Hunyuan3D 3.0 System Card

Tencent Hunyuan*

September 16, 2025

Figure 1: High quality 3D assets generated by Hunyuan3D 3.0.

1 Introduction

In recent years, 3D generation has undergone rapid development. Early works such as CLAY [1], Trellis [2], Hunyuan3D-2 [3, 4, 5], and TripoSG [6] established diffusion models as the mainstream approach for 3D generation. The introduction of Hunyuan3D 2.5 marked a significant milestone, elevating this line of work to a new level of capability. Subsequently, methods like Direct3D-S2 [7] and Sparc3D [8] have also achieved remarkable progress. The field of 3D generation is now advancing rapidly toward ultra-high-fidelity synthesis.

*Full authorship contribution statements appear at the end of the document.

Figure 2: Illustration of better geometry details and image-mesh following in Hunyuan3D 3.0.

Hunyuan3D 3.0 introduces a new geometry diffusion generation model with better detail preserving, image-mesh following, and stability. The model uses the same architecture as LATTICE, which is initially introduced in Hunyuan3D 2.5 [9]. Compared to its previous version, Hunyuan3D 3.0 is trained with a larger-scale high-quality 3D dataset, using more strict data filtering, and better data processing. Thanks to the scalable architecture of LATTICE, we continue scaling up the compute to build a larger model that achieves consistent improvement on geometry generation quality. This report highlights the key features and provides a thorough evaluation on the various properties, including image-mesh following, mesh resolution, surface quality, etc.

2 Highlights

Better Geometry Details. Hunyuan3D 3.0 significantly enhances the capability to generate intricate and high-fidelity geometric structures. Through the use of a larger and more rigorously curated 3D dataset, combined with improved training strategies and scaled model size, the model produces geometries with finer details and higher structural accuracy, making it suitable for applications requiring high-resolution and semantically rich 3D assets. Figure 2 demonstrates the comparison between Hunyuan3D 2.5 [9] and Hunyuan3D 3.0 for single-image-to-3D and multi-images-to-3D.

Better Image-Mesh Following. The new model demonstrates superior alignment between input images and generated 3D meshes. This is not only attributes to the improvement on detail generation, but also thanks to the enhancements in the conditioning mechanism, which ensure that the output mesh more faithfully reflects the geometric and semantic cues present in the reference image, as shown in Figure 2.

Kept Surface Smoothness. While enhancing geometric details, Hunyuan3D 3.0 maintains excellent surface smoothness and overall mesh quality, as revealed in Figure 3. The use of advanced LATTICE architecture mitigates artifacts and produces clean, well-structured surfaces suitable for both visualization and downstream applications.

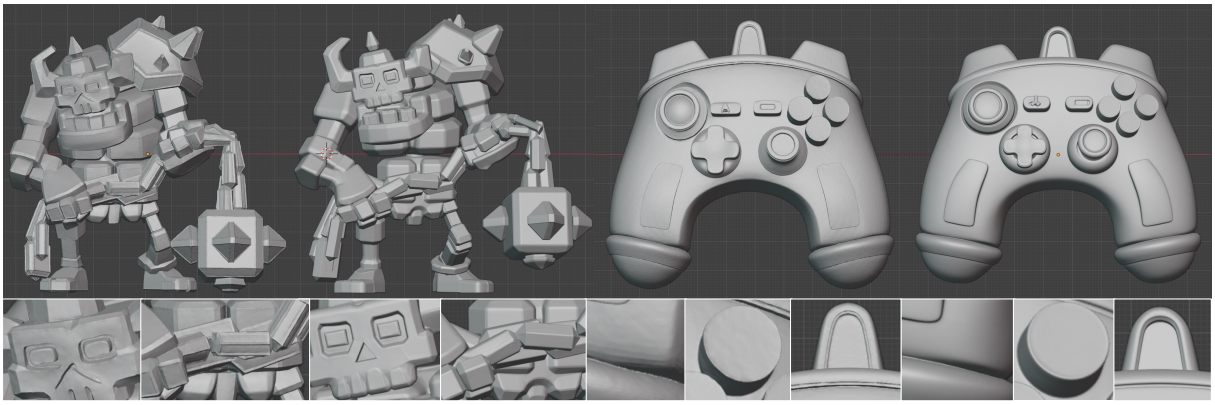

Figure 3: Illustration of surface smoothness in Hunyuan3D 3.0.

