has come a long way in a few short years. The first version of [Midjourney](#) came out in February 2022. It could barely create a realistic-looking human. Now, almost three years on its output is indistinguishable from reality.

Recent updates

Update Dec 17: We put Google's impressive [Imagen 3 AI image model to the test](#) with a series of 7 prompts.

In a rapidly evolving market, a good model isn't enough anymore though. What makes a platform stand out is the additional features including tools like [Ideogram's Canvas](#), upscaling in Freepik or [Midjourney's Editor](#) which also includes editing external images. Working out which model to use, especially when you are parting with money to get the best capabilities, can be a challenge. Hopefully, this guide will help. As well as simply using the image generation platforms, I also [put them to the test with a range of prompt styles](#), image types and criteria to see not only how they perform, but also how they compare to one another.

LATEST VIDEOS FROM TOMSGUIDE

I've run realism tests, which image generator [animated best with Runway](#), which creates the best text and which produces the most lifelike people. All of these helped me come up with this guide to the best AI image generators available today.

How we test AI image generators

Why you can trust Tom's Guide Our writers and editors spend hours analyzing and reviewing products, services, and apps to help find what's best for you. [Find out more about how we test, analyze, and rate.](#) The most simple test of any tool is to use it — a lot. I create images as part of my everyday work in all of the big AI image generators, especially the ones I've included here. I then run comparison tests, and prompt tests, and try to see what they can't do.

I play with different tools and features not part of the core model and see how the output interacts with other tools such as [AI-generated video or lip-synching](#). Each model has its own unique prompting quirks as well, and I've tried to find the best of those.

For this guide, I've focused on platforms rather than underlying models. Some models, like Flux and Stable Diffusion, are widely available on services like FAL, NightCafe, Leonardo and Freepik. Others are single-platform models. However, that is changing with the rise in APIs. For example, you can get Imagen 3, DALL-E 3 and Ideogram on NightCafe now.

For each of the generators below I came up with five prompts to put them to the test. You can see the resulting images in a gallery for each section. For Leonardo, these were made with Phoenix, NightCafe images were made using Flux 1.1 Pro and Freepik images were generated using its own Mystic model.

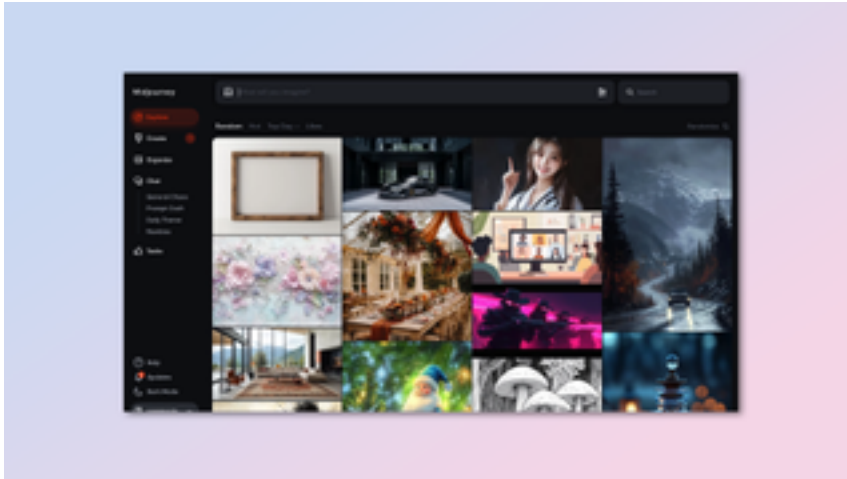
REVIEWED BY

Ryan Morrison

AI Editor

Ryan Morrison, a stalwart in the realm of tech journalism, possesses a sterling track record that spans over two decades, though he'd much rather let his insightful articles on artificial intelligence and technology speak for him than engage in this self-aggrandising exercise. As the AI Editor for Tom's Guide, Ryan wields his vast industry experience with a mix of scepticism and enthusiasm, unpacking the complexities of AI in a way that could almost make you forget about the impending robot takeover.

