

View Transformation Robustness for Multi-View 3D Object Reconstruction with Reconstruction Error-Guided View Selection

Qi Zhang¹, Zhouhang Luo^{2,1}, Tao Yu¹, Hui Huang^{1*}

¹College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China

²Guangdong Laboratory of Artificial Intelligence and Digital Economy (Shenzhen), Shenzhen, China
qi.zhang.opt@gmail.com, {luozhouhang2022, yutao2023}@email.szu.edu.cn, hhzhian@gmail.com

Abstract

View transformation robustness (VTR) is critical for deep-learning-based multi-view 3D object reconstruction models, which indicates the methods' stability under inputs with various view transformations. However, existing research seldom focuses on view transformation robustness in multi-view 3D object reconstruction. One direct way to improve the models' VTR is to produce data with more view transformations and add them to model training. Recent progress on large vision models, particularly Stable Diffusion models, has provided great potential for generating 3D models or synthesizing novel view images with only a single image input. To fully utilize the power of Stable Diffusion models without causing extra inference computation burdens, we propose to generate novel views with Stable Diffusion models for better view transformation robustness. Instead of synthesizing random views, we propose a reconstruction error-guided view selection method, which considers the reconstruction errors' spatial distribution of the 3D predictions and chooses the views that could cover the reconstruction errors as much as possible. The methods are trained and tested on sets with large view transformations to validate the 3D reconstruction models' robustness to view transformations. Extensive experiments demonstrate that the proposed method can outperform state-of-the-art 3D reconstruction methods and other view transformation robustness comparison methods.

Introduction

Voxel-based multi-view 3D object reconstruction outputs 3D voxels of the object by fusing various viewpoints with deep-learning neural networks. It plays an important role in computer vision, robotics, augmented reality, and other domains. Naturally, multi-view 3D object reconstruction shall perform stably with inputs of different viewpoints, i.e., the view transformation robustness (VTR) is vital. However, few research works have studied the view transformation robustness for multi-view 3D object reconstruction.

Instead, the VTR issue has been dealt with in other areas, such as in multi-view classification (Wei, Yu, and Sun 2020) or point cloud recognition (Qi et al. 2017; Wang, Yang, and Tao 2022). However, these methods are not specially designed to deal with the view transformation robustness is-

*Corresponding author.




Figure 1: The method's main idea is to utilize the proposed reconstruction error-guided view selection method for selecting and generating viewpoints covering the most errors.

sue for multi-view 3D object reconstruction. So, proposing novel methods for boosting view transformation robustness is important for multi-view 3D object reconstruction.

Recent progress in large vision models (LVMs), especially stable diffusion models (SD) (Rombach et al. 2022; Hu et al. 2024; Shi et al. 2023; Qi et al. 2024; Yoo et al. 2024; Zou et al. 2024; Kwak et al. 2024; Bauer et al. 2024), has provided the extra potential for multi-view 3D object reconstruction (Liu et al. 2024, 2023c,a), or novel view synthesis. While utilizing LVMs for 3D reconstruction causes heavy computation burdens at the inference stage and their view transformation robustness (Liu et al. 2023b) is still limited. An effective way of utilizing stable diffusion models (Liu et al. 2023b) for boosting the view transformation robustness of multi-view 3D reconstruction, is to use them as a data augmentation platform to aid the model training by randomly synthesizing novel views (see Figure 1). However, even though these randomly generated views are useful for increasing the model's robustness to view transformations, the repeated or similar views are also being synthesized with no aid for the model's robustness, reducing the approach's effectiveness. Thus, **instead of generating random views, and considering the error spatial distributions of 3D reconstruction predictions, we propose a novel reconstruction error-guided view selection method to choose the most effective views to improve the 3D reconstruction models' view transformation robustness (VTR).**

Overall, the proposed method consists of three components (see Figure 2): *the original 3D object reconstruction*

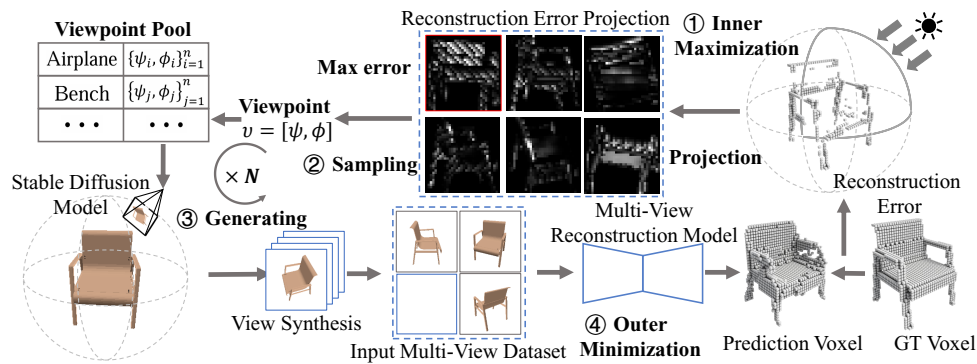


Figure 2: The pipeline of the method: first, the reconstruction error-guided view selection module chooses the views covering most reconstruction errors; then, the stable diffusion model works as a strong data platform for producing multi-view images of selected viewpoints; moreover, the new view images are added to finetune the 3D object reconstruction model. These steps are alternatively conducted and trained to increase the 3D reconstruction models’ robustness to view transformations.

model, the view selection module, and the stable diffusion model-based view synthesis module. The original 3D object reconstruction model could be any deep learning-based multi-view 3D object reconstruction model (e.g., Xie et al. (2020), or Yang et al. 2023). It is first trained on the training set with limited view angles, thus it cannot perform well on the test set with large view transformations. The view selection module is guided by the 3D reconstruction error, where the view angles which can cover more 3D object reconstruction errors are more likely to be selected. Then, the selected view parameters are fed into the view synthesis model to generate new views. The view synthesis module generates novel new views based on the stable diffusion model with the learned view parameters from the view selection module. Later, the newly synthesized views are added to the training set to fine-tune the 3D object reconstruction model further. These steps are recycled to gradually improve the 3D reconstruction model’s view transformation robustness.

