REPORT

# The best Al tools for 2D game art

ROOM8 GROUP



### Contents

SETTING THE STAGE

Contributing team

Setting the stage

Al in game art generation

Room 8 Group's approach to Al

The AI art landscape

Methodology

GENERATIVE AI 2D ART IN LATE 2024

Comparing the output of four AI art models

What about Midjourney?

The best tools on the market: Stable Diffusion XL (SDXL)

The best tools on the market: FLUX.1 Dev CONCLUSIONS

How you can use our results

Putting the results into action

How Room 8 Group is using Flux and SDXL to save time and money on real projects

Next steps

About Room 8 Group



## Contributing team









CONTRIBUTOR



Svitlana Sokolovska

ART PRODUCER

### Oleksii Tyshchenko

2D TECHNICAL DIRECTOR

### Billy Studholme

SENIOR COPYWRITER











Andrey Rudnichenko

COMPETENCE LEAD

Tetyana Kaminer

LEGAL COUNSEL

Oleksii

LEGAL COUNSEL

■ SETTING THE STAGE

### Setting the stage

Most companies in game development are actively conducting research to find the best ways of integrating AI solutions into their production pipeline. One area in which it is actively put to use today is generative AI for 2D and concept art.

That is the focus of this report. We at Room 8 Group have conducted a review of dozens of genAl tools for 2D and concept art in the context of game development.

GenAl can complement human skill and creativity in making outstanding games, and it is a fast-moving space. Our commitment to delivering copyright-compliant game assets of the highest quality for our partners drives us to continually explore new tools and methods. We approach the tools reviewed in this report as supplementary elements within a broader, comprehensive production pipeline. Their role is carefully defined to support specific stages of the workflow, while the human touch remains central to creativity, final polish, and overall quality. This initiative also reflects our dedication to knowledge-sharing and professional growth, empowering our specialists to enhance their skills and adapt to new technologies.

As with any new technology, it is vital to understand both the potential and limitations of genAl. We prioritize transparency and compliance, enabling us to offer our partners informed and legally sound solutions for integrating these tools into their workflows, whether online or offline.

### This report will:

- 1. Demonstrate where is best and worst—to use genAl in the 2D/concept art production pipeline
- 2. The best genAl tools for 2D/ concept art generation in the context of game development
- 3. How Room 8 Group is already using these findings in real projects, and how we have saved—and helped our partners save—significantly on time and cost



















### Al in 2D game art generation

Best stages for AI integration into art pipeline

					*		
Reference search	Moodboard creation/ ideas/rough sketches/ art direction	Basic composition sketch	Creation of line art/ shape sketch	Coloring of BW (better and quicker by hand)	2D/concept rendering	Polishing/cu animation/e	
<ul> <li>Nicely applicable</li> <li>Somewhat applicable</li> <li>Can be applicable with some conditions</li> <li>Hardly applicable</li> </ul>		$\uparrow$	$\uparrow$	Ŷ		Polishing is availa	
		Possible to integrate into pipeline as generation of additional sketch; additionally we can create a model that generates BW compositions for artists		Low quality of AI to date		complex FLUX 1; partially with img: or with ControlNe	

During our review we clarified the areas of the 2D art pipeline for which AI is clearly useful, and where it is less so.

Al has the potential to make just about every aspect of game art and development more efficient. Given current limitations, however, it is presently most useful while drafting and rendering 2D and concept art.

### Draft concepting

Al art tools significantly reduce the amount of time needed to create drafts of concept art—a key stage of the art production process. Using AI, artists can easily iterate on their ideas and fine-tune them to a specific format. It's important to note that concepting is a highly intellectual task, meaning AI can't do it all that effectively at present. But it can be used to decent effect to draft concept art and generate ideas.

#### 2D/concept rendering

Rendering can be time-consuming work, but AI is perfect for this kind of task. In addition to speeding up the final rendering of art, AI can re-render games to help develop fun variants or alternative game modes. You may have seen wacky Al-generated videos of a photo-realistic Super Mario, or of popular games converted into a cartoonish or noir-infused style. With the power of AI, one game environment can be quickly and seamlessly rendered into another.









 $\times$ 

Right now, if you want both speed and quality, not every part of the 2D and concept art pipeline should involve Al. For example, in the latter stages of development most models are simply incapable of helping devs adequately polish their images, although more advanced models such as SDXL and FLUX.1 Dev are beginning to offer this functionality. For now, even advanced models are still marked by occasional small errors—extra fingers, unrealistic physical positions, and so on—and thus require a human touch during the final

#### How to use AI

Broadly speaking, you need a strong reference base to use genAl art tools effectively. This will allow you to create a custom style with the right visual identity, objects, characters, and so on, so it's important to have references that highlight elements of the desired result. Five references should be a minimum; with 10 you could achieve a specific style; 15 or more is optimal.