BEST AI IMAGE GENERATOR OVERALL



(Image credit: Midjourney)

1. Midjourney

Impressive photorealism from image generation

SPECIFICATIONS

Free plan: N/A

Paid plans: Start at \$10/month

REASONS TO BUY

Text on image

Granular control

Impressive photorealism

Powerful web UI

REASONS TO AVOID

No free plan

While [Midjourney](#) doesn't have the range of models or video capabilities of the big platforms like FAL, NightCafe and Leonardo, it is tried and tested. It is the power user AI image generator with high levels of realism and an impressive degree of control over the output.

It took me a while to get to grips with the huge number of parameters and features available through Midjourney, especially while it was still inside Discord, but once you've mastered the fundamentals it is a remarkable application.

Some of the more talented users can get it to create pictures that look like they're straight from a phone camera. Midjourney was one of the first to crack the finger problem and has consistently real-looking people. With version 6.1 the company added new skin textures to further improve human realism.

Midjourney is somewhat controversial for its refusal to discuss the source of its training data. Many suspect much of it comes from scraping any publicly available images it could find regardless of whether it had permission from the image creators.

What makes Midjourney stand out for me though is the level of control you have over every aspect of the generation. You can use parameter commands to reference the style or a character within another image or use others to completely change how an image appears.

Image 5 of 5

(Image credit: Midjourney/Future AI)

Prompt: *A vibrant, candid street photo of people in a busy London market during autumn, focusing on diversity and interactions between friends and families. The scene should capture realistic clothing, natural smiles, and market stalls brimming with fresh produce and handmade items. Golden autumn light softly illuminates the scene, creating a warm atmosphere.*

(Image credit: Midjourney/Future AI)

Prompt: *A futuristic eco-friendly skyscraper in central Tokyo. The building incorporates lush vertical gardens on every floor, with cascading plants and trees lining glass terraces. Solar panels and wind turbines are integrated into the structure's design, reflecting a sustainable future. The Tokyo Tower is visible in the background, contrasting the modern eco-architecture with traditional city landmarks.*

(Image credit: Midjourney/Future AI)

Prompt: *A beautifully plated vegan gourmet dish featuring a colourful arrangement of roasted vegetables, microgreens, and edible flowers on a minimalist ceramic plate. Each ingredient should have intricate detail, capturing a mix of textures like roasted, crispy, and creamy. The backdrop is soft and blurred, making the vibrant colours of the food the central focus.*

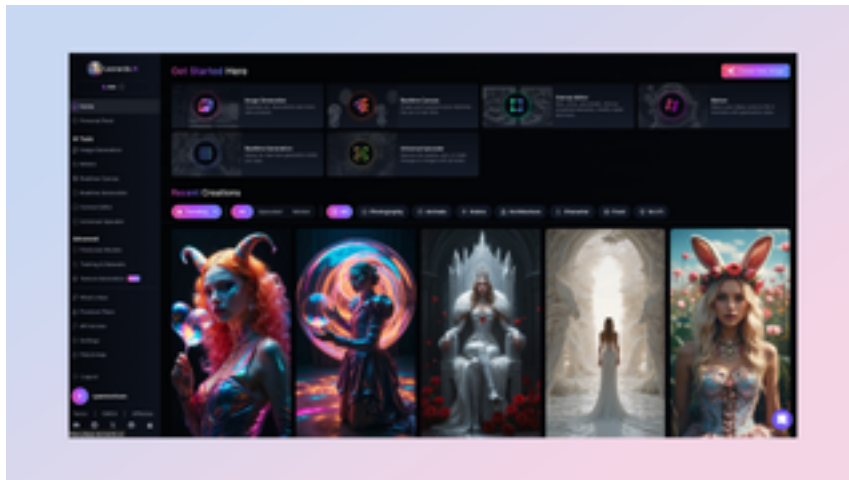
(Image credit: Midjourney/Future AI)

Prompt: *An elegant snow leopard perched on a cliff in the Himalayan mountains, surrounded by swirling snow. The animal's fur is intricately detailed with distinctive patterns and a thick winter coat. The scene captures the majesty and isolation of the leopard's habitat, with mist and mountain peaks fading into the background.*

(Image credit: Midjourney/Future AI)

Prompt: *A retro-futuristic sci-fi poster with a neon-coloured space scene featuring a large, glowing planet in the foreground and multiple smaller planets in orbit. Include a spaceship flying towards the largest planet, with a trail of light behind it. Add stylised text at the top saying, "Journey Beyond" in a bold, vintage font, and at the bottom, a tagline: "Explore the Stars - 2089." The image should evoke a mix of nostalgia and futuristic exploration.*