Besides, since existing 3D object reconstruction datasets (e.g., ShapeNet (Chang et al. 2015)) are shot from view angles roughly around the objects with small viewpoint ranges (denoted as ‘Aligned’ data), to study the view transformation robustness, we generate a new dataset *ShapeNet-VTR* with more view angle distributions: the ‘Hemispherical’ and ‘Spherical’ set (see Figure 5). **The 3D reconstruction model is trained on the training set of ‘Aligned’ data and tested on the ‘Hemispherical’ and ‘Spherical’ sets, to validate the method’s robustness to view transformations** (see details in Sec. Experiments). The paper’s contributions are summarized as follows.

- This is the first study on view transformation robustness (VTR) for multi-view 3D object reconstruction. Seldom research has focused on the issue of view transformations in this area. Besides, we propose a new dataset specially designed for studying the issue.
- We propose a novel reconstruction error-guided view selection method for choosing the most effective views, which is much more effective compared to randomly generating views or data augmentation methods.

- We use existing stable diffusion models for boosting the view transformation robustness of current 3D object reconstruction models, without training a new one or deploying them at the inference stage, without causing extra computation burdens. Thorough experiments demonstrate the advantages of the proposed method over 3D reconstruction SOTAs and other view transformation robustness comparison methods.

Related Work

Multi-View 3D Object Reconstruction

Deep learning-based multi-view 3D object reconstruction methods have achieved remarkable performance. Early methods focus on the fusion of multi-view features, such as (Su et al. 2015; Paschalidou et al. 2018; Huang et al. 2018; Yao et al. 2018), where the multi-view features are reduced into a fixed size of feature maps. Later methods tried recurrent neural network fusion methods (Choy et al. 2016; Kar, Häne, and Malik 2017; Dosovitskiy et al. 2020; Yao et al. 2019), regarding input views as a sequence. However, RNN fusion methods are not invariant to input order permutations and cannot handle a large number of views efficiently. Attention-based fusion methods (Yang et al. 2020; Xie et al. 2019; Zhu et al. 2023c; Xie et al. 2020) are also proposed to fuse multi-views according to the attention maps estimated from the attention subnet. Recent methods utilize transformer networks for more complicated fusion between the views, such as (Shi et al. 2021; Tiong, Sigmund, and Teoh 2022; Wang et al. 2021; Yagubbayli et al. 2021; Arshad and Beksi 2023; Yang et al. 2023; Zhu et al. 2023b). LRGT (Yang et al. 2023) proposed long-range grouping attention for grouping tokens from all views with separate attention operations. *However, none of these methods have ever paid attention to the multi-view 3D object reconstruction methods’ robustness to the input view transformations.*

3D View Transformation/Rotation Robustness

The issue of view transformation and rotation robustness has been studied in various areas, such as point cloud classifica-

tion and multi-view image classification. Instead of using rotation invariant descriptors (Spezialetti, Salti, and Stefano 2019; Zhu et al. 2023a) as inputs or designing rotation equivariant networks (Shen et al. 2020; Cohen et al. 2021), ART-Point (Wang, Yang, and Tao 2022) improved rotation robustness by training the point classifier on inputs with adversarial rotations. Considering that the camera viewpoints are often fixed for all shapes in multi-view 3D shape classification, Hamdi et al. (2022a) proposed the Multi-View Transformation Network (MVTN), which uses differentiable rendering to determine optimal view-points for 3D shape recognition. Recent methods have started to focus on novel view synthesis for boosting the model viewpoint robustness. View-Fool (Dong et al. 2022) found adversarial viewpoints that mislead 3D recognition models with an entropic regularizer. VIAT (Ruan et al. 2023) proposed to improve the viewpoint robustness of multi-view image classification with the inner diverse adversarial viewpoints and the outer viewpoint invariant classifier training. *Overall, synthesizing novel viewpoints and selecting the most optimal views are the keys for these methods to improve the models’ view transformation or rotation robustness. Note that none of these methods have focused on the view transformation robustness issue for multi-view 3D object reconstruction.*

Method

The main idea of the method is to propose a novel reconstruction error-guided view selection method to select viewpoints that influence the reconstruction results most, then these viewpoint parameters are fed into an existing Stable Diffusion (SD) model to generate the selected view images for finetuning the original 3D reconstruction model for better view transformation robustness. The proposed method consists of **the multi-view 3D reconstruction model**, **the view selection module**, and **the view synthesis module**, whose details are as follows.

