DI-PCG: Diffusion-based Efficient Inverse Procedural Content Generation for High-quality 3D Asset Creation

Supplementary Material

6. More Implementation Details

We use six procedural generators from Infinigen and Infinigen Indoors, namely chair, table, vase, basket, flower and dandelion generators. They contain 48, 19, 12, 14, 9, 15 controllable parameters, respectively. These are also the input token lengths of each diffusion models, as the procedural parameters directly serve as the denoising variables. **Our code will be released once the paper is public**.

7. More ablations

We show more ablation studies on view numbers and data augmentations in Table 3. Using multiple views and data augmentations are both crucial, and more views give better results. Binary masking is useful to improve robustness over colors. Edge map augmentation does not affect the quantitative results, while in practice it improves the sketch input performance.

8. More Qualitative Results

Here we show more qualitative results of DI-PCG. The generation results for the chair, table, and vase categories are shown in Figure 9. DI-PCG can handle complex shape variations and details, generating high-quality 3D models from input single images. The results for the basket, flower, and dandelion are shown in Figure 10. These categories intrinsically have a bit fewer variations due to the somewhat limited generality of these three procedural generators from Infinigen. Despite that, our method can capture the geometric details and recover the appropriate parameters for the input images, generating fine 3D geometries.

DI-PCG can effectively handle sketch input as conditions. We show qualitative examples in Figure 11. In our experiments, we observe that DI-PCG works just as well on sketch inputs as on RGB image inputs. This provides DI-PCG more flexibility and less burden to cooperate with artists.

We also provide some visual examples of our quantitative evaluations on DI-PCG's test split and ShapeNet chair subset. As shown in Figure 12 and 13, compared to existing SOTA reconstruction and generation methods, DI-PCG delivers much better 3D models with neat geometry.

9. Discussions and Failure Cases

As discussed in the main paper, DI-PCG is limited by the generality and granularity of the given procedural generators. Although the adopted generator from Infinigen can cover a wide range of common variations of the corresponding category, it still has obvious boundaries. Figure 8 shows some failure cases. The input chair images are out-of-domain samples for Infinigen chair generators, thus DI-PCG can not generate precisely aligned 3D models. Instead, it outputs the closest parameter sets to approximate the images. Although bounded by the procedural generator, DI-PCG focus on the efficient inverse ability of PCG, and represents a general tool to easily and effectively control any existing procedural generator, facilitating their usage in 3D content creation. As proceudral generators are getting increasing attention and become mature to develop thanks to the modern design softwares, the available number and cover range of existing procedural generators are rapidly growing, which can further benefit DI-PCG. DI-PCG can be applied for any procedural generator, to greatly enhance its controllability. Moreover, in the future, utilizing AI techniques, such as Large Language Model (LLM), to generate procedural generation programs could be possible and exciting. AI-generated procedural generators and DI-PCG can naturally work together, to form a new paradigm of 3D content generation.

Figure 8. Some failure cases.

View number	1	1	6	6	12	12	24	24	No	No
Augmentation		✓		✓		✓		✓	Binary	Edge
CD↓	0.139	0.136	0.135	0.109	0.129	0.093	0.125	**0.092**	0.098	0.093
EMD↓	0.140	0.138	0.138	0.123	0.131	**0.108**	0.127	**0.108**	0.111	**0.108**
F-Score↑	0.321	0.329	0.331	0.402	0.353	0.452	0.409	**0.456**	0.443	0.453

Table 3. Ablation studies on ShapeNet chairs subset.

Figure 9. More qualitative results for chair, table, and vase generations. Input images are collected from the internet. DI-PCG can handle diverse input images with various styles, views and textures. It accurately captures geometric details in the input images and generates high fidelity 3D models, facilitating downstream applications.

