

Magic3D: High-Resolution Text-to-3D Content Creation

Chen-Hsuan Lin* Jun Gao* Luming Tang* Towaki Takikawa* Xiaohui Zeng* Xun Huang Karsten Kreis Sanja Fidler† Ming-Yu Liu† Tsung-Yi Lin

*†: equal contributions



JUNE 18-22, 2023
CVPR VANCOUVER, CANADA



Overview

Magic3D is a new text-to-3D content creation tool that generates high-quality 3D meshes.

We provide users with new ways to control 3D synthesis, opening up new avenues to various creative applications.

Major benefits

- Fast (40 minutes, 2× faster than DreamFusion^[1])
- Uses high-resolution (512 × 512) diffusion priors

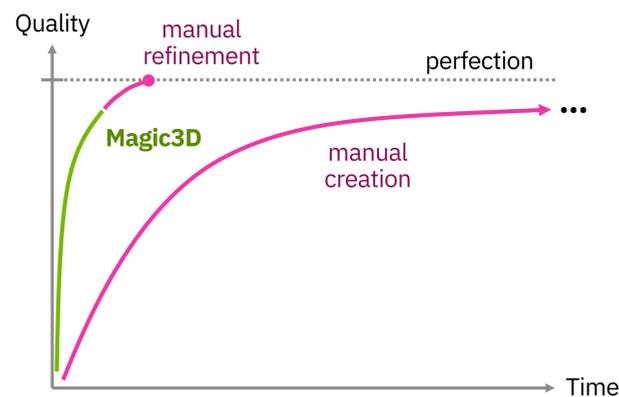
Potential applications

- Movie/game production
- Metaverse
- Robot/self-driving simulation
- 3D designs



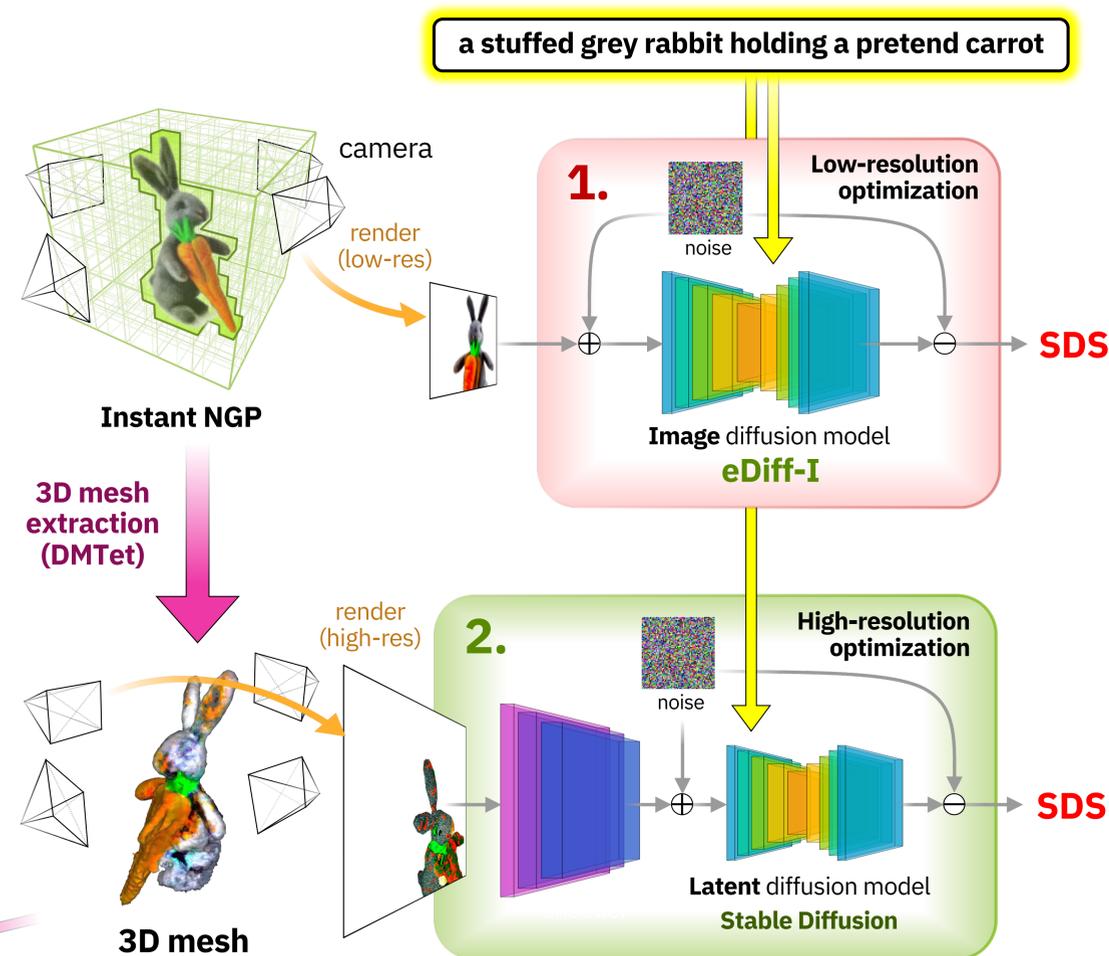
Our goal

- Turbocharge expert 3D artists
- Facilitate 3D content creation for novices



High-resolution 3D mesh

Approach