Multi-View 3D Reconstruction Model

A multi-view 3D object reconstruction model f usually consists of the following parts (Choy et al. 2016): First, an encoder to extract the feature representation for reconstruction from the image set (I_1, \dots, I_n) of Object I ; Second, multiple 2D image features or 3D volume features are fed into a merger to fuse features from different views; Third, a decoder to predict the corresponding voxel-base 3D shape V from the features map. The overall prediction process is formulated as: $V = f(I_1, \dots, I_n)$. We adopt two typical kinds of multi-view 3D object reconstruction models, *CNNs-based Pix2Vox++* (Xie et al. 2020) and *transformer-based LRGT* (Yang et al. 2023). Denote the 3D voxel ground truth as V^{gt} , and the reconstruction loss is $Loss(V, V^{gt})$ (details in Supp.). Following (Xie et al. 2020) and (Yang et al. 2023), we use binary cross entropy loss in Pix2Vox++ while Dice loss (Milletari, Navab, and Ahmadi 2016) in LRGT. The model is first trained on the training set of ‘Aligned’ data.

View Selection Module

The view selection module first projects the 3D reconstruction error $E = |V - V^{gt}|$ to a series of viewpoints, and




Figure 3: The view selection module: select views covering most errors and then sample based on selected views.

then selects viewpoints containing the most reconstruction errors to be fed into the view synthesis module (see Figure. 3). Specifically, We denote $v = [\psi, \theta, \phi] \in \mathbb{R}^3$ as the viewpoint parameters, where $[\psi, \theta, \phi]$ is the camera rotation along the z-y-x axes using the Tait-Bryan angles, and roll θ is set as zero. Then, our target is to estimate the $v = [\psi, \phi]$ from the 3D reconstruction error E , which indicates the areas with bad reconstruction results of an object. Naturally, if we could provide more information (views) about the areas with large reconstruction errors and try to rectify these errors, the performance of the reconstruction model could be boosted. Particularly, if a camera view can cover the reconstruction error regions more clearly, it is supposed to aid the robustness of the 3D reconstruction model.

To reduce computation, we first divide yaw ψ and pitch ϕ into discrete values with K degree interval to obtain a set of dense discrete 3D viewpoints. The yaw is divided into $\{[-180^\circ + (i-1) \times K^\circ, -180^\circ + i \times K^\circ]\}_{i=1}^{n_\psi}$ and the pitch is divided into $\{[-90^\circ + (i-1) \times K^\circ, -90^\circ + i \times K^\circ]\}_{i=1}^{n_\phi}$, where $n_\psi = 360^\circ/K^\circ$ and $n_\phi = 180^\circ/K^\circ$ refer the number of interval for the yaw ψ and the pitch ϕ , respectively. Then we traverse each interval and use the median value of each interval as the viewpoint position for that interval: $\psi = -180^\circ + (i - \frac{1}{2}) \times K^\circ$ and $\phi = -90^\circ + (i - \frac{1}{2}) \times K^\circ$. As K gets smaller, the traversed viewpoints $v = [\psi, \phi]$ are more continuous. By rotating the voxel grid E by the corresponding degree v , we set the viewpoint direction to the negative direction along the X-axis, denoted as $E_{rot} = Rot(E, v)$, where Rot is the rotation operation.

We then perform the orthographic projection on E_{rot} to each discrete viewpoint to obtain the reconstruction error projection maps P , conducted as follows:

$$P((y, z), E_{rot}) = E_{rot}(x^*, y, z), \quad (1)$$

where the $x^* = \min_{E_{rot}(x, y, z) \neq 0} X$, $E_{rot}(x, y, z)$ represents the integer 3D coordinate and the X represents the range of x . Intuitively, we get the first occupancy value along each line of sight, i.e., the surface of the error voxel.

After we get the error projection map P under each viewpoint v , we obtain the sum of the pixels of each projection map: $sum_v = \sum_{i=1}^S P_i$, where S is the number of a projection map pixels and P_i is the i -th pixel of the projection map P . Then, we select the maximum n views $\{v_s\}$ based on the sum_v . Because larger sum_v means the viewpoint covers more reconstruction error regions. Rather than directly using the selected viewpoints for reconstruction, which will de-




Figure 4: The 3D reconstruction results of Zero-1-to-3, which has weak robustness to view transformations.

crease the diversity of inputting viewpoints, we model each selected viewpoint v_s as a Gaussian distribution and sample views among the distribution: $v \sim \mathcal{N}(v_s, \sigma)$, where σ refers the standard deviation, which is set to $K/6$ in our experiments. Finally, we can obtain the n corresponding viewpoints $v = [\psi, \phi]$ from the 3D reconstruction error E for each object for n -view reconstruction.

Viewpoint pool. During experiments, we found that we construct a viewpoint set for each object under each category, the learned viewpoint distribution will degenerate with the increase of training iterations, that is, the diversity of selected viewpoints will decrease in the later training period, resulting in the phenomenon of overfitting in iterative training. A simple solution is to construct multiple viewpoint sets with multiple iterations for each object of each category. However, it will be very time-consuming. To solve this problem, according to (Wang, Yang, and Tao 2022), it is observed that the selected viewpoint distributions of objects within the same category are highly similar, which has a strong transferability. Based on this, we construct a viewpoint pool by saving the viewpoints selected on each object by category:

$$v = [\{v_{i,1}\}_{i=1}^{n_1}, \dots, \{v_{i,k}\}_{i=1}^{n_k}, \dots, \{v_{i,N}\}_{i=1}^{n_M}], \quad (2)$$

where $v_{i,k}$ refers to the viewpoint selected on object i of category k . We will save the viewpoints corresponding to all n_k objects in the category k and traverse all M categories to construct the viewpoint pool. In the process of iterative optimization, it is only necessary to sample the viewpoint pool according to the category and convert it into the corresponding view. Due to the transferability of the viewpoint distribution of the same type of object, high reconstruction loss can also be induced.

