Ethical AI art generation? Adobe Firefly may be the answer

Enlarge / An Adobe Firefly AI image generator example.

reader comments

51 with

Share this story

On Tuesday, Adobe unveiled Firefly, its new AI image synthesis generator. Unlike other AI art models such as Stable Diffusion and DALL-E, Adobe says its Firefly engine, which can generate new images from text descriptions, has been trained solely on legal and ethical sources, making its output clear for use by commercial artists. It will be integrated directly into Creative Cloud, but for now, it is only available as a beta.

Since the mainstream debut of image synthesis models last year, the field has been fraught with issues around ethics and copyright. For example, the AI art generator called Stable Diffusion gained its ability to generate images from text descriptions after researchers trained an AI model to analyze hundreds of millions of images scraped from the Internet. Many (probably most) of those images were copyrighted and obtained without the consent of their rights holders, which led to lawsuits and protests from artists.

To avoid those legal and ethical issues, Adobe created an AI art generator trained solely on Adobe Stock images, openly licensed content, and public domain content, ensuring the

generated content is safe for commercial use. Adobe goes into more detail in its news release:

Advertisement

> *Adobe Firefly will be made up of multiple models, tailored to serve customers with a wide array of skillsets and technical backgrounds, working across a variety of different use cases. Adobe's first model, trained on Adobe Stock images, openly licensed content and public domain content where copyright has expired, will focus on images and text effects and is designed to generate content safe for commercial use. Adobe Stock's hundreds of millions of professional-grade, licensed images are among the highest quality in the market and help ensure Adobe Firefly won't generate content based on other people's or brands' IP. Future Adobe Firefly models will leverage a variety of assets, technology and training data from Adobe and others. As other models are implemented, Adobe will continue to prioritize countering potential harmful bias.*

In addition to its commitment to a more ethical form of AI generator, Adobe is doubling down on ethics with a "Do Not Train" tag for creators who do not want their content used in model training. According to Adobe, this tag will "remain associated with content wherever it is used, published, or stored."

In terms of performance, Adobe Firefly looks similar to OpenAI's DALL-E, although Ars has not been able to evaluate its performance for itself yet. Demos on the Firefly website show various capabilities, including "text to image" (creating unique images from a text description), "text effects" (applying styles or textures with a text prompt), and "recolor vectors" (creating unique variations of a work from a text description).

Whether artists will adopt Firefly into their workflows

remains to be seen, but the new AI model seems, on its face, like a positive step forward for AI ethics advocates. As always, we'll need to take Adobe's claims with a grain of salt, and we'll keep you updated as new details emerge.

https://arstechnica.com/?p=1925844