We optimize the 3D content over pretrained text-to-image diffusion models with SDS (Score Distillation Sampling) loss in a coarse-to-fine fashion, with the following procedure:

- Use a **low-resolution** image diffusion prior and optimize a neural field representation to obtain the coarse model.
- Extract a textured 3D mesh from the density/color neural fields.
- Fine-tune the mesh with a **high-resolution** latent diffusion model.

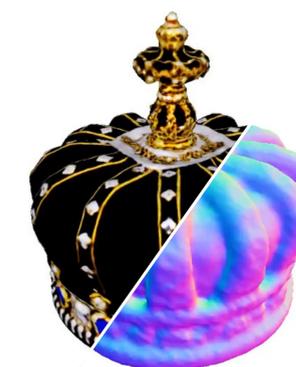
Results



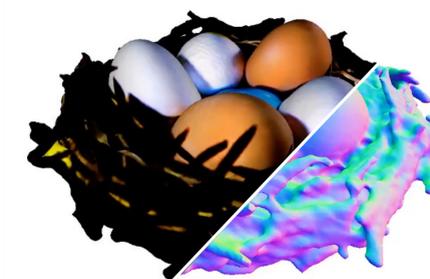
Neuschwanstein castle, aerial view.



A car made out of sushi.



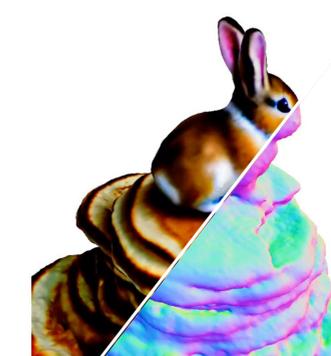
An imperial state crown of England.



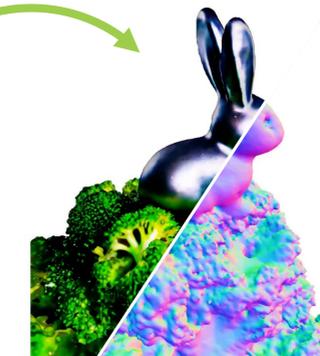
A nest with a few white eggs and one golden egg.



An iguana holding a balloon.



A baby bunny sitting on top of a stack of pancakes.



A metal bunny sitting on top of a stack of broccoli.



A sphinx sitting on top of a stack of chocolate cookies.

continue optimization with modified prompt

[1] Poole et al. "DreamFusion: Text-to-3D using 2D Diffusion." ICLR 2023.