

MARKET RESEARCH REPORT



In the realm of generative AI, OpenAI's DALL-E 2, MidJourney, and Stable Diffusion each bring a unique perspective to the canvas. DALL-E 2 excels in precision, MidJourney pioneers new horizons of creativity, and Stable Diffusion offers unwavering reliability. These three pillars collectively redefine the future of AI, where possibilities are boundless and the extraordinary is now within reach.

Table of Contents

Table of Contents	1
Executive Summary	2
OpenAI's DALL-E 2	3
Midjourney	5
STABLE DIFFUSION XL by Stability.ai	7
Conclusion and Recommendations	9
Thank you!	10

Executive Summary

In the ever-evolving landscape of generative AI, we embark on a journey that unfurls the remarkable distinctions between OpenAI's triumvirate of DALL-E 2, MidJourney, and Stable Diffusion.

Generative AI is revolutionizing commerce by enhancing creativity, automating tasks, and fostering innovation. It's a pivotal force in shaping the future of businesses as they harness its capabilities to streamline processes and unlock creative potential.

Overview

In the dynamic landscape of generative AI images, OpenAI's DALL-E 2, Stable Diffusion, and MidJourney APIs emerge as revolutionary tools for commercial use. DALL-E 2 masterfully balances precision and fine art, making it an indispensable asset for businesses seeking to elevate their visual content. The unwavering reliability of Stable Diffusion ensures consistent quality, while MidJourney paves the way for limitless creativity, redefining commercial artistry. Together, these APIs redefine the boundaries of imagination and innovation, providing businesses with a dynamic and unprecedented palette for unparalleled visual storytelling and artistic vision.

In the realm of generative AI, OpenAI's DALL-E 2, MidJourney, and Stable Diffusion each bring a unique perspective to the canvas. DALL-E 2 excels in precision, MidJourney pioneers new horizons of creativity, and Stable Diffusion offers unwavering reliability. These three pillars collectively redefine the future of AI, where possibilities are boundless and the extraordinary is now within reach.

Scope and limitations

Generative AI APIs offer a broad scope for enhancing creativity and content generation across industries, streamlining processes, and promoting innovation. They excel at automating tasks and maintaining quality but should be used as creative aids, considering the need for human refinement. Ethical and privacy concerns also warrant careful attention when integrating these APIs. Balancing their potential with an understanding of their limitations is essential for their responsible use.

OpenAI's DALL-E 2

DALL-E 2, an evolution of the original, generates detailed, high-resolution images from text descriptions. Trained on vast datasets, it understands the visual representation of various concepts, resulting in both realistic and creative images.

Build DALL-E directly into your apps to generate and edit novel images and art. Our image models offer three tiers of resolution for flexibility.

Generation & Data collection methods

DALL-E 2 can create original, realistic images and art from a text description. It can combine concepts, attributes, and styles.

- [Image generation](#) - DALL-E 2 can create original, realistic images and art from a text description. It can combine concepts, attributes, and styles.
- [Outpainting](#) - DALL-E 2 can expand images beyond what's in the original canvas, creating expansive new compositions.
- [Inpainting](#) - DALL-E 2 can make realistic edits to existing images from a natural language caption. It can add and remove elements while taking shadows, reflections, and textures into account.
- [Variations](#) - DALL-E 2 can take an image and create different variations of it inspired by the original.

#	Resolution	Price
1.	1024×1024	\$0.020 / image
2.	512×512	\$0.018 / image
3.	256×256	\$0.016 / image

71.7%
preferred for caption matching

88.8%
preferred for photorealism

What are the rate limits on OpenAI's DALL-E 2 API?

We can view the rate limits for our organization under the rate limits section of the account management page. Rate limits are automatically adjusted based on the history of good use.

They enforce rate limits at the organization level, not the user level, based on the specific endpoint used as well as the type of account you have. Rate limits are measured in three ways: RPM (requests per minute), RPD (requests per day), and TPM (tokens per minute). Rate limit can be hit by any of the three options depending on what occurs first. For example, you might send 20 requests with only 100 tokens to the Completions endpoint and that would fill your limit (if your RPM was 20), even if you did not send 150k tokens (if your TPM limit was 150k) within those 20 requests.

Usage tiers :

Your rate limits and spending limit (quota) are automatically increased as you establish a track record of successful payments and API usage. The qualifications and limits associated with each usage tier are shown below:

TIER	QUALIFICATION	MAX CREDITS	REQUEST LIMITS	TOKEN LIMITS
Free	User must be in an allowed geography	\$100	3 RPM 200 RPD	20K TPM (GPT-3.5)
Tier 1	\$5 paid	\$100	500 RPM 10K RPD	40K TPM (GPT-3.5) 10K TPM (GPT-4)
Tier 2	\$50 paid and 7+ days since first successful payment	\$250	5000 RPM	80K TPM (GPT-3.5) 20K TPM (GPT-4)
Tier 3	\$100 paid and 7+ days since first successful payment	\$500	5000 RPM	160K TPM (GPT-3.5) 40K TPM (GPT-4)
Tier 4	\$250 paid and 14+ days since first successful payment	\$1000	10K RPM	1M TPM (GPT-3.5) 150K TPM (GPT-4)
Tier 5	\$1000 paid and 30+ days since first successful payment	\$1000	10K RPM	1M TPM (GPT-3.5) 150K TPM (GPT-4)

✦ How do rate limits work?