With strong references and precise text prompts, you can test different ideas and perhaps even start to piece together the art direction of your game. GenAl is also great for rendering at the latter stages of 2D and concept art generation.







## Room 8 Group's approach to Al

Before we get into the results of our review, it's important to explain how we at Room 8 Group think about and use AI.

It's very possible that today's leading AI art models will become an afterthought in the near future. However, we are teaching our specialists to use the technology in a way that will be applicable in the future, regardless of the specific tools they use.

We are actively applying the findings of this report to set up AI art pipelines that we have already put to use for partners on real game projects, though only with their explicit approval (see page #25). These pipelines are designed to be useful and relevant for years to come, even if the specific AI tools used are subject to change.

Our focus is on best practices for genAl as a medium—not on the specific strengths or weaknesses of any individual model.



Our approach to AI is fundamentally rooted in the open-source models, which act as the base for the overwhelming majority of AI models on the market right now, even the proprietary ones. If any open-source tool becomes proprietary or ceases to exist, we are prepared to use alternatives for just about every function. And if our partners want to use the proprietary tools, we're well-versed in those, too.

### We use genAl models as we would any other tool: they form part of a wider art production effort for 2D and concepting.

Much like traditional artists have never been beholden to a specific variety of oil paint or a particular brand of pencils, our specialists are taught to be proficient with genAl tools as one particularly effective medium. Such tools are not the core driver of their creativity.



## The 2D Al art landscape in early 2025

At the moment, there are three types of generative AI models.

### 1. Open-source solutions

Includes Stable Diffusion XL, Stable Diffusion XL Turbo, and Stable Diffusion 3. These models are effective and can be used locally, but there's no guarantee that they were trained on legally licensed data.

### 2. Solutions based on Stable Diffusion

Such as Vizcom, Leonardo, and Krea.Al. These models can arguably provide more customization than the open-source Stable Diffusion solutions, but once again provide no guarantees that they were trained on licensed data.

### 3. Proprietary solutions

Includes OpenAl's Dall-E 3, BRIA, and Amazon Titan. These solutions claim to be trained on licensed data, meaning their images don't contain copyrighted content. However, it's not possible to ensure this data is 100% safe because the datasets used to train many proprietary Al solutions are not made public, so one cannot openly download them. Lastly, these solutions can't be used locally and struggle to customize their output for specific goals.





We have conducted, and interpreted the results of, our review from the perspective of a leading external game developer; we at Room 8 Group are a partner to the world's biggest game-makers.

Therefore, the results reflect how suitable we think different AI art tools are for corporate use by game developers and publishers.

The aim of our review was threefold:

- 1. To explore the latest genAl art solutions
- 2. To find out which AI art tools are most appropriate for corporate use by game developers and publishers
- 3. To explore how we can use AI solutions to reduce costs and save time in the 2D art pipeline

**PROJECT SUMMARY** 

collecting and reviewing data

chosen for consideration, which we are already using in production

### Nethodology



### 2 superior models





Some tools, such as Midjourney v5 and v6, are listed together. This is because while they are technically different versions, they produced very similar results during review. Plus, because someone would almost certainly not use v5 over v6 if the results are the same, it made little sense to list them separately here, because the purpose of this review is to show game developers and publishers which tools are best for them to use.

Members of our team across legal, technology and engineering, and art production spent six months reviewing 34 AI tools for art generation.



E

S

Ο

Adobe Firefly 2

**Amazon Titan 1** 

BRIA V.1

Dall-E 3

**FLUX.1 Dev** 

Bing Image Creator

FAL Auraflow 0.2, 0.3

FLUX.1 Pro, FLUX.1.1 Pro

AlphaVLLM Lumina Next

- FLUX.1 Schnell



Ideogram.AI



- iStock (Getty Images)
- - Kandinsky 2, 3
- **KREA**



Kwai Kolors



Leonardo.Al



Midjourney

- Pixart Alpha

- Pixart Sigma
- PlaygroundAl 1, 2
- Segmind SSD-1B, Tiny 众
- ٢, Shutterstock
  - **Stable Cascade Medium, Lite**
- **Stable Diffusion 1.4, 1.5, 1.6**
- Stable Diffusion 3 Large
- Stable Diffusion 3 Medium  $(\cdot)$
- Stable Diffusion 3.5 Medium
- **Stable Diffusion XL**
- Tencent HunyuanDit 1.2
  - THUDM CogView 3
- **Topaz Al**



Vizcom

Warp AI Wuerstchen 1.2

■ SETTING THE STAGE, METHODOLOGY





The first review stage involved preliminary testing of each model's legal context and basic art capabilities.

