Enhancing High-resolution 3D Generation through *Pixel-wise Gradient Clipping*

Zijie Pan[1] Jiachen Lu[1] Xiatian Zhu[2] Li Zhang[1]
[1] Fudan University [2] University of Surrey

Contact: lizhangfd@fudan.edu.cn
Paper, code, and demo are available: https://fudan-zvg.github.io/PGC-3D

Motivation

In 3D generation with Latent Diffusion Model (LDM) and Score Distillation Sampling (SDS) , gradient propagation pathway is ambiguious between SDS and VAE training.

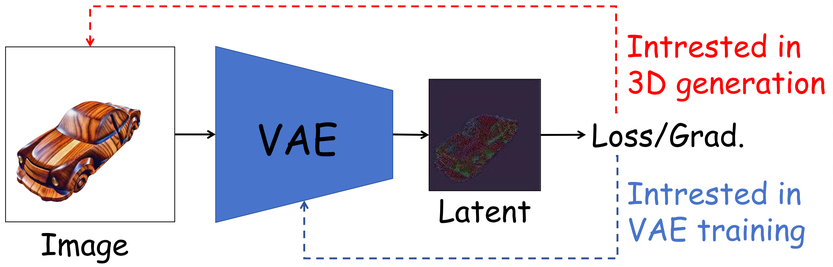

Solution exploring

We find the gradient issue from failure of vanilla SDS through VAE encoder and success of latent/linear approximated gradients. By exploring gradient suppression methods, we find the solution: **Pixel-wise Gradient Clipping**

Text prompt: a wooden car. Diffusion guidance: SDXL

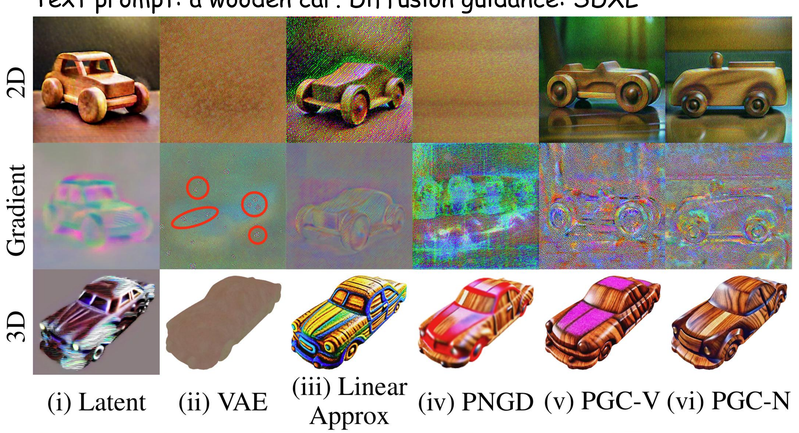

Pixel-wise Gradient Clipping (PGC)

PGC: Given clipping value $c \in \mathbb{R}$ and a pixel's gradient $g \in \mathbb{R}^3$,

$$PGC(g) = \frac{\min(\|g\|, c)}{\|g\|} g.$$

Based on a noise assumption, PGC effectively remove noises.

- highlights noise.

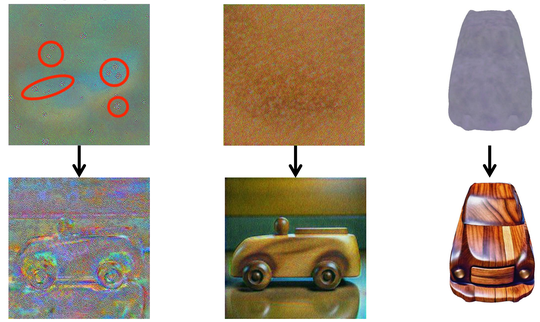

Application: mesh painting

Experiments

PGC exhibits amazing improvements in **fine stage (mesh optimization)**.

PGC also exhibits improvement in **coarse stage**.

* Videos corresponding to all the images above are available at: *https://fudan-zvg.github.io/PGC-3D*.