View Synthesis from Stable Diffusion Model

With the camera viewpoint $v = [\psi, \phi]$ selected from the view selection module, we synthesize novel images of an object from the Stable Diffusion model Zero-1-to-3 (Liu et al. 2023b), which is a viewpoint-conditioned diffusion approach and can generate an image under the novel viewpoint by providing an image of the object and the viewpoint transformation matrix. Given a dataset of paired images and their relative camera extrinsic (\hat{I}, I, \hat{v}, v) , the model is trained by solving for the following objective:

$$\min_{\theta} \mathbb{E}_{z \sim \varepsilon(\hat{I}), t, \epsilon \sim \mathcal{N}(0,1)} \| \epsilon - \epsilon_{\theta}(z_t, t, c(\hat{I}, I, \hat{v}, v)) \|_2^2, \quad (3)$$

where ε denote the encoder, ϵ_{θ} is a denoiser U-Net, t is the diffusion time step, and z is the latent representation of




Figure 5: The viewpoint distribution examples of our ShapeNet-VTR dataset. The blue dot is the object center.

the input image encoded by the encoder. After the model is trained, we can generate the novel image from any viewpoint by performing iterative denoising from a Gaussian noise image conditioned on $c(\hat{I}, \hat{v}, v)$.

Use Zero-1-to-3 for direct 3D reconstruction. Directly using Zero-1-to-3 for 3D reconstruction reports bad performance as in Figure 4. It suffers drastically from the view transformations and the 3D reconstruction inference process is very time-consuming (see running time in the Supplemental). Thus, we utilize Zero-1-to-3 for novel view synthesis and rely on the new views to boost the view transformation robustness of existing 3D reconstruction models, instead of using it for direct 3D reconstruction.

Training Method

The training adopts an iterative optimization scheme containing the inner maximization and the outer minimization. Specifically, in the first iteration, we evaluate the pre-trained 3D reconstruction model to select the viewpoints for inner maximization, and then re-train the model on new view images generated by the view synthesis module from the previously selected viewpoints before obtaining a robust model for outer minimization. The process will be repeated until the model converges to the most robust state. Our goal is formulated as: $\min_W \sum_{i=1}^N \mathbb{E}_{v_i} [\max_{v_i} L(W, \mathcal{R}(v_i), y_i)]$, where W denotes the parameters of the reconstruction model f , L is a reconstruction loss function, $\mathcal{R}(v_i)$ is the rendered images of the i -th object given the viewpoint v_i and the y_i is the GT of the corresponding object. We then detail how we improve the inefficiency of the iterative optimization scheme with the random update strategy.

Random update strategy. A random update strategy is adopted to reduce the training time. Instead of generating new views for all objects in each epoch, we randomly select 5% objects and generate new views by Zero-1-to-3 for them, while the remaining 95% objects are still trained with the views from the ‘Align’ training set. At the same time, to make up for the diversity of viewpoints and improve the effectiveness of training, the newly generated view images will be added to the original view images set after each fine-tuning epoch, so that each epoch of iterative training can fully and effectively obtain the prior information left by the previous epoch, and help the 3D reconstruction model learn more efficiently. **Note that only the multi-view 3D reconstruction model is needed at inference.**

	Test set Method	Aligned		Hemispherical		Spherical	
		mIoU \uparrow	F_score \uparrow	mIoU \uparrow	F_score \uparrow	mIoU \uparrow	F_score \uparrow
None-VTR	Pix2Vox++ (Xie et al. 2020)	64.63%	42.02%	47.02%	26.49%	36.71%	18.44%
	GARNet (Zhu et al. 2023c)	65.14%	46.50%	49.69%	30.03%	37.46%	20.41%
VTR	Nature (Pix2Vox++)	65.54%	46.69%	49.00%	30.20%	37.13%	21.03%
	Random (Pix2Vox++)	64.89%	46.67%	50.89%	31.77%	41.62%	24.45%
	VIAT (Ruan et al. 2023) (Pix2Vox++)	65.66%	47.30%	51.50%	31.88%	40.36%	23.47%
	MVTN (Hamdi et al. 2022b) (Pix2Vox++)	65.67%	47.03%	49.58%	30.20%	37.30%	21.36%
	Ours (Pix2Vox++)	65.52%	47.22%	53.53%	33.53%	46.12%	27.01%

Table 1: The view transformation robustness comparison of CNNs-based 3D reconstruction methods. None-VTR/VTR indicates methods without/with view transformation robustness techniques, respectively. All VTR use CNNs-based Pix2Vox++ as the 3D reconstruction model. Our method achieves the best results on ‘Hemispherical’ and ‘Spherical’ with view transformations.