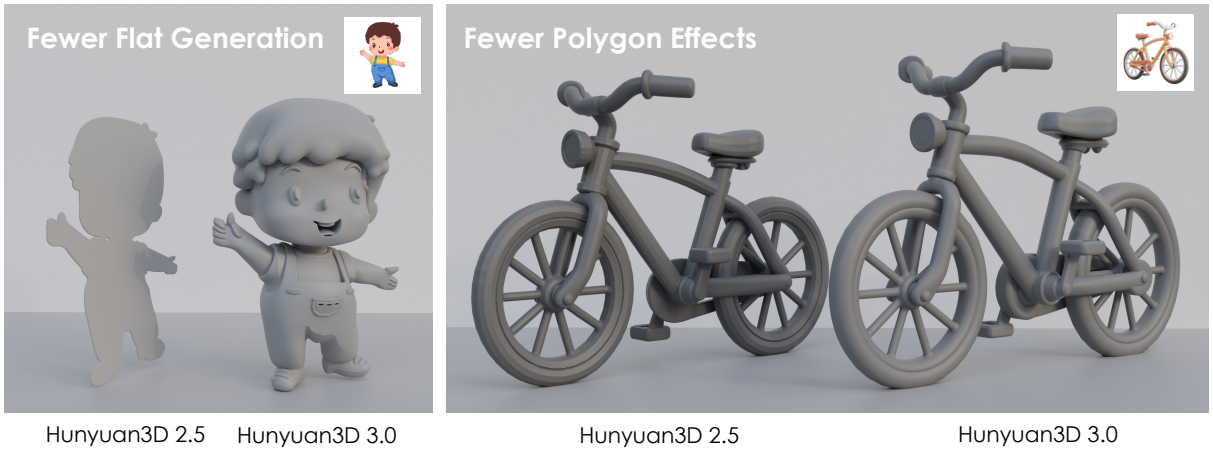

Figure 4: Illustration of the improved stability of Hunyuan3D 3.0.

Improved Stability. In Hunyuan3D 3.0, we pay great attention to the generation stability. In Figure 4, we show two representative corner cases that our previous model may occurs, including flat generation and polygon effects. In Hunyuan3D 3.0, we improve the training strategies and data processing, greatly reducing the probability of occurrence.

More Realistic Texture Performance: Hunyuan3D 3.0 delivers refined textures that more faithfully reproduce object appearances. By seamlessly integrating high-resolution texture training data with physically-based rendering algorithms, the model produces texture maps rich in material details, precise specular reflections, and natural color transitions. As illustrated in the Figure 5, these enhancements significantly elevate the model's expressive power and visual fidelity.

3 Evaluation

User Study. We conducted a large-scale user study to evaluate the perceptual quality of results produced by Hunyuan3D V3.0 against its predecessor V2.5. The study covered 7 representative sub-categories, including human characters, animals, daily objects, natural/virtual scenes, and etc. For each comparison, users were presented with paired 3D models generated by both versions and asked to indicate their preference. As shown in Figure 6, V3.0 achieves consistently higher win rates across all categories, with particularly significant improvements observed in characters

Figure 5: Illustration of the texture improvements of Hunyuan3D 3.0.