BEST FOR FEATURES



(Image credit: Leonardo)

2. Leonardo

A great all round platform for generative content

SPECIFICATIONS

Free plan: 150 images

Paid plans: Start at \$10/month

REASONS TO BUY

Free version

Video mode

Powerful Phoenix model

Granular control over image generation

REASONS TO AVOID

Complex user interface

[Leonardo](#) started life as a well-done wrapper for the variety of Stable Diffusion models but it has evolved into so much more than that with the addition of the powerful new Phoenix model. It still has the SD models but offers them with custom styles and fine-tuned versions of models.

Leonardo has also joined Affinity and become part of the Canva family. The foundation models built by the Australian image startup will likely be integrated into Canva products and the sale will give Leonardo the resources to expand and build more impressive generative AI features.

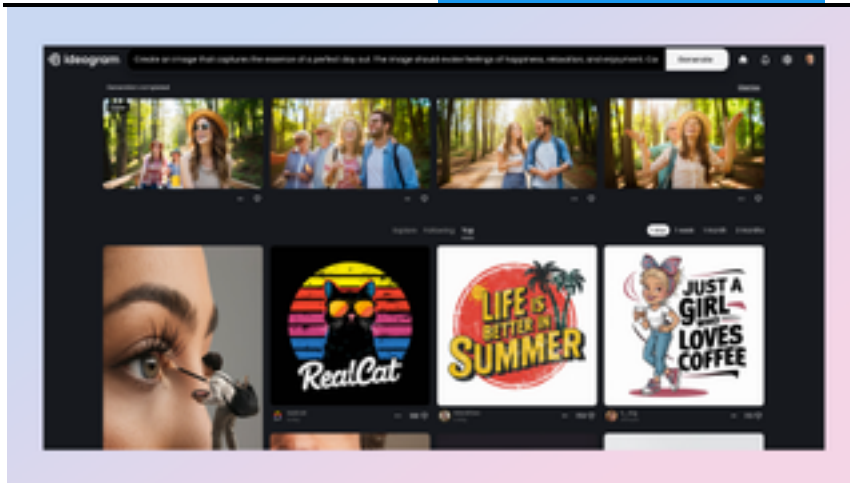
It is already up there with the best. Leonardo's ability to generate photorealistic images, thanks to the fine-tuned PhotoReal model, is nearly up to Midjourney standards and it can create a range of styles with the Elements feature.

These elements are a fine-tuned model that can be applied before generation and steer the image to be created with a particular look such as a sketch or sculpture. You also have the ability set a style such as cinematic, food or long exposure.

What really makes Leonardo stand out for me is the combination of an easy to use UI, mixed with an incredible amount of control. You can add reference images and set how the AI should use them, control the size and layout, even add a transparent background.

Most of these features are available in other platforms, but Leonardo has them all as well as a range of other features such as image upscaling, live image generation and one of the most creative tools — the ability to draw a sketch and have the AI turn it into a full image.

BEST FOR TEXT ON IMAGES



(Image credit: Ideogram)

3. Ideogram

Create perfect text on images most of the time

SPECIFICATIONS

Free plan: 25 prompts per day

Paid plans: Start at \$8/month

REASONS TO BUY

Impressive photorealism

Excellent prompt adherence

Consistent text on images

REASONS TO AVOID

Limited fine control

[Ideogram](#) is one of my favorite AI image generators for my own personal use. Version 2 gave us a new set of features that put it on par with Midjourney, as well as some not available elsewhere such as color pallets and a 'design' style.

It follows a prompt really well and can add text like no other model. I've been able to generate full movie posters, flyers and greeting cards with accurate text.

With the new Canvas feature you can edit and manipulate generated images using extended text prompts. This is great for correcting text or graphic design.