	Test set Method	Aligned		Hemispherical		Spherical	
		mIoU \uparrow	F_score \uparrow	mIoU \uparrow	F_score \uparrow	mIoU \uparrow	F_score \uparrow
None-VTR	UMIFormer (Zhu et al. 2023b)	69.87%	50.56%	53.04%	33.79%	41.19%	23.51%
	LRGT (Yang et al. 2023)	70.01%	51.19%	53.57%	34.71%	41.88%	24.46%
VTR	Nature (LRGT)	70.90%	52.92%	56.06%	35.54%	43.99%	25.16%
	Random (LRGT)	70.92%	52.68%	56.02%	35.28%	44.18%	25.32%
	VIAT (Ruan et al. 2023) (LRGT)	70.86%	52.33%	56.07%	35.29%	44.70%	25.27%
	MVTN (Hamdi et al. 2022b) (LRGT)	71.09%	52.97%	55.76%	34.92%	43.58%	24.07%
	Ours (LRGT)	70.95%	52.78%	57.11%	35.89%	44.97%	26.02%

Table 2: The view transformation robustness comparison of Transformer-based 3D reconstruction methods. All VTR methods use Transformer-based LRGT as the 3D reconstruction model. Our proposed method achieves the best results.

Experiments

Experiment Settings

Dataset generation. Since no existing dataset is suitable for evaluating view transformation robustness, we generate a new dataset *ShapeNet-VTR*, based on ShapeNet (Chang et al. 2015), which consists of larger viewpoint ranges than the original version via rendering new views from the 3D CAD models in ShapeNet. ShapeNet-VTR consists of 13 categories and 39239 objects in total. The objects in each category are divided according to 7:1:2 for training, validation, and testing, respectively. ShapeNet-VTR consists of 3 sets: ‘Aligned’, ‘Hemispherical’, and ‘Spherical’, sharing the same object dividing way but differing in the viewpoint range as shown in Figure 5. ‘Aligned’ set is rendered at a fixed ϕ of 60° and every 15° in ψ ; In the ‘Hemispherical’ set, the range of views is set as $\phi \in [0^\circ, 90^\circ]$, $\psi \in [-180^\circ, 180^\circ]$, and thus images are randomly rendered in the upper hemisphere. In ‘Spherical’, the range of views is $\phi \in [-90^\circ, 90^\circ]$, $\psi \in [-180^\circ, 180^\circ]$, and thus images are randomly rendered in the spherical space. In each set, each object contains 24 views but from different camera angle ranges. We train the 3D reconstruction models on the training set of ‘Aligned’ at first and **use the models’ performance on ‘Hemispherical’ and ‘Spherical’ sets to evaluate the robustness of 3D object reconstruction models to view transformations.**

Implementation details. We utilize two 3D reconstruction models, CNN-based Pix2Vox++ (Xie et al. 2020) and Transformer-based LRGT (Yang et al. 2023), respectively. The input view number is 3 and the image resolution is $256 \times 256 \times 3$ and the size of the voxel output is $32 \times 32 \times 32$. We use a threshold of 0.4 to obtain the occupancy voxel grid and set the interval of degree $K = 30^\circ$ in the experiments.

Comparison methods. We compare with SOTA 3D object reconstruction methods Pix2Vox++ (Xie et al. 2020), GARNet (Zhu et al. 2023c), UMIFormer (Zhu et al. 2023b), and LRGT (Yang et al. 2023) (without VTR approaches, denoted as None-VTR) and four view transformation robustness comparison methods (denoted as VTR). *Note that there are no existing VTR methods for multi-view 3D object reconstruction, and we propose reasonable baselines or adopt methods from other areas for comparison:*

- Nature: Data augmentation with the most common viewpoint renderings from training objects’ natural states (e.g. cars are usually viewed in the side view).
- Random: Data augmentation with random viewpoints, with view range $\phi \in [-90^\circ, 90^\circ]$, $\psi \in [-180^\circ, 180^\circ]$.
- VIAT (Ruan et al. 2023): By regarding viewpoint transformation as an attack, VIAT solves the inner maximization problem to parameterize the Gaussian Mixture distribution of adversarial viewpoints with trainable parameters. In our experiment setting, VIAT keeps the roll θ fixed and learns the parameters of the Gaussian Mixture distribution of yaw ψ and pitch ϕ to get an adversarial viewpoint. Next, we select the other two adversarial viewpoints at azimuth intervals of 120° and 240° , starting from the first adversarial viewpoint.
- MVTN (Hamdi et al. 2022b): A viewpoint selection algorithm that uses differentiable rendering to determine optimal viewpoints for 3D shape recognition through the gradient descent algorithm. We changed its downstream task to multi-view 3D reconstruction. We changed the 3D representation from point clouds to voxels and set the number of the selected viewpoints to 3 as the view amount of the input.

K ($^\circ$)	Aligned	Hemispherical	Spherical
15	65.37%	51.44%	43.27%
30	65.52%	53.53%	46.12%
60	65.40%	52.72%	45.03%
90	65.38%	52.32%	44.26%

Table 3: The ablation study on the interval of degree K . 30° achieves the best mIoU results.

Viewpoint Pool	Aligned	Hemispherical	Spherical
with	65.52%	53.53%	46.12%
without	65.28%	52.07%	42.39%

Table 4: The ablation study on whether the viewpoint pool is used in the proposed method. With the viewpoint pool, the method achieves much better mIoU results.

Besides, we also evaluate the method’s single-view reconstruction robustness, comparing with single-view reconstruction state-of-the-art methods (Mescheder et al. 2019), which does not incorporate prior knowledge of point clouds or depth maps to assist in reconstruction as ours.

Metrics. The evaluation metrics include the mean Intersection-over-Union (mIoU) and F_score (see details in Supp.), whose higher values indicate better performance.