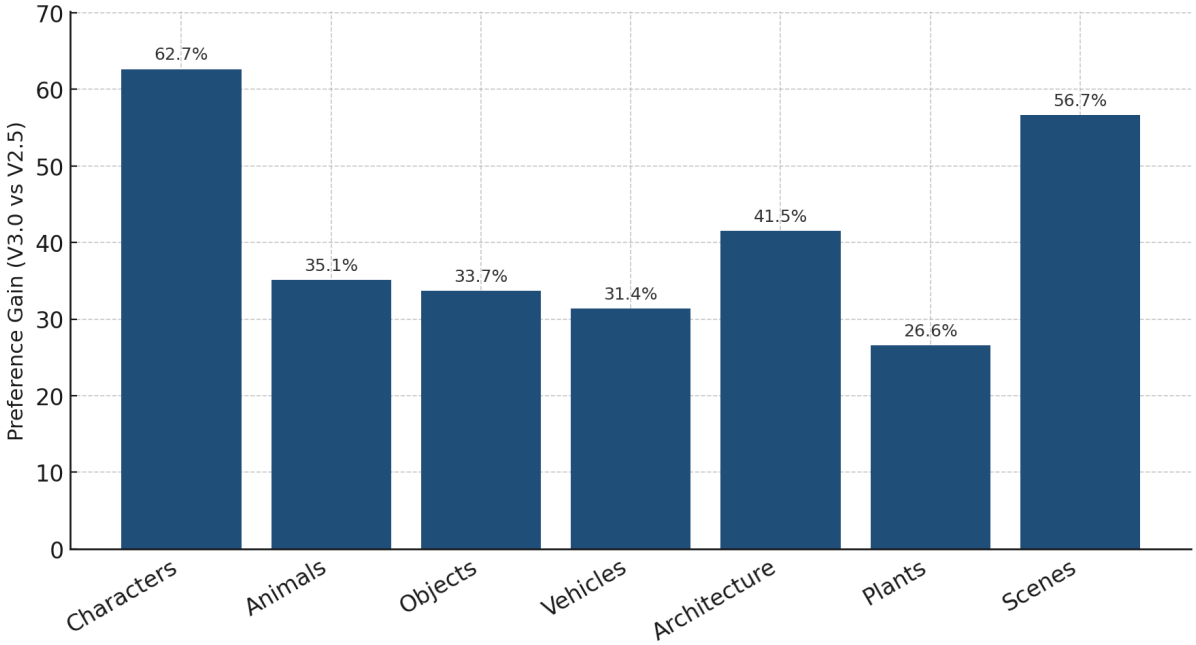

Figure 6: Win rate difference (%) of Hunyuan3D V3.0 compared to V2.5 across 7 sub-categories, showing consistent improvements in most categories.

(+62.7%), scenes (+56.7%), and architecture (+41.5%). These results demonstrate that the proposed improvements in details generation and human characters.

4 Best Practice for Hunyuan3D 3.0

Hunyuan3D 3.0 is an image-to-3D generation model. To fully unleash model capabilities, it is important to keep the input image high quality. Here are some recommendations:

- **High resolution**: The image should have a resolution of at least 512×512 pixels—the higher, the better. Images with excessively low resolution may lose important details.
- **Clear object without background**: Hunyuan3D 3.0 is an object generation model. Although it can generate multiple objects or small scenes at once, for optimal results, it is recommended to use an image with a clean background and clearly visible objects. The

object should be fully visible and unobstructed. If necessary, use image editing tools to extract the main subject from the original image.

- **Remove flat textures**: In some cases, it may be difficult to determine whether geometry should be generated from a flat texture. To avoid ambiguity, it is advisable to remove any unwanted textures using image editing software beforehand.
- **Multiview generation**: Whenever possible, multiview generation is highly recommended, as it typically yields better results—especially for the back views. If such views are not available, consider generating additional images using image editing tools.
- **Recommend using images without heavy shadows**: Strong shadows may result in incomplete separation of lighting information during the PBR decomposition process, leading to residual light variations in the final output. To achieve cleaner material results, it is advisable to ensure even lighting when capturing or selecting source images, or preprocess the images using editing tools to minimize shadow interference.

