

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



stability.ai

[API](#) [Research](#)

[Blog](#) [FAQ](#)

[in](#)



Stable Diffusion Public Release



It is our pleasure to announce the public release of stable diffusion following our release for researchers [<https://stability.ai/blog/stable-diffusion-announcement>]

Over the last few weeks we all have been overwhelmed by the response and have been working hard to ensure a safe and ethical release, incorporating data from our beta model tests and community for the developers to act on.

In cooperation with the tireless legal, ethics and technology teams at HuggingFace and amazing engineers at CoreWeave, we have incorporated the following elements:

i) The model is being released under a Creative ML OpenRAIL-M license [<https://huggingface.co/spaces/CompVis/stable-diffusion-license>]. This is a permissive license that allows for commercial and non-commercial usage. This license is focused on ethical and legal use of the model as your responsibility and must accompany any distribution of the model. It must also be made available to end users of the model in any service on it.

ii) We have developed an AI-based Safety Classifier included by default in the overall software package. This understands concepts and other factors in generations to remove outputs that may not be desired by the model user. The parameters of this can be readily adjusted and we welcome input from the community how to improve this. Image generation models are powerful, but still need to improve to understand how to represent what we want better.



As these models were trained on image-text pairs from a broad internet scrape, the model may reproduce some societal biases and produce unsafe content, so open mitigation strategies as well as an open discussion about those biases can bring everyone to this conversation. Learn more about the model strengths and limitations in the model card.

This release is the culmination of many hours of collective effort to create a single file that compresses the visual information of humanity into a few gigabytes.



We hope everyone will use this in an ethical, moral and legal manner and contribute both to the community and discourse around it. Please carefully read the model card for a full outline of the limitations of this model and we welcome your feedback in making this technology better.

You can join our dedicated community for Stable Diffusion here: [\[https://discord.gg/stablediffusion\]](https://discord.gg/stablediffusion) where we have areas for developers, creatives and just anyone inspired by this.

You can find the weights, model card and code here: [\[https://huggingface.co/CompVis/stable-diffusion\]](https://huggingface.co/CompVis/stable-diffusion)

An optimized development notebook using the HuggingFace diffusers library: [\https://colab.research.google.com/github/huggingface/notebooks/blob/main/

[diffusers/stable_diffusion.ipynb](#)]

A public demonstration space can be found here:

[\[https://huggingface.co/spaces/stabilityai/stable-diffusion\]](https://huggingface.co/spaces/stabilityai/stable-diffusion)


For more control and rapid generation you can try our DreamStudio beta here:

<http://beta.dreamstudio.ai>.

Additional functionality and API access will be activated shortly, including local GPU support, animation, logic-based multi-stage workflows and many more.

We are also happy to support many partners through our API and other programs and will be posting on these soon.

The recommended model weights are v1.4 470k, a few extra training steps from the v1.3 440k model made available to researchers. The final memory usage on release of the model should be 6.9 Gb of VRAM.





In the coming period we will release optimized versions of this model along with other variants and architectures with improved performance and quality. We will also release optimisations to allow this to work on AMD, Macbook M1/M2 and other chipsets. Currently NVIDIA chips are recommended.

We will also release additional tools to help maximize the impact and reduce potential adverse outcomes from these tools with amazing partners to be announced in the coming weeks.

This technology has tremendous potential to transform the way we communicate and we look forward to building a happier, more communicative and creative future with you all.

Please contact info@stability.ai with any queries, much more soon.

Stable Diffusion Launch Announcement