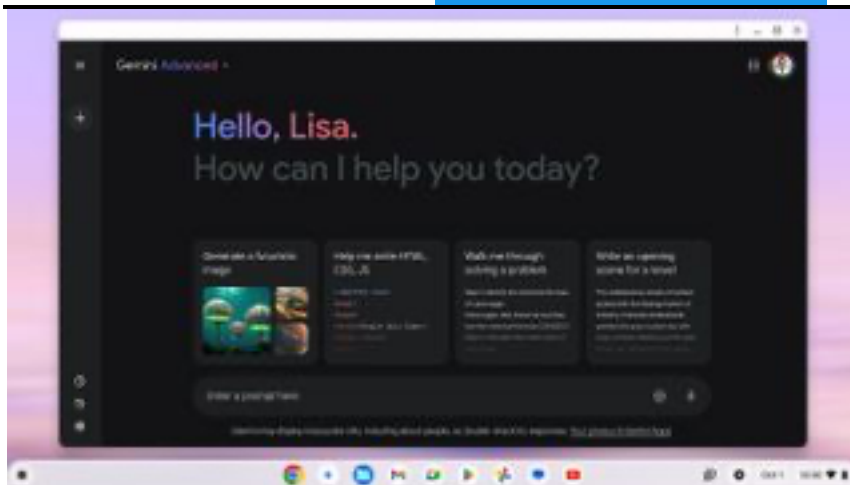
Accessed through a wonderfully simple prompt box, with the option to automatically enhance your prompt to get a better image, it is both easy to use and powerful.

While it is best for adding text on to images, it also has a bit of Midjourney flare to the style of the works it generates. You can turn off magic prompt and create more artistically simple images, or even add customized style tags.

Ideogram's Magic Prompt is a compelling feature. If its turned on then a large language model will analyze your prompt and re-write to be much more descriptive to get closer to your vision.

You can see your original prompt as well as the magic prompt for any image, adapt it or use it to create a new picture. You can also use any generated image as a source for a new image.

BEST FOR COLLABORATION



(Image credit: Google)

4. Google Gemini

Free to use with a Google account

SPECIFICATIONS

Free plan: Yes

Paid plan: Additional generations for \$20/month

REASONS TO BUY

Easy editing

Follow-up prompting

Can use natural language prompts

REASONS TO AVOID

Minimal control over generations

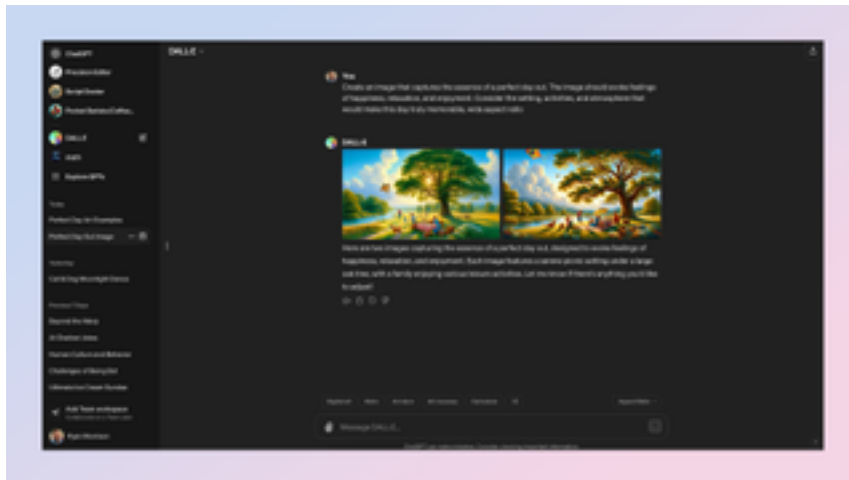
[Google](#) Gemini has had a rocky road when it comes to image generation. It uses the Google DeepMind developed Imagen family of AI image models to create pictures with the Gemini chatbot acting as a go-between, enhancing prompts and making changes.

This led to some issues where it was generated unexpected images of people. Google blocked Gemini from making pictures of people — real or fake — for a while and limited the rollout of image creation globally until it was fixed.

With the launch of Imagen 3 earlier this year Gemini was given its image creation ability back and, while it can only make square format pictures, its collaboration capabilities are what make it stand out.

Once you get the first round of pictures from a prompt you can then ask Gemini to change any single aspect. In one test I was able to ask it to generate a grand piano, then add a cat at the keys and finally turn it into a full band with a giraffe on bass and dog singer.