5 Conclusion

Hunyuan3D 3.0 represents a stable upgrade over its predecessor, offering improved detail generation, tighter alignment between image and mesh, and greater overall stability. Nevertheless, certain issues remain, such as difficulties in generating very fine details and occasional inaccuracies in perspective structures, which we will continue to improve in the future.

Authorship

Project Sponsors
Jie Jiang, Linus

Project Leaders
Chunchao Guo, Jingwei Huang, Zeqiang Lai

Core Contributors
Zeqiang Lai, Yunfei Zhao, Jingwei Huang, Haolin Liu, Qingxiang Lin, Zibo Zhao, Zhihao Liang, Bowen Zhang, Fan Yang, Xianghui Yang, Yifei Feng, Xin Huang, Li Wang, Xin Yang

Contributors
Lifu Wang, Sicong Liu, Yixuan Tang, Yulin Cai, Tian Liu, Yuhong Liu, Kaiyi Zhang

References

[1] L. Zhang, Z. Wang, Q. Zhang, Q. Qiu, A. Pang, H. Jiang, W. Yang, L. Xu, and J. Yu, "Clay: A controllable large-scale generative model for creating high-quality 3d assets," *ACM Transactions on Graphics (TOG)*, vol. 43, no. 4, pp. 1–20, 2024.

[2] J. Xiang, Z. Lv, S. Xu, Y. Deng, R. Wang, B. Zhang, D. Chen, X. Tong, and J. Yang, "Structured 3d latents for scalable and versatile 3d generation," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 21469–21480, 2025.

[3] Z. Zhao, Z. Lai, Q. Lin, Y. Zhao, H. Liu, S. Yang, Y. Feng, M. Yang, S. Zhang, X. Yang, *et al.*, "Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation," *arXiv preprint arXiv:2501.12202*, 2025.

[4] Z. Lai, Y. Zhao, Z. Zhao, H. Liu, F. Wang, H. Shi, X. Yang, Q. Lin, J. Huang, Y. Liu, *et al.*, "Unleashing vecset diffusion model for fast shape generation," *arXiv preprint arXiv:2503.16302*, 2025.

[5] T. Hunyuan3D, S. Yang, M. Yang, Y. Feng, X. Huang, S. Zhang, Z. He, D. Luo, H. Liu, Y. Zhao, *et al.*, "Hunyuan3d 2.1: From images to high-fidelity 3d assets with production-ready pbr material," *arXiv preprint arXiv:2506.15442*, 2025.

[6] Y. Li, Z.-X. Zou, Z. Liu, D. Wang, Y. Liang, Z. Yu, X. Liu, Y.-C. Guo, D. Liang, W. Ouyang, *et al.*, "Triposg: High-fidelity 3d shape synthesis using large-scale rectified flow models," *arXiv preprint arXiv:2502.06608*, 2025.

[7] S. Wu, Y. Lin, F. Zhang, Y. Zeng, Y. Yang, Y. Bao, J. Qian, S. Zhu, X. Cao, P. Torr, *et al.*, "Direct3d-s2: Gigascale 3d generation made easy with spatial sparse attention," *arXiv preprint arXiv:2505.17412*, 2025.

[8] Z. Li, Y. Wang, H. Zheng, Y. Luo, and B. Wen, "Sparc3d: Sparse representation and construction for high-resolution 3d shapes modeling," *arXiv preprint arXiv:2505.14521*, 2025.

[9] Z. Lai, Y. Zhao, H. Liu, Z. Zhao, Q. Lin, H. Shi, X. Yang, M. Yang, S. Yang, Y. Feng, *et al.*, "Hunyuan3d 2.5: Towards high-fidelity 3d assets generation with ultimate details," *arXiv preprint arXiv:2506.16504*, 2025.