Experiment Results

Multi-view 3D reconstruction performance Table 1 shows the view transformation robustness results of CNNs-based multi-view 3D object reconstruction methods. Our method utilizes Pix2Vox++ (Xie et al. 2020) as the 3D reconstruction model. Compared with SOTA 3D reconstruction methods without view transformation robustness techniques (None-VTR), including Pix2Vox++ (Xie et al. 2020) and GARNet (Zhu et al. 2023c), the proposed method outperforms them a lot on ‘Hemispherical’ and ‘Spherical’ with much larger view transformations than ‘Aligned’. **The proposed method also outperforms all VTR comparisons** (all use Pix2Vox++ as the 3D reconstruction model as ours) on both ‘Hemispherical’ and ‘Spherical’, indicating the proposed method is more robust to view transformation than all comparisons. ‘Nature’ and ‘Random’ try to improve the model robustness via a data augmentation fashion, but also introduce redundant views to the model training, limiting the efficiency and effectiveness. VIAT (Ruan et al. 2023) and MVTN (Hamdi et al. 2022b) are not specially designed for multi-view 3D reconstruction tasks: VIAT ignores the reconstruction error’s spatial relation to the target views and MVTN relies on purely model learning for target view estimation without an in-depth rationale for the view selection step, and thus their performance is reduced. In contrast, the proposed method considers the spatial connections between the 3D reconstruction error and the selected views, where the selected views could cover the reconstruction errors as much as possible to finetune the multi-view 3D reconstruction models on these views in the next circle.

Table 2 shows the VTR of Transformer-based multi-view 3D object reconstruction methods. VTR methods uti-

Test set Method	Aligned		Hemispherical		Spherical	
	mIoU	F_score	mIoU	F_score	mIoU	F_score
OCCNet	52.89%	33.46%	38.31%	21.49%	26.62%	13.74%
Pix2Vox++	56.30%	38.93%	39.12%	23.90%	30.06%	17.39%
LRGT	62.62%	42.63%	45.17%	26.63%	34.70%	19.52%
Ours (Pix2Vox++)	59.47%	40.32%	43.81%	25.84%	36.54%	20.71%
Ours (LRGT)	63.21%	43.35%	47.14%	28.12%	36.13%	20.20%

Table 5: The experiments on single-view 3D object reconstruction methods on the test set with view transformations.

Method	mIoU	F_score
Pix2Vox++	16.24%	7.76%
Ours (Pix2Vox++)	16.45%	7.77%
LRGT	17.08%	7.79%
Ours (LRGT)	17.47%	8.59%

Table 6: The testing results on the Pix3D Chairs dataset.

lize LRGT (Yang et al. 2023) as the 3D reconstruction model. Similarly, the proposed method also outperforms all Transformer SOTAs (Zhu et al. 2023b; Yang et al. 2023) (None-VTR) and view transformation robustness comparisons (VTR) on ‘Hemispherical’ and ‘Spherical’, which proves the advantage of the proposed method and its effectiveness of boosting the robustness of Transformer-based 3D reconstruction methods. The performance gain of our method over comparisons is larger for CNNs-based 3D reconstruction models, due to the weak view fusion ability of CNNs models compared to Transformer 3D reconstruction models. **Overall, from Table 1 and 2, we can conclude the proposed method could boost the robustness of both CNNs-based and Transformer-based 3D reconstruction models, indicating our method’s generalization ability.**

Visualization results. We show the visualization examples of the proposed method and comparisons in Figure 6 (CNNs-based) and Figure 7 (Transformer-based). In each figure, for each object example, the left three columns show the input 3 views, and the middle seven columns indicate the voxel predictions of the methods, and the last column is the ground-truth voxel. The first and the second row of each object example represent results on ‘Hemispherical’ and ‘Spherical’, respectively. In Figure 6 and Figure 7, our method can recover the rough shape of the plane and the chair, while comparisons ignore a lot of details, like the wings of the plane in Figure 6 and the chair legs in Figure 7. We can also observe that the results of the same object of each method on ‘Spherical’ are worse than on ‘Hemispherical’, because ‘Spherical’ contains larger view transformations and is more difficult. Transformer-based methods can achieve better visualization results than CNN-based methods because Transformer methods have stronger inter-view fusion ability. No matter whether using which 3D reconstruction model or on which testing set, the proposed method always achieves the best results, which indicates the stronger robustness of the proposed method to view transformations.

Ablation study Ablation study on the interval of degree K . Table 3 presents the reconstruction result of the model

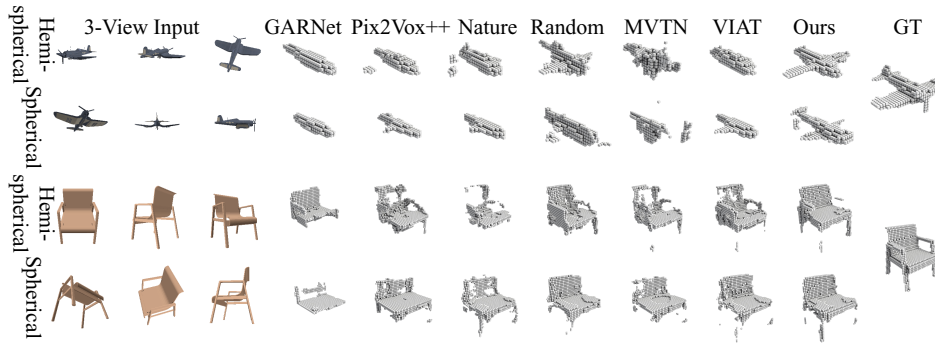


Figure 6: The visualization results of CNNs-based methods, and we use Pix2Vox++ as the 3D reconstruction model.