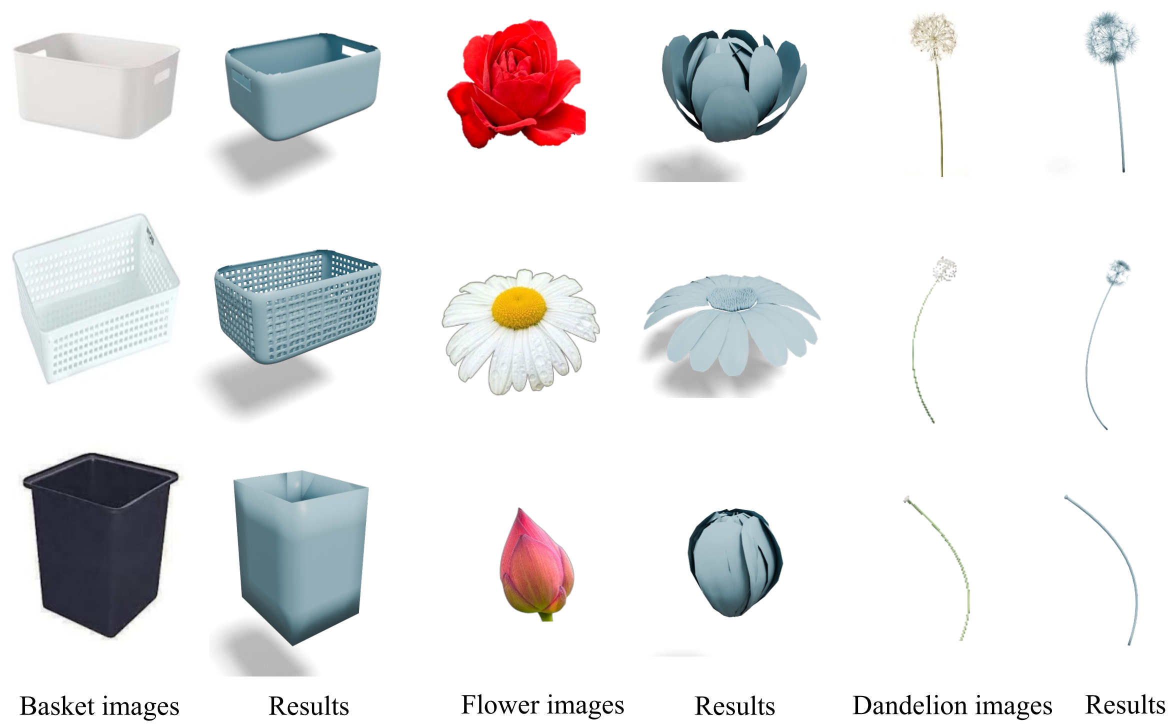

Figure 10. More qualitative results for basket, flower and dandelion generations.

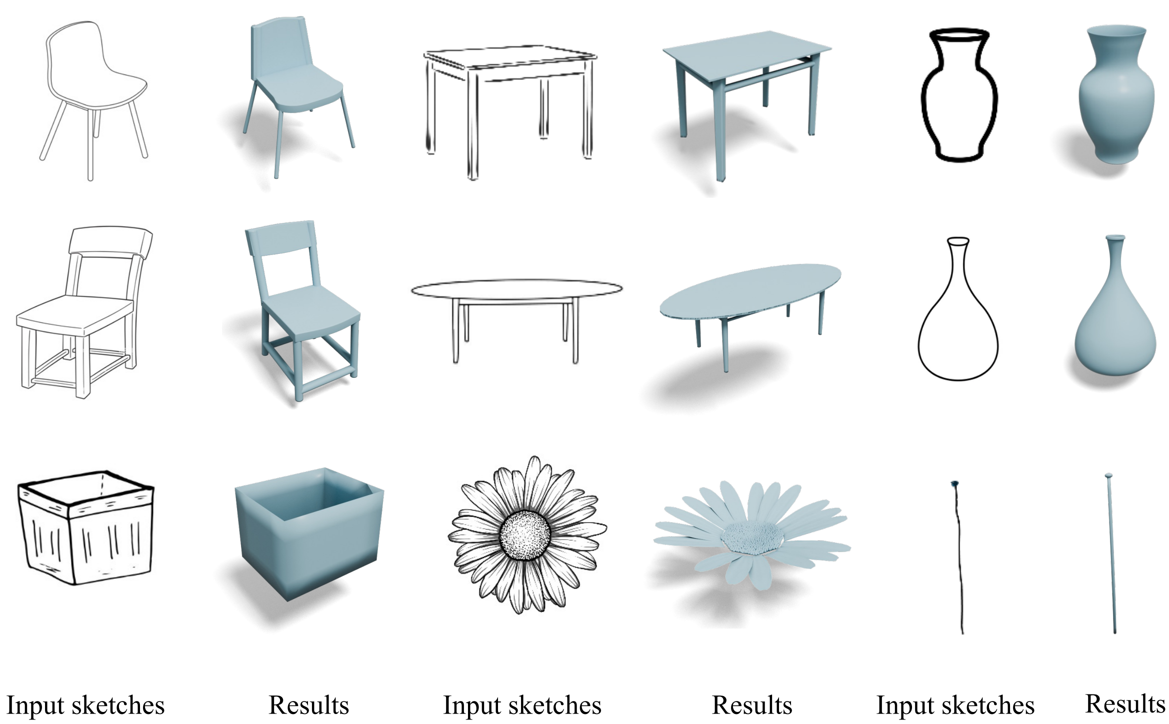

Figure 11. More qualitative results for sketch inputs. DI-PCG can effectively process sketch inputs, offering a convenient way to design and edit objects.