If your rate limit is 60 requests per minute and 150k tokens per minute, you'll be constrained by reaching the requests/min cap or depleting tokens first. For instance, with a max of 60 requests/min, you can send one request per second. Beyond this, just a 200ms pause is required after hitting the rate limit to send another request. Defaulting to 3,000 requests/min allows sending requests every 20ms or every 0.02 seconds.

Midjourney

Midjourney AI is a powerful new tool for generating creative images, using a combination of artificial intelligence and human input. It is still in beta testing, but has already been used to create stunning and original artwork, from digital paintings to concept art to architectural renderings.

Midjourney's API opens up even more possibilities for commercial usage, allowing businesses to integrate Midjourney's image generation capabilities into their own products and services. For example, a fashion brand could use Midjourney to generate new designs for clothing or accessories, or a marketing agency could use it to create visually appealing and engaging ads.

#	Period	Basic Plan	Standard Plan	Pro Plan	Mega Plan
1.	Monthly Subscription Cost	\$10	\$30	\$60	\$120
2.	Annual Subscription Cost	\$96 (\$8 / month)	\$288 (\$24 / month)	\$576 (\$48 / month)	\$1152 (\$96 / month)
3.	Fast GPU Time	3.3 hr/month	15 hr/month	30 hr/month	60 hr/month
4.	Relax GPU Time	-	Unlimited	Unlimited	Unlimited
5.	Purchase Extra GPU Time	\$4/hr	\$4/hr	\$4/hr	\$4/hr
6.	Work Solo In Your Direct Messages	Yes	Yes	Yes	Yes
7.	Stealth Mode	-	-	Yes	Yes
8.	Maximum Concurrent Jobs	3 Jobs 10 Jobs waiting in queue	3 Jobs 10 Jobs waiting in queue	12 Fast Jobs 3 Relaxed Jobs 10 Jobs in queue	12 Fast Jobs 3 Relaxed Jobs 10 Jobs in queue
9.	Rate Images to Earn Free GPU Time	Yes	Yes	Yes	Yes
10.	<u>Usage Rights</u>	General Commercial Terms*	General Commercial Terms*	General Commercial Terms*	General Commercial Terms*

*[*Click here to view Usage Terms & Conditions](#)*

Our Rights

Subject to the above license, You own all Assets You create with the Services, provided they were created in accordance with this Agreement. This excludes upscaling the images of others, which images remain owned by the original Asset creators. Midjourney makes no representations or warranties with respect to the current law that might apply to You. Please consult Your own lawyer if You want more information about the state of current law in Your jurisdiction. Your ownership of the Assets you created persists even if in subsequent months You downgrade or cancel Your membership. However, You do not own the Assets if You fall under the exceptions below.

- If You are an employee or owner of a company with more than \$1,000,000 USD a year in gross revenue and You are using the Services on behalf of Your employer, You must purchase a "Pro" or "Mega" membership for every individual accessing the Services on Your behalf in order to own Assets You create. If You are not sure whether Your use qualifies as on behalf of Your employer, please assume it does.
- If You are not a Paid Member, You don't own the Assets You create. Instead, Midjourney grants You a license to the Assets under the Creative Commons Noncommercial 4.0 Attribution International License (the "Asset License").

The full text is accessible as of the Effective Date here: [Link](#)

- Please note: Midjourney is an open community which allows others to use and remix Your images and prompts whenever they are posted in a public setting. By default, Your images are publically viewable and remixable. As described above, You grant Midjourney a license to allow this. If You purchase a "Pro" or "Mega" plan, You may bypass some of these public sharing defaults.
- If You purchased the Stealth feature as part of Your "Pro" or "Mega" subscription or through the previously available add-on, we agree to make best efforts not to publish any Assets You make in any situation where you have engaged stealth mode in the Services.
- Please be aware that any image You make in a shared or open space such as a Discord chatroom, is viewable by anyone in that chatroom, regardless of whether Stealth mode is engaged.

*[*Click here to view A Complete Usage Terms & Conditions](#)*

STABLE DIFFUSION XL

by **Stability.ai**

With Stable Diffusion XL, you can create descriptive images with shorter prompts and generate words within images. The model is a significant advancement in image generation capabilities, offering enhanced image composition and face generation that results in stunning visuals and realistic aesthetics.

Stable Diffusion XL is a state-of-the-art text-to-image diffusion model that can generate high-quality images from text descriptions. It is one of the most powerful and versatile image generation models available, and its API has the potential to revolutionize the way businesses create and use images.