16 of the 34 tools we reviewed passed this preliminary check.

Those 16 were then reviewed in more depth, according to four criteria: quality, control, legal, and cost.

Of those 16 tools, there were four standouts, which we then put through a more rigorous round of review, mostly in terms of their quality and control.

Of those four, there were two superior tools that we recommend—which we are already using in production.



### 1. Quality

Quality refers to the overall standard of the model's output according to Room 8 Group's artists; in other words, whether the resulting art is of a high-enough quality to be used in real game projects.

### 2. Control

Control refers to how precisely the AI tool generates artwork according to a prompt—in other words, its controllability for the inputting party. Also, control includes how capably the model adheres to visual references like sketches, drafts, line art, and so on. Control is an important factor for us as a game developer to measure because professional work requires a higher level of control than most AI solutions provide.

### 3. Cost

is not included.

### 4. Legal

Legal considerations such as whether the model: guarantees no copyright infringement; involves considerations such as whether the tool guarantees the inputter's ownership of outputs (and for how long); can be installed locally (thus making it appropriate for use when an NDA has been signed); states explicitly that output art will not be used for the model's training purposes (such as machine-learning); and so on.

This concerns only the cost of the AI solutions themselves. The price of hardware needed to run AI models



### Comparing the output of four Alart models Once we had whittled it down to the four best options, we conducted fur testing by giving each the same prompts. Take a look at this example: Once we had whittled it down to the four best options, we conducted further testing by giving each the same prompts. Take a look at this example:

Prompt

Crying cartoon orange cat, sad animal, poor cat, high detailed, tears from cats eyes, poor stray cat on the street, half closed eyes, crying cartoon cat at a bar with beer

A girl in Pixar style, running with dog in Pixar city, detailed 3d cartoon

Tesla Cybertruck with Japanese stickers









### What SDXL produced is very close to the required level of detail for professional use.

It generates images faster than most, meaning you can make many more variants than other tools in the same amount of time, thus iterating quicker and more effectively. It's flexible in terms of style: it can easily be trained to create images in a specific style or genre—and in the shortest time—and it has a vast number of integrated styles and the highest number of communitytrained custom models. Because it is perhaps the most popular AI art tool in the world, it offers the widest set of professional instruments.

MORE DETAIL ON SDXL IN THE NEXT SECTION.





### FLUX.1 Dev currently yields the best results in terms of image quality, resolution, and the ability to follow text prompts.

In the above example, it was the only tool whose cat actually cried. Its main flaw is that it is slower than models like SDXL. It also requires more resources to train, though users are able to train it better than other AI tools (meaning better quality and control). It is suitable for more complex projects and requires less additional overpaint to tidy images. Essentially, FLUX may take longer than other models to reach a desired outcome, but with enough work, it can produce the best results.

MORE DETAIL ON SDXL IN THE NEXT SECTION.





OpenAl's art tool, Dall-E 3, is powered by the best large language model (LLM) in the world, GPT-4.

It's great at coming up with details such as poses, lines, and color without requiring direct references. However, there is no guarantee that Dall-E's output is legally licensed and safe—and the platform self-censors its results more heavily than others, sometimes requiring more time to obtain desired results.









### BRIA is trained exclusively on stock images, making it the most legally permissive Al art model of the bunch.

It doesn't require a powerful computer to function, and allows for high customization of details. However, this model requires custom training to create art that's even halfway good—and at \$7,500 per month, it comes with a bigger price tag than many of its competitors.







### What about Midjourney?



ROOM 8 GROUP 2025 Worth a mention given its popularity, Midjourney was one of the first tools to demonstrate Al's power to create high-quality 2D images. If you're even vaguely familiar with Al art, you've probably heard of Midjourney.

But although this model is capable of generating detailed and lifelike AI images, we do not recommend it for use because it is nearly impossible to keep the results of Midjourney prompts private. At the moment, Midjourney uploads all of its results to a public gallery. If you're working under a non-disclosure agreement, use Midjourney with caution.