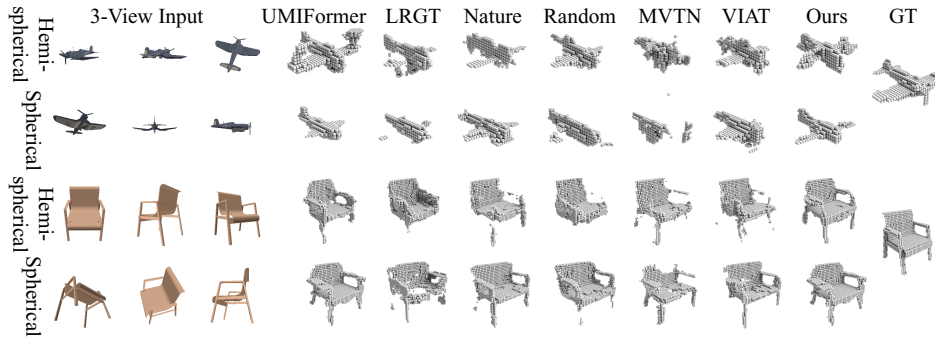


Figure 7: The visualization results of Transformer-based methods, and we use LRGT as the 3D reconstruction model.

after iterative training with different intervals of degree K in the view selection, where the best results are achieved by $K = 30^\circ$. It shows a suitable interval of degree benefits the model in achieving better results. If the degree is too large, the viewpoint selection will tend to be random; while if it is too small, it will affect the viewpoint diversity and hinder the reconstruction robustness. Thus, we experiment with $K = 30^\circ$ in our methods.

Ablation study on the viewpoint pool. We conduct the ablation study on the effectiveness of the viewpoint pool in Table 4. It shows that the method with the viewpoint pool can achieve better performance than without the viewpoint pool on all testing sets, especially on the ‘Spherical’ set with much larger view transformations. This indicates the effectiveness of the viewpoint pool in boosting the multi-view 3D reconstruction model’s robustness to view transformations.

Single-view 3D reconstruction performance. We also validate the proposed method’s robustness in single-view 3D reconstruction and compare it with single-view SOTA methods (Mescheder et al. 2019; Xie et al. 2020; Yang et al. 2023) in Table 5. It shows that our method achieves better results than all comparisons on all sets, no matter whether using Pix2Vox++ (Xie et al. 2020) or LRGT (Yang et al. 2023) as 3D reconstruction models, which proves the effectiveness of the proposed method for single-view 3D reconstruction.

Evaluation on the real-world Dataset Pix3D. We also test the proposed method on Pix3D (Sun et al. 2018) to verify its single-view reconstruction performance on real-world

data with more complicated object viewpoint distributions. Following (Xie et al. 2020), we use the data from the Chair category in ShapeNet to generate a training set and render images with random backgrounds in the SUN dataset (Xiao et al. 2010), and each object has 60 synthesized images. Table 6 shows that our method’s performance is better than SOTA methods, suggesting that our method helps the 3D reconstruction model generalize better to real-world datasets by improving the model’s view transformation robustness.

Conclusion

In this paper, we propose a novel reconstruction error-guided view selection method together with view synthesis via a Stable Diffusion model to improve the view transformation robustness of existing multi-view 3D object reconstruction methods. Instead of randomly synthesizing new views from Stable Diffusion models and adding them to model training as a data augmentation approach, we consider the spatial distributions of the 3D reconstruction errors and use them to guide the view selection for choosing the most effective views covering the errors as most as possible. Our proposed method shows the best view transformation robustness compared to the latest multi-view 3D reconstruction SOTAs and various view transformation robustness comparison methods. We provide a new perspective on incorporating large vision models into existing relatively ‘small’ 3D object reconstruction models for robustness gains without increasing the model deploying cost.

Acknowledgements

This work was supported in parts by NSFC (62202312, U21B2023), Guangdong Basic and Applied Basic Research Foundation (2023B1515120026), Shenzhen Science and Technology Program (KQTD 20210811090044003, RCJC20200714114435012), and Scientific Development Funds from Shenzhen University.

