

We are proud to announce the first stage of the release of Stable Diffusion. [Click here to learn more.](#) ×



API Research News FAQ

Stable Diffusion Launch Announcement



Stability AI and our collaborators are proud to announce the first stage of release of Stable Diffusion to researchers via [this form](#), the model weights are hosted by our friends at Hugging Face once you get access. [The code is available here](#) and the model card [here](#). We are working together towards a public release soon.

This has been led by Patrick Esser from [Runway](#) and Robin Rombach from the [Machine Vision & Learning](#) research group at LMU Munich (formerly [CompVis](#) lab at Heidelberg University) building on their prior work on [Latent Diffusion Models](#) at CVPR'22, combined with support from communities at [Eleuther AI](#), [LAION](#) and our own generative AI team.

Stable Diffusion is a text-to-image model that will empower billions of people to create stunning art within seconds. It is a breakthrough in speed and quality meaning that it can run on consumer GPUs. You can see some of the amazing output that has been created by this model without pre or post-processing on this page.

The model itself builds upon the work of the team at CompVis and Runway in their widely used latent diffusion model combined with insights from the conditional diffusion models by our lead generative AI developer Katherine Crowson, Dall-E 2 by Open AI, Imagen by Google Brain and many others. We are delighted that AI media generation is a cooperative field and hope it can continue this way to bring the gift of creativity to all.



The core dataset was trained on LAION-Aesthetics, a soon to be released subset of LAION 5B. LAION-Aesthetics was created with a new CLIP-based model that filtered LAION-5B based on how “beautiful” an image was, building on ratings from the alpha testers of Stable Diffusion. LAION-Aesthetics will be released with other subsets in the coming days on <https://laion.ai>.

Stable Diffusion runs on under 10 GB of VRAM on consumer GPUs, generating images at 512x512 pixels in a few seconds. This will allow both researchers and soon the public to run this under a range of conditions, democratizing image generation. We look forward to the open ecosystem that will emerge around this and further models to truly explore the boundaries of latent space.

The model was trained on our 4,000 A100 Ezra-1 AI ultracluster over the last month as the first of a series of models exploring this and other approaches.

We have been testing the model at scale with over 10,000 beta testers that are creating 1.7 million images a day.



This output has given us numerous insights as we prepare for a public release soon. This will provide the template for the release of many open models we are currently training to unlock human potential. We will also be releasing open synthetic datasets based on this output for further research.

We aim to set new standards of collaboration and reproducibility for the models that we create and support and will share our learnings in the coming weeks.

We hope to progressively increase the number of collaborators for our benchmark models. If you would like to help, please join one of the communities we support and/or reach out to info@stability.ai

Some comments by various folks:

"EleutherAI has spent the past two years advancing open source large-scale AI research. We are thrilled to be working with and supporting like-minded researchers to enable scientific access to these emerging technologies" - Stella Biderman, Lead Researcher at EleutherAI

"With this project we continue to pursue our mission to make state of the art machine learning accessible for people from all over the world. 100% open. 100% free." - Christoph,

Organizational Lead & researcher at LAION e.V.

"We are excited to see what will be built with the current models as well as to see what further works will be coming out of open, collaborative research efforts!" - Patrick (Runway) and Robin (LMU)

"We're excited that state of the art text-to-image models are being built openly and we are happy to collaborate with CompVis and Stability.ai towards safely and ethically release the models to the public and help democratize ML capabilities with the whole community" - Apolinário, ML Art Engineer, Hugging Face

"We are delighted to release the first in a series of benchmark open source Stable Diffusion models that will enable billions to be more creative, happy and communicative. This model builds on the work of many excellent researchers and we look forward to the positive effect of this and similar models on society and science in the coming years as they are used by billions worldwide".
- Emad, CEO, Stability AI

p.s. "GPUs go brrr." - Robin

Stable Diffusion Public Release

[API](#)

[Research](#)

[Blog](#)

[Press](#)

[FAQ](#)

[HPC Center](#)

[Privacy Policy](#)

[Terms of Use](#)

[Simulacrabot Terms of Service](#)

[Stable Diffusion Terms of Service](#)

Contact Us

info@stability.ai

press@stability.ai

careers@stability.ai

© STABILITY AI LTD, 2022