Here are some specific examples of how Stable Diffusion XL and its API can be used for commercial purposes:

- **Product design and development:** Stable Diffusion XL can be used to generate new product ideas, create early prototypes, and visualize final designs. For example, a fashion company could use Stable Diffusion XL to generate concepts for new clothing or accessories, or a furniture company could use it to design new chairs or tables.
- **Marketing and advertising:** Stable Diffusion XL can be used to create visually appealing and engaging marketing materials, such as ads, social media posts, and website banners. For example, a food company could use Stable Diffusion XL to generate images of delicious-looking dishes, or a travel agency could use it to generate images of beautiful landscapes and attractions.
- **Entertainment and gaming:** Stable Diffusion XL can be used to create concept art, character designs, and backgrounds for video games, movies, and TV shows. For example, a game developer could use Stable Diffusion XL to create new worlds and characters for their game, or a film studio could use it to design the sets and costumes for their next movie.
- **Education and research:** Stable Diffusion XL can be used to create educational materials, such as diagrams, illustrations, and simulations. For example, a school could use Stable Diffusion XL to create visuals for science textbooks, or a research lab could use it to generate images of complex scientific data.

Stable Diffusion XL simplifies various image tasks like editing, classification, and segmentation. It enhances workflow quality, saving time and resources. This versatile tool has the potential to transform image use in business, improving products, expanding customer reach, and boosting profits. It's open-source, providing flexibility and room for innovative applications.

How does pricing work?

Most API usage requires credits, where the total price is in part based on the amount of compute required to fulfill the request.

Buying Credits:

After depleting your initial 25 free credits, additional credits can be purchased via the [account page](#). Credits are priced at \$10 per 1,000 credits, which is enough credits for roughly 5,000 SDXL 1.0 images.

Stable Diffusion XL 1.0 -

Building upon the success of the Stable Diffusion XL v0.9 release, Stable Diffusion XL v1.0 brings with it a number of image detail improvements. Specifically, Stable Diffusion XL v1.0 offers enhanced vibrancy and overall color tone accuracy, including deeper black and brighter white tones.

★ What is a credit?

A credit is the unit of currency used to pay for AI image generation and editing. Credit usage depends on to the compute required by the tool used.

Find out more on the number of credits needed per call or per image generation on the table above.

*Credit Cost - 50 Steps (0.25, \$0.003) [All purchased credits are valid for 2 years, starting from the date of their purchase.]

#	Tool	Credit cost
1.	Cleanup	1 credit / call
2.	Remove text and logo	1 credit / call
3.	Depth estimation	1 credit / call
4.	Normals estimation	1 credit / call
5.	Remove background	1 credit / call
6.	Replace background	1.5 credit / call
7.	Text to image	1.6 credit / call
8.	Upscaler	From 0.2 credit / output megapixel

#	Credits	Credit cost
1.	5000 Credits	\$144
2.	50000 Credits	\$961
3.	500 000 credits	\$7089

Conclusion and Recommendations

Ultimately, the choice of API service depends on the specific requirements and priorities of the user. For those seeking photorealism and adherence to detailed instructions, DALL-E 2 emerges as the frontrunner. Midjourney caters to those seeking artistic flair and the ability to visualize abstract ideas. Stable Diffusion XL, with its flexibility and openness, caters to users who value customization and control over the image generation process.

DALL-E 2 is the **Best Overall** text-to-image AI model, with the best balance of image quality, versatility, and features. However, it is currently in beta testing and has limited access.

Stable Diffusion XL is a **Good Alternative** to DALL-E 2, especially if you need high-quality images with a wide range of details or if you need an open-source solution.

When choosing an API service from DALL-E 2, Midjourney, and Stable Diffusion XL, you should consider the following factors:

#	Factors	Priority
1	Image style: What kind of image style are you looking for? If you need realistic and creative images, then DALL-E 2 is a good choice. If you need artistic and dreamlike images, then Midjourney is a good choice. If you need high-quality images with a wide range of details, then Stable Diffusion XL is a good choice.	HIGH
2	Image consistency: How important is it for the generated images to be consistent with the text description? If image consistency is very important, then Stable Diffusion XL is a good choice.	MEDIUM
3	API features: What kind of API features do you need? For example, do you need an API that allows you to generate images with different resolutions or that allows you to specify the image style?	LOW
4	Pricing: How much are you willing to pay for the API service? DALL-E 2 and Midjourney are currently in beta testing and have limited access. Stable Diffusion XL is open-source and can be used freely.	HIGH

#	Model	Image quality	Versatility	Features	Access	Pricing
1.	DALL-E 2	Best	Good	Good	Beta testing	Limited access
2.	Stable Diffusion XL	Good	Very good	Good	Open source	Free
3.	Midjourney	Good	Limited	Limited	Beta testing	Limited access



Thank you!

For taking the time to explore this report. Your attention and consideration are greatly appreciated. As we wrap up this document, I extend my heartfelt gratitude for your interest and involvement. Your support and engagement make our endeavors truly meaningful. If you have any questions or further insights to share, please feel free to reach out. Once again, thank you for being a part of this journey.

SAVANNAH-LABS