BEST FOR INTERACTION



(Image credit: ChatGPT)

5. OpenAI ChatGPT (DALL-E 3)

As simple as talking to ChatGPT

SPECIFICATIONS

Free plan: N/A

Paid plan: \$20/month with ChatGPT Plus

REASONS TO BUY

Text on images

Prompt-based editing

Natural language prompting

REASONS TO AVOID

No free plan

Doesn't always follow prompts

DALL-E 3 is only available in [ChatGPT](#) to those with a Plus account. There are a few ways to use DALL-E within ChatGPT. You can just access it through the main interface, through the DALL-E GPT custom chatbot, or by tagging DALL-E in the main chat.

The original DALL-E was one of the first high profile commercial generative AI image tools. Initially available as an API or through a dedicated DALL-E page, [OpenAI](#) has since bundled it with its chatbot. This is also its primary selling point — the ability to talk through an image.

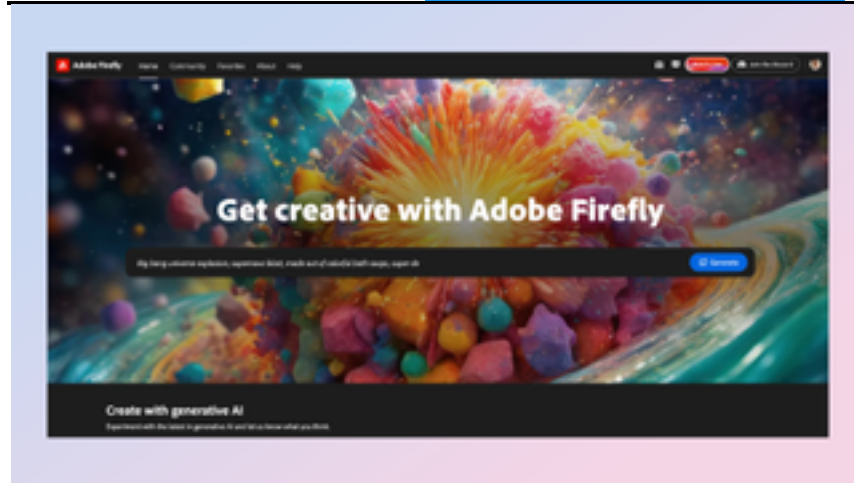
Everything is based on text prompts and it uses completely natural language for generation. For example you can tell it to generate a picture of a cat, then follow up with asking it to add a hat.

With the most recent update you can now click on the image itself and make edits within the picture. This is done by drawing over the part you want to change and telling ChatGPT how to change it — once again relying on the conversational nature of the editor.

I don't think DALL-E is the best AI image generator but it is a good all rounder. It can do text, it can create photorealistic images (with a slight uncanny valley) and produce artistic work but the ability to reason and rationalize over the image with text is the big benefit.

It is in urgent need of an AI image update. Technically GPT-4o can create images on its own without the need for a third model (the way Gemini uses Imagen 3 or ChatGPT currently uses DALL-E 3) but it isn't currently allowed to do so because of safety concerns.

BEST FOR ETHICAL TRAINING



(Image credit: Adobe Firefly)

7. Adobe Firefly

A more ethical approach with impressive artistic flare

SPECIFICATIONS

Free plan: Free with Adobe account

Paid plan: \$4.99/month

REASONS TO BUY

Prompt suggestions

Flexible controls

Licensed images in training data

REASONS TO AVOID

Struggles with photorealism

[Adobe Firefly](#) has some impressive tools including prompt suggestions, deep customizations for the image generation and a training dataset almost exclusively trained on Adobe Stock pictures.

This last point means it has a more ethical training set than most image generators on the market, even prompting Adobe to offer financial indemnity from copyright claims against images generated using Firefly.