### Generative AI 2D art: Our picks



#### THE BEST TOOLS ON THE MARKET:

### Stable Diffusion XL (SDXL)

Stable Diffusion originally launched in 2022. It's one of the most accessible AI art models due to its use of diffusion technology to significantly reduce processing requirements. Stable Diffusion can run on consumergrade graphics cards and allows for rapid iterations on concepts and ideas.

Stable Diffusion XL 1.0, or SDXL, is the latest and greatest version of Stable Diffusion. SDXL builds on previous versions of Stable Diffusion by allowing for the generation of legible text, better human anatomy, and a much wider range of available art styles, ranging from photorealistic images to anime or comic-book art.



### Strengths:

- + Good art
- + Best professional control for now; however, FLUX will most likely have better tools for professional control in the near future
- + Largest number of professional features such as Controlnet models, which allow the use of visual references in art generation
- + Constant updates from a vibrant community, which will likely continue given SDXL's relatively low hardware requirements
- + Vast number of supported interfaces and many custom fine-tunes and customized mini-models (LoRAs); most effective when using LoRAs made with databases of owned images
- + LoRAs can reduce costs by 15%–30%

### Legal status:

SDXL was not necessarily trained on licensed data, creating potential legal challenges when not used locally

Must only be used internally to replace intermediate pipeline stages before finalization

#### Weaknesses:

- Good art, but not the best of the tools we reviewed
- LoRAs can require several days of training; saved
- competitors (such as FLUX.1 Dev)
- correct mistakes
- source development

### Pricing:

Open source, with internal costs for training and hardware

fine-tunes for complex projects can take more than a week, potentially negating the time

– If the SDXL community slows its support, the tool may become out of date soon, and newer models by Stability.Al are worse than some - Unable to provide fully finished results; human artists are still required to avoid plagiarism and

Very resource-hungry and buggy due to open-

Require more expertise than proprietary tools



THE BEST TOOLS ON THE MARKET:

### FLUX1 Dev

FLUX.1 Dev was launched on August 1, 2024, by Black Forest Labs—some of the original creators of Stable Diffusion. Better known simply as FLUX, it is one of the most advanced genAl art tools on the market right now.

FLUX bills itself as the largest state-of-the-art opensource text-to-image AI model yet, capable of creating exceedingly high-quality images from simple text descriptions. Compared to some other AI art models, FLUX is particularly good at creating high-resolution and photorealistic images of humans, all at a faster speed than competitors such as SDXL.



ROOM 8 GROUP | 2025

### Strengths:

- + The strongest AI art tool in both speed and quality
- + Allows for better control when using descriptions than other models, with less variance
- + An open-source, fully professional tool capable of generating finished art

#### Weaknesses:

- referencing), and so on
- Requires an extremely strong GPU to run effectively



 Relatively new—the latest model was launched in August by a group of former Stable Diffusion staffers; this is important because new models need time to be incorporated in interfaces, new tools need to be created for more professional control (such as models to be used for visual

- Model relies heavily on providing an accurate,

### Legal status:

The best open-source model on the market right now; relatively safe from a legal standpoint

### Pricing:

Basic pricing is free; unlimited projects start at \$12 per editor per month

### HOW YOU can use our results

Our findings are useful for game-makers who want to incorporate AI tools into their art production pipelines but lack the time to conduct their own comprehensive review. The recommendations made here—FLUX.1 Dev and SDXL—have been made from the perspective of a leading game developer.

We conducted this review for ourselves at Room 8 Group first: its purpose was to find efficiencies in the 2D art production pipeline (and by extension, for our partners) using whatever genAl tools are most suitable.

In fact, we are already putting these findings to use on real projects (see next page).

As with any new type of software, training and onboarding is required to use genAl to its full potential. Across the board, we found that artists who had rarely taken advantage of AI tools in the past are struggling not only to use them in 2024, but to even understand how to use them within the broader creative pipeline. Artists who were already familiar with photobashing and overpainting techniques were able to use AI more effectively.

As a leading external game developer, our insights are directly applicable to developers and publishers who want to improve an important part of their art production pipelines by making them more efficient, ensuring legal safety, and maintaining a high quality.

While genAl is clearly a disruptive technology for gaming, it will not replace people anytime soon. That said, given industry pressures affecting companies big and small, staying up-to-date with the latest technologies is essential. But for us, it is equally important we do so ethically, in full consideration of our incredibly talented and hard-working artists.







### Putting the results into action

Room 8 Group has already applied the findings of this review—we've created a 10-step training course to teach our specialists internally about best practices in 2D genAl.

The goal: To empower our creatives to work more efficiently, augmenting their artistic skills with a healthy dusting of AI magic.

Our training process is designed in a general, forward-looking manner that can easily be applied to all kinds of AI models in the future. Although SDXL and FLUX.1 Dev are our current tools of choice, we acknowledge they might not be even in a few months' time. Because the AI space is evolving rapidly, our course trains people on only the most important parts of AI tools and pipelines.

Here's how we teach our specialists to become masters of AI art:

- 1. We begin with an introductory lecture covering responsible practices for AI
- 2. We train our artists on the basics of using AI art tools, giving them practical assignments to understand the best ways to prompt the model
- 3. For artists who lack the technology to work with AI locally, we provide lessons about the best ways to apply cloud servers to Al art
- 4. We hold Q&A sessions to answer practical questions and receive feedback

We will continue building out our AI art pipeline by adding advanced courses on the practical challenges of AI art projects, including the differences between positive and negative prompts, as well as the basic composition prompt writers should use in order to best take advantage of AI tools. Our internal workflow also ensures that our specialists are skilled at adjusting aspects such as resolution and seed, as well as developing more advanced minimodels to boost the efficiency of the entire process.

Perhaps the most important part of our internal training is our practical AI policy for artists, which outlines the ethical and legal responsibilities of creatives engaging with AI art. Ultimately, we want our generated results to be as far from generally available public art as possible, meaning the process of generating images should be unique and customized from beginning to end. We want our specialists to make the highest-quality art possible—and we want them to do it in a legally safe and ethical way.







### How Room 8 Group is using FLUX and SDXL to save time and money on real projects

#### LESS THAN 7.5 DAYS TO COMPLETE

ANTICIPATED 12 DAYS

For one partner, we used a custom SDXL mini-model to create new and dynamic illustrations of a group of pre-existing fictional characters.

After integrating AI, we found that we were able to cut the amount of time needed for the project by 30%–50%, with the final project taking less than 7.5 days to complete, rather than an anticipated 12.

ONLY 60% OF TIME NEEDED TO COMPLETE

For another partner, Room 8 Group used both SDXL mini-models as well as FLUX.1 Dev to help design environment concept art in the form of surreal fantasy landscapes.

With FLUX, we found we were able to cut the amount of time needed by over 40%, while maintaining our ability to control the art style and the image's level of abstraction.

The information above isn't just theory. At Room 8 Group, we've already started to apply AI art to our work with partners in the game development space (with their approval)—and so far, the results are very promising.

22 HOURS TO COMPLETE

ANTICIPATED 32 H

Another test project, during which we designed a custom SDXL mini-model to improve a partner's existing character designs before rendering them in a designated style, we found that an asset that normally required 32 hours to design could be finished in only 22 hours using AI. While the training of the AI itself required about 10 hours, this model could be used for future character modeling down the line, saving roughly 30% of the time needed for additional character models.







NEXT STEPS

### Looking to take better advantage of generative AI in your game development pipeline? Room 8 Group can help.

The information in this report is just a fraction of the data our research team collected during their comprehensive review. We are available to apply our findings to your projects.

At Room 8 Group we work with many of the industry's leading publishers and developers, including Microsoft, Nintendo, Ubisoft, Scopely, and dozens more. If you're interested in learning more, let's talk!

#### CONTACT US

ROOM 8 GROUP | 2025



### About Room 8 Group

Room 8 Group is an end-to-end strategic partner in external game development. Working across all platforms, we provide creative and technical expertise across game development, technology, art, trailers, and QA for AAA and AA games.

Since 2011, we've built creative partnerships with world-leading publishers such as Microsoft,

Nintendo, Ubisoft, Sony, Gameloft, Take2, EA, and more. While leveraging our own cutting-edge tools and R&D capabilities, we've co-created a multitude of award-winning projects for video game IPs and franchises including Call of Duty, Diablo, Assassin's Creed, Star Trek, The Walking Dead, Doctor Who, and many more.

Through primarily organic growth, we are proudly independent, now home to around 1,300 specialists based in Europe, North America, and South America.









#### DISCOVER MORE ON WEB

#### FOLLOW US ON LINKEDIN



### Company Awards

WINNER **Best Service Provider** 

WINNER Commitment to Environmental, Social, and Governance

SHORTLISTED **Co-Development Partner** 

Service Provider

Art, Animation, and Trailer Supplier

SHORTLISTED **Best Creative Provider** 



SHORTLISTED Service Provider of the Year