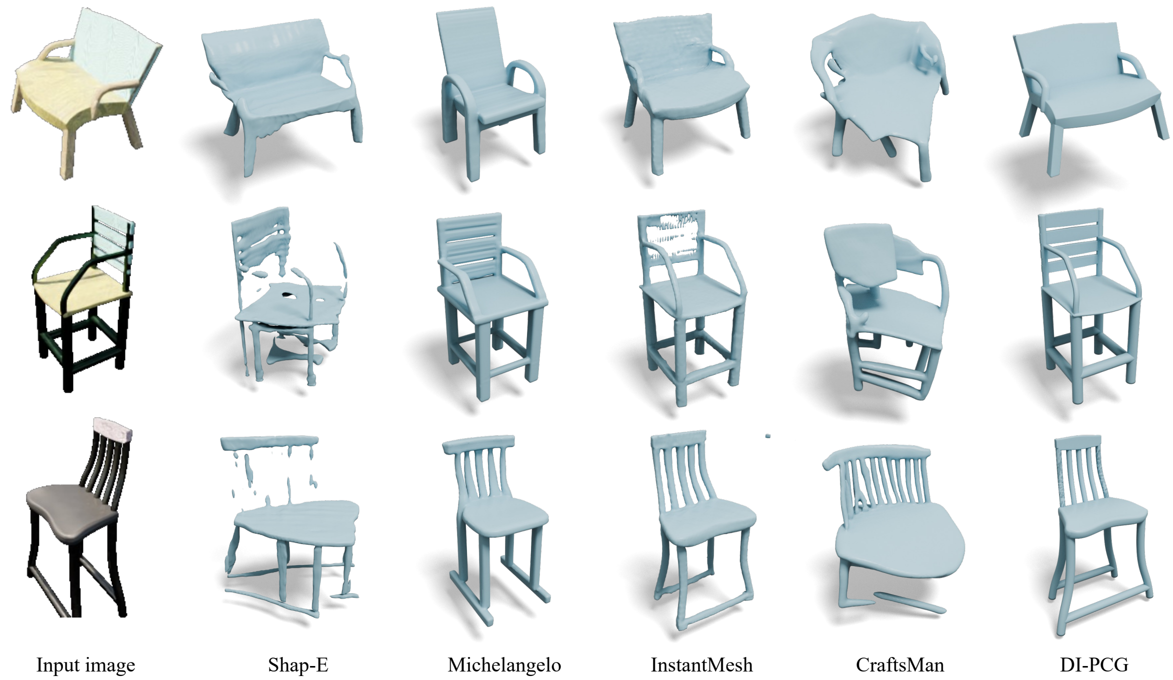

Figure 12. Example comparisons on DI-PCG's test split of chairs. Only DI-PCG generates aligned and clean 3D models.

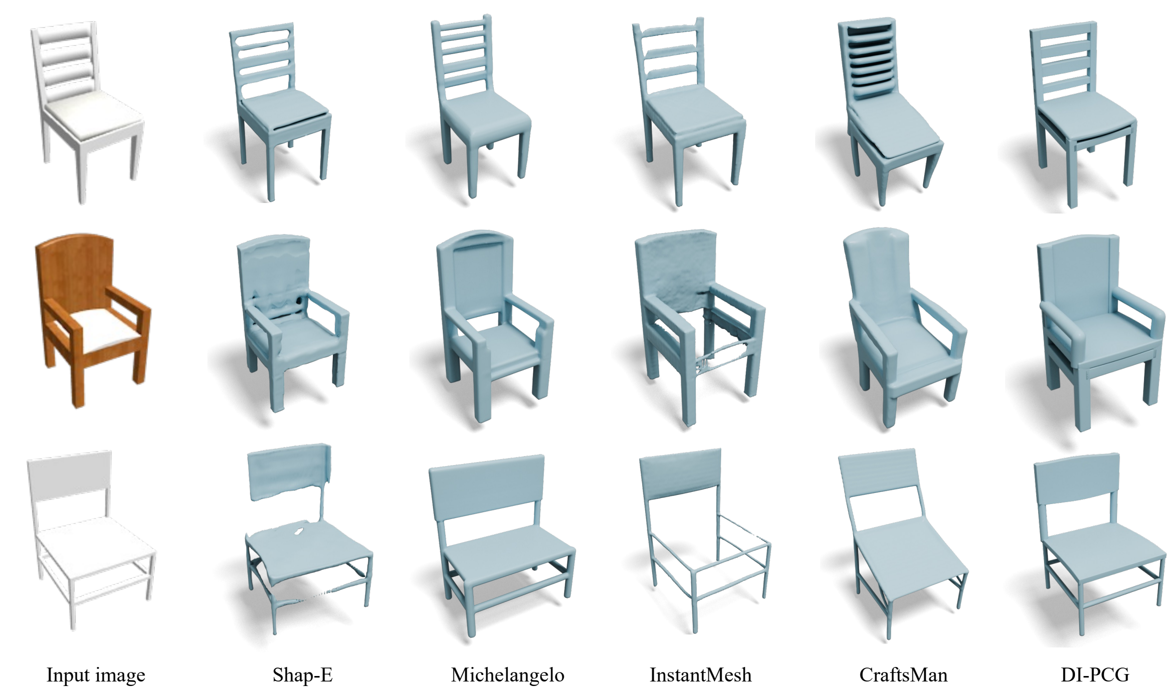

Figure 13. Example comparisons on ShapeNet chair subset. DI-PCG generalizes well and produce high quality 3D meshes.