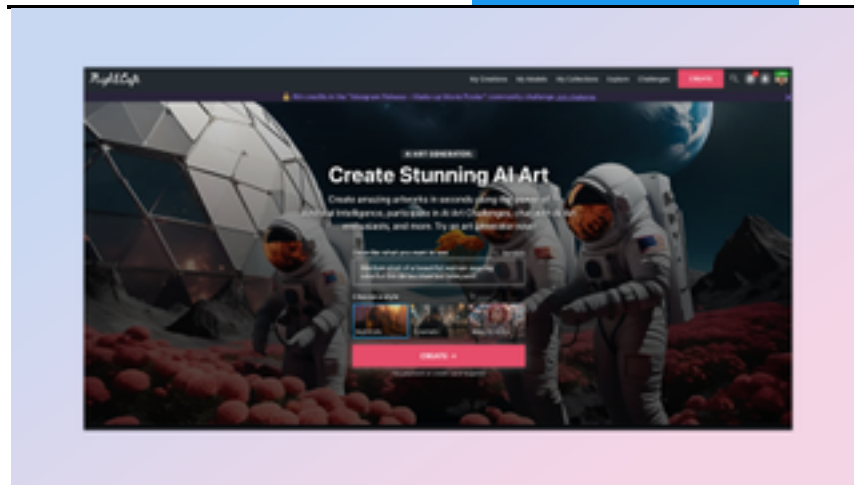
Firefly was given a major upgrade recently with the arrival of the Firefly 3 model. With this came impressive degrees of photorealism compared to the previous model and even better artistic skills — which were already among the best.

Adobe offers a number of generative AI features including vector generation, template creation and generative fill in Photoshop, all powered by the Firefly model.

One of the best features of Firefly is also one of the most recent. It is called Structural Reference and lets you carry the layout of one image over to another.

There is even a video model coming soon.

BEST FOR COMMUNITY



(Image credit: NightCafe)

8. NightCafe

A fantastic community with games, challenges and a range of models

SPECIFICATIONS

Free plan: 5 free daily credits plus credits for challenges and sharing images

Paid plan: Starts at \$6 a month for 100 credits per month in addition to the free credits

REASONS TO BUY

Beginners and experts

Community support

Range of model choices

REASONS TO AVOID

Initial learning curve

Busy UI

NightCafe was one of my first introductions to the world of generative AI. It was initially built on some custom models and later expanding to use Stable Diffusion-based models that have since been fine-tuned and customized.

As well as offering access to a wide variety of models, including Flux, Google's Imagen 3, Stable Diffusion, OpenAI's DALL-E 3 and now Ideogram, NightCafe makes it relatively easy to fine-tune or train your own models and lets easily manage and share your creations.

It even allows you to animate the individual photos using Stable Video Diffusion and upscale images using a range of upscaler models. You can even animate with Runway Gen-3 Turbo and Luma Labs Dream Machine within the single platform.

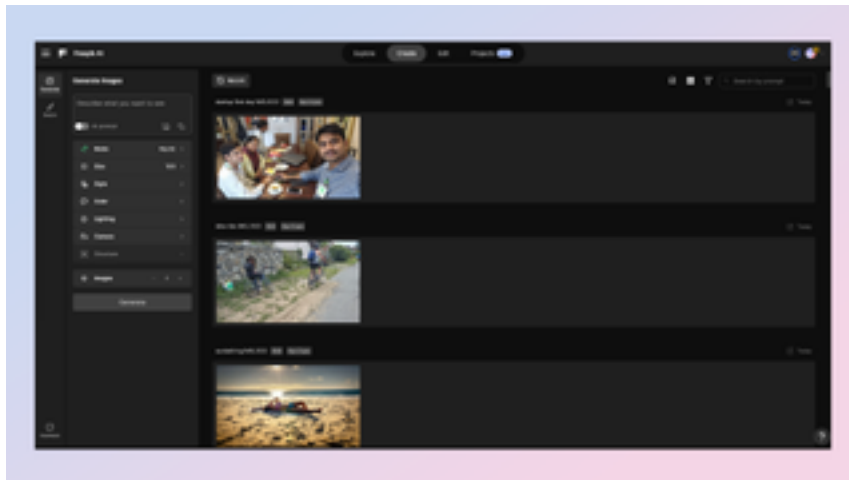
Its community is the real secret sauce behind why I keep going back to NightCafe again and again, regardless of which new models or services launch. There are regular competitions, chatrooms and creator jams for new ideas.

The company works to support that community, even engaging in partnerships with third party companies to offer prizes and free credits.

The credit system is also particularly helpful. Other services require a fixed monthly payment or have costly credit packs but NightCafe allows you to buy what you need.

It isn't perfect. The UI can be busy and the initial user journey can be daunting, but once you've engaged with it a while you'll find yourself coming back again and again.

BEST FOR INNOVATION



(Image credit: Freepik)

9. Freepik

Home of Mystic

SPECIFICATIONS

Free Plan: 20 images/day (restricted models)

Paid Plan: Starts at \$7/month (up to 7000 images a month)

Freepik is an interesting platform. Originally known for its library of images, it has very quickly become a major force in the artificial intelligence space. Some of this was driven by its acquisition of AI-upscaler Magnific earlier this year.

The AI generation part of the platform is incredibly easy to use and intuitive. It also includes a prompt enhancer that uses a large language model to improve your simplistic prompts.

You can add filters and elements to the image generation such as custom styles, colors and lighting. You can also specify the camera type if you are taking a photo. When exporting you have an easy upscale button that happens automatically before offering the download.

The feature set goes beyond simple image generation. One of these is sketch enhancement using AI. You can doodle on the screen and the AI will turn your doodle into an artwork or photograph.

Outside of generation, editing and sketching you can also upscale images from outside of Freepik, expand canvas in a similar way to Photoshop and use AI to retouch existing images.

Freepik has become a good all-round AI powerhouse, a way to play with image creation, including offering the latest models. It has all versions of the Flux family of models including Flux 1.1 Pro, the new rapid and realistic version.

The big stand out though is Mystic, built by the Magnific team. This is a fine-tuned and customized version of Flux that is perfect for photorealistic depictions of people.

BEST FOR ORIGINALITY



(Image credit: Luma Labs Dream Machine)

10. Luma Photon

A brilliant collaboration tool with a lot of control

SPECIFICATIONS

Free plan: Limited use, non-commercial

Paid plans: \$9.99, 3200 credits, non-commercial

REASONS TO BUY

Easy to collaborate with the AI

Convert image to video

Multiple image sizes

Chat-based prompting

REASONS TO AVOID

Non-commercial on cheap plan

Luma's Dream Machine is more traditionally an AI video platform, at least that is how it started. With its most recent update, and the addition of the Photon image model it has become one of the most innovative approaches to AI creation.

The Photon model can generate compelling images and is surprisingly creative and personalizable through complex prompting. It is also higher resolution for a lower price than other similar-sized models from Luma's competitors. Its big differentiator is its

prompt adherence and integration with the new look Dream Machine, which is now both image and video.


In Dream Machine a user enters a prompt and then collaborates with the platform to enhance both the prompt and adapt the image. [I tested this](#) in a review and found that the customization makes it very easy to get the exact image you are looking for.

One of my favorite features of Dream Machine is the interface. It is more like a chatbot than a traditional AI image generator. After entering your prompt you get the full prompt as well as four images. Keywords within the prompt are highlighted and can be quickly customized to change specific elements within the image.

You can then also opt to show more images, brainstorm on the side, and even reply to the response to that prompt to get specific changes. Clicking reply will load the generated images into the chat box as an attachment where you can send it to the video editor or change the image using as simple natural language prompt.

Luma Dream Machine works by creating “boards” for each topic. You can then iterate through the idea, brainstorming with the AI and keeping it all collected in one location. Image 1 of 5

BEST FOR VIDEO LINKUP



(Image credit: Kling KOLORS)

Kling's KOLORS is its dedicated AI image model. A lot of the big video platforms are launching custom image models as they use them to enhance their video platforms. Luma has Photon, Runway is launching Frames and Haiper already has an image version of its model.

KOLORS is impressive for visual consistency, human realism and portrait aesthetics but I've found it struggles with text rendering, something other models now largely do very well.

One area it is particularly good is in generating portraits. It captures the light and shadow required from a prompt particularly well and follows a prompt accurately.

KOLORS also comes with a virtual “try-on” where you can generate a virtual model and then either upload or generate an outfit for the model to wear. It comes with a series of example images you can use for inspiration and each image costs 1 credit.

The true value of KOLORS though comes from the deep integration with the Kling AI video model. The user generates an image and then clicks to animate it. Kling is one of the best AI video models currently available. It can create hyperrealistic video with accurate motion.

Final thoughts

Generating AI images is a fun process. You get to see an idea you've had in your head almost immediately and with new features like Canvas you can edit and make changes to that image to fit your exact concept.

I've used Ideogram to create posters to remind the children to turn off lights, artwork for app icons and even entire app designs. I've used Midjourney to create photos worthy of printing and Adobe Firefly regularly contributes to cover images for my stories.

I hope this guide has given you some insight into the best tools for creating AI art. If you've tried something I've not included send me an email, find me on X or leave a comment and I'll give it a review.