References

- Arshad, M. S.; and Beksi, W. J. 2023. LIST: Learning Implicitly from Spatial Transformers for Single-View 3D Reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9321–9330.
- Bauer, D.; Hönig, P.; Weibel, J.-B.; García-Rodríguez, J.; Vincze, M.; et al. 2024. Challenges for Monocular 6D Object Pose Estimation in Robotics. *IEEE Transactions on Robotics*.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*.
- Choy, C. B.; Xu, D.; Gwak, J.; Chen, K.; and Savarese, S. 2016. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, 628–644. Springer.
- Cohen, T.; et al. 2021. *Equivariant convolutional networks*. Ph.D. thesis, Taco Cohen.
- Dong, Y.; Ruan, S.; Su, H.; Kang, C.; Wei, X.; and Zhu, J. 2022. ViewFool: Evaluating the Robustness of Visual Recognition to Adversarial Viewpoints. *36th Conference on Neural Information Processing Systems (NeurIPS)*, abs/2210.03895.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Hamdi, A.; AlZahrani, F.; Giancola, S.; and Ghanem, B. 2022a. MVTN: Learning Multi-View Transformations for 3D Understanding. *ArXiv*, abs/2212.13462.
- Hamdi, A.; AlZahrani, F.; Giancola, S.; and Ghanem, B. 2022b. MVTN: Learning multi-view transformations for 3D understanding. *arXiv preprint arXiv:2212.13462*.
- Hu, J.; Hui, K.-H.; Liu, Z.; Li, R.; and Fu, C.-W. 2024. Neural wavelet-domain diffusion for 3d shape generation, inversion, and manipulation. *ACM Transactions on Graphics*, 43(2): 1–18.
- Huang, P.-H.; Matzen, K.; Kopf, J.; Ahuja, N.; and Huang, J.-B. 2018. Deepmvs: Learning multi-view stereopsis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2821–2830.
- Kar, A.; Häne, C.; and Malik, J. 2017. Learning a multi-view stereo machine. *Advances in neural information processing systems*, 30.
- Kwak, J.-g.; Dong, E.; Jin, Y.; Ko, H.; Mahajan, S.; and Yi, K. M. 2024. Vivid-1-to-3: Novel view synthesis with video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6775–6785.
- Liu, M.; Shi, R.; Chen, L.; Zhang, Z.; Xu, C.; Wei, X.; Chen, H.; Zeng, C.; Gu, J.; and Su, H. 2023a. One-2-3-45++: Fast Single Image to 3D Objects with Consistent Multi-View Generation and 3D Diffusion. *arXiv preprint arXiv:2311.07885*.
- Liu, M.; Xu, C.; Jin, H.; Chen, L.; Varma T, M.; Xu, Z.; and Su, H. 2024. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems*, 36.
- Liu, R.; Wu, R.; Hoorick, B. V.; Tokmakov, P.; Zakharov, S.; and Vondrick, C. 2023b. Zero-1-to-3: Zero-shot One Image to 3D Object. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9264–9275.
- Liu, Y.; Lin, C.; Zeng, Z.; Long, X.; Liu, L.; Komura, T.; and Wang, W. 2023c. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*.
- Mescheder, L.; Oechsle, M.; Niemeyer, M.; Nowozin, S.; and Geiger, A. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4460–4470.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, 565–571. Ieee.
- Paschalidou, D.; Ulusoy, O.; Schmitt, C.; Van Gool, L.; and Geiger, A. 2018. Raynet: Learning volumetric 3d reconstruction with ray potentials. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3897–3906.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Qi, Z.; Yu, M.; Dong, R.; and Ma, K. 2024. Vpp: Efficient conditional 3d generation via voxel-point progressive representation. *Advances in Neural Information Processing Systems*, 36.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ruan, S.; Dong, Y.; Su, H.; Peng, J.; Chen, N.; and Wei, X. 2023. Towards Viewpoint-Invariant Visual Recognition via Adversarial Training. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 4686–4696.
- Shen, W.; Zhang, B.; Huang, S.; Wei, Z.; and Zhang, Q. 2020. 3d-rotation-equivariant quaternion neural networks.

- In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, 531–547. Springer.
- Shi, R.; Chen, H.; Zhang, Z.; Liu, M.; Xu, C.; Wei, X.; Chen, L.; Zeng, C.; and Su, H. 2023. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*.
- Shi, Z.; Meng, Z.; Xing, Y.; Ma, Y.; and Wattenhofer, R. 2021. 3d-retr: End-to-end single and multi-view 3d reconstruction with transformers. *arXiv preprint arXiv:2110.08861*.
- Spezialetti, R.; Salti, S.; and Stefano, L. D. 2019. Learning an effective equivariant 3d descriptor without supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6401–6410.
- Su, H.; Maji, S.; Kalogerakis, E.; and Learned-Miller, E. 2015. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, 945–953.
- Sun, X.; Wu, J.; Zhang, X.; Zhang, Z.; Zhang, C.; Xue, T.; Tenenbaum, J. B.; and Freeman, W. T. 2018. Pix3d: Dataset and methods for single-image 3d shape modeling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2974–2983.
- Tiong, L. C. O.; Sigmund, D.; and Teoh, A. B. J. 2022. 3D-C2FT: Coarse-to-fine Transformer for Multi-view 3D Reconstruction. In *Proceedings of the Asian Conference on Computer Vision*, 1438–1454.
- Wang, D.; Cui, X.; Chen, X.; Zou, Z.; Shi, T.; Salcudean, S.; Wang, Z. J.; and Ward, R. 2021. Multi-view 3d reconstruction with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5722–5731.
- Wang, R.; Yang, Y.; and Tao, D. 2022. ART-Point: Improving Rotation Robustness of Point Cloud Classifiers via Adversarial Rotation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14351–14360.
- Wei, X.; Yu, R.; and Sun, J. 2020. View-gcn: View-based graph convolutional network for 3d shape analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1850–1859.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, 3485–3492. IEEE.
- Xie, H.; Yao, H.; Sun, X.; Zhou, S.; and Zhang, S. 2019. Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2690–2698.
- Xie, H.; Yao, H.; Zhang, S.; Zhou, S.; and Sun, W. 2020. Pix2Vox++: Multi-scale context-aware 3D object reconstruction from single and multiple images. *International Journal of Computer Vision*, 128(12): 2919–2935.
- Yagubbayli, F.; Wang, Y.; