



EQUITY RESEARCH

UPDATED

06/03/2025

ComfyUI

TEAM

Jan-Erik Asplund
Co-Founder
jan@sacra.com

Marcelo Ballve
Head of Research
marcelo@sacra.com

DISCLAIMERS

This report is for information purposes only and is not to be used or considered as an offer or the solicitation of an offer to sell or to buy or subscribe for securities or other financial instruments. Nothing in this report constitutes investment, legal, accounting or tax advice or a representation that any investment or strategy is suitable or appropriate to your individual circumstances or otherwise constitutes a personal trade recommendation to you.

This research report has been prepared solely by Sacra and should not be considered a product of any person or entity that makes such report available, if any.

Information and opinions presented in the sections of the report were obtained or derived from sources Sacra believes are reliable, but Sacra makes no representation as to their accuracy or completeness. Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information, opinions and estimates contained in this report reflect a determination at its original date of publication by Sacra and are subject to change without notice.

Sacra accepts no liability for loss arising from the use of the material presented in this report, except that this exclusion of liability does not apply to the extent that liability arises under specific statutes or regulations applicable to Sacra. Sacra may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect different assumptions, views and analytical methods of the analysts who prepared them and Sacra is under no obligation to ensure that such other reports are brought to the attention of any recipient of this report.

All rights reserved. All material presented in this report, unless specifically indicated otherwise is under copyright to Sacra. Sacra reserves any and all intellectual property rights in the report. All trademarks, service marks and logos used in this report are trademarks or service marks or registered trademarks or service marks of Sacra. Any modification, copying, displaying, distributing, transmitting, publishing, licensing, creating derivative works from, or selling any report is strictly prohibited. None of the material, nor its content, nor any copy of it, may be altered in any way, transmitted to, copied or distributed to any other party, without the prior express written permission of Sacra. Any unauthorized duplication, redistribution or disclosure of this report will result in prosecution.



ComfyUI

Open-source, node-based application for generating images, videos, and audio using AI

#ai #generative-ai

[Visit Website](#)

Details

HEADQUARTERS

San Francisco, CA

CEO

Yoland Yan



FUNDING

\$16,200,000

2025

Valuation

ComfyUI raised \$16.2 million across multiple funding rounds, with their most recent extension round in May 2025 bringing the total to this figure. The company is backed by notable investors including Pace Capital, Chemistry, Cursor VC, and Guillermo Rauch from Vercel.

The initial seed round in Q4 2024 was led by Pace Capital with participation from Chemistry, Cursor VC, and Stratus Ventures. A subsequent extension round in May 2025 added additional capital from the existing investor base.

While the exact valuation has not been publicly disclosed, the funding history and investor profile suggest the company is positioned for significant growth despite its current pre-revenue status.

Product

ComfyUI is like LEGO blocks for AI image and video creation. Instead of typing a single prompt and hoping for the best, users drag visual building blocks called nodes onto a canvas and connect them with wires to create complex AI workflows.

Each node performs a specific function - one might load an AI model, another adds text prompts, a third applies style controls, and a fourth upscales the final image. Users can branch their workflows to create multiple variations, add video generation nodes, or incorporate 3D asset creation. When they hit run, ComfyUI processes only the parts of the workflow that have changed, making iteration fast and efficient.

The tool runs locally on users' computers and supports virtually any graphics card, from high-end NVIDIA GPUs down to older hardware with just 1GB of memory. It can load any Stable Diffusion model or checkpoint that users provide, and the entire workflow can be saved as a JSON file and shared with others.

Power users love ComfyUI because it gives them precise control over every step of the AI generation process. VFX artists can create reproducible pipelines for generating consistent character designs. Game developers can automate asset creation workflows. Researchers can experiment with new AI models by plugging them into existing node graphs.

The platform supports hundreds of built-in nodes covering everything from basic Stable Diffusion to cutting-edge video models like Stable Video Diffusion, plus a thriving ecosystem of over 10,000 community-created custom nodes that add specialized functionality.

Business Model

ComfyUI operates as an open-source project with a freemium expansion strategy. The core software remains completely free under a GPL-3 license, with monetization coming through optional premium services and cloud hosting.

The go-to-market model is B2C for individual creators and B2B for enterprises. Individual users download and run ComfyUI locally for free, while businesses can purchase managed cloud hosting, enterprise support, and team collaboration features through partners like Comfy Deploy.

The company is introducing API nodes that let users call commercial AI services like Google's Veo 2 or OpenAI's models directly within their workflows. This creates a usage-based revenue stream where ComfyUI takes a margin on API calls while maintaining the free local experience.

The business model creates a flywheel where free users contribute to the ecosystem through custom nodes and workflows, attracting more users and eventually converting some to paid services. The open-source nature builds trust and adoption, while premium features target professional use cases that require reliability, collaboration, and integration with proprietary AI models.

Revenue diversification comes through multiple channels: cloud hosting subscriptions, API usage fees, enterprise licensing, and potential marketplace revenue sharing from the custom node ecosystem.

Competition

Traditional stable diffusion interfaces

Automatic1111 remains the install base leader with over 90,000 GitHub stars but faces criticism for slower feature development and lack of native support for newer models like Flux and SD3.5. The interface uses an extension system rather than ComfyUI's visual node approach, making complex workflows more difficult to build and maintain.

InvokeAI positions itself as an enterprise-focused alternative with multi-user support, SSO integration, and role-based permissions. Their commercial cloud tier directly competes with ComfyUI's enterprise ambitions by offering collaboration features that individual installations cannot match.

Newer entrants like StableSwarmUI and Forge promise auto-scaling and AI-assisted prompt building, showing how quickly the open-source ecosystem fragments and how differentiation increasingly comes from DevOps capabilities rather than just user interface design.

Ease-of-use challengers

Foocus takes a one-button approach that automates sampler, LoRA, and refiner settings, appealing to mainstream users who find ComfyUI's node-based interface intimidating. This creates a funnel problem where ComfyUI might lose potential users who never graduate from simpler tools.

DiffusionBee and Easy Diffusion offer native Mac and Windows applications with drag-and-drop installers, significantly reducing the technical barrier to entry. These tools target the same local-first philosophy as ComfyUI but prioritize simplicity over power-user features.

The risk is that these simplified tools continue adding advanced features while maintaining their accessibility advantage, potentially commoditizing ComfyUI's power-user moat.

Commercial cloud platforms

Midjourney, Adobe Firefly, and DALL-E 3 compete by moving up-market into professional workflows while offering superior convenience and legal indemnification. These platforms eliminate the need for local GPU hardware and model management entirely.

Runway, Leonardo, and Stability AI's DreamStudio offer Stable Diffusion as a service with professional features like team collaboration, version control, and enterprise billing. They compete directly with ComfyUI's cloud hosting partners by providing managed experiences without requiring technical setup.

The competitive threat comes from these platforms gradually adding more customization and control features while maintaining their ease-of-use advantage and eliminating hardware requirements.

TAM Expansion

New products and capabilities

ComfyUI is expanding beyond static images into video generation with Stable Video Diffusion support and 3D asset creation pipelines. The September 2024 release added animation capabilities and latent interpolation, roughly doubling the addressable creator market by moving into motion graphics and 3D workflows.

The roadmap includes becoming an operating system for generative AI that handles environment management and exposes ComfyUI as an API backend. This positions the platform to power third-party applications, mobile frontends, and game engine integrations, extending beyond individual artists to any product embedding generative media.

Professional team features like multi-user workflow management, live preview scaling, and node-level permissions open opportunities for paid ComfyUI Studio offerings targeting agencies, VFX houses, and game studios that already license tools like Houdini or Unreal Engine.

Customer base expansion

Cloud-hosted wrappers like RunComfy and Comfy Deploy remove local GPU requirements and provide managed options for corporate IT departments. This brings in non-technical marketing, design, and product teams who want ComfyUI's capabilities without hardware management complexity.

The marketplace ecosystem with over 10,000 custom nodes creates opportunities for revenue sharing similar to Unity's Asset Store model. Instead of purely open-source distribution, ComfyUI could capture value from workflow sales and premium node licensing while maintaining the free core platform.

Low-specification GPU support and efficient model options make ComfyUI accessible in cost-sensitive regions and mobile-first markets where desktop GPUs are scarce, expanding the global user base significantly.

Adjacent industry integration

Partnerships with model distribution platforms like Hugging Face and CivitAI could create new reseller channels while embedding ComfyUI deeper in the generative AI stack. Bundling with GPU cloud credits from providers like Lambda or CoreWeave reduces barriers to entry for new users.

Toolchain consolidation through acquiring leading workflow marketplaces or cloud hosting frontends would provide instant SaaS revenue while keeping the core engine open source. This hybrid approach balances community growth with commercial sustainability.

The exploding demand for video and 3D assets, combined with falling GPU costs and erosion of traditional Adobe monopolies, creates opportunities for ComfyUI to become the standard creative pipeline for next-generation content production.

Risks

Ecosystem fragmentation: The open-source nature that drives ComfyUI's adoption also enables competitors to fork the codebase and create specialized versions. As hardware vendors like Intel, NVIDIA, and AMD actively court UI projects with optimized backends, ComfyUI faces the risk of its community and developer mindshare being split across multiple incompatible forks, weakening the network effects that currently drive its dominance.

Commercial platform competition: Cloud-based AI image generators like Midjourney and Adobe Firefly are rapidly adding professional features while maintaining superior ease of use and legal indemnification. If these platforms successfully move upmarket with advanced customization tools while eliminating hardware requirements, they could capture ComfyUI's target professional users who prioritize convenience over local control.

Supply chain security: The June 2024 incident involving malicious nodes in the custom node ecosystem highlighted the security risks of ComfyUI's open plugin architecture. As enterprises increasingly adopt the platform, any major security breach through compromised custom nodes could damage trust and slow enterprise adoption, particularly in security-conscious industries that represent ComfyUI's highest-value expansion opportunities.

DISCLAIMERS

This report is for information purposes only and is not to be used or considered as an offer or the solicitation of an offer to sell or to buy or subscribe for securities or other financial instruments. Nothing in this report constitutes investment, legal, accounting or tax advice or a representation that any investment or strategy is suitable or appropriate to your individual circumstances or otherwise constitutes a personal trade recommendation to you.

This research report has been prepared solely by Sacra and should not be considered a product of any person or entity that makes such report available, if any.

Information and opinions presented in the sections of the report were obtained or derived from sources Sacra believes are reliable, but Sacra makes no representation as to their accuracy or completeness. Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information, opinions and estimates contained in this report reflect a determination at its original date of publication by Sacra and are subject to change without notice.

Sacra accepts no liability for loss arising from the use of the material presented in this report, except that this exclusion of liability does not apply to the extent that liability arises under specific statutes or regulations applicable to Sacra. Sacra may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect different assumptions, views and analytical methods of the analysts who prepared them and Sacra is under no obligation to ensure that such other reports are brought to the attention of any recipient of this report.

All rights reserved. All material presented in this report, unless specifically indicated otherwise is under copyright to Sacra. Sacra reserves any and all intellectual property rights in the report. All trademarks, service marks and logos used in this report are trademarks or service marks or registered trademarks or service marks of Sacra. Any modification, copying, displaying, distributing, transmitting, publishing, licensing, creating derivative works from, or selling any report is strictly prohibited. None of the material, nor its content, nor any copy of it, may be altered in any way, transmitted to, copied or distributed to any other party, without the prior express written permission of Sacra. Any unauthorized duplication, redistribution or disclosure of this report will result in prosecution.

Published on Jun 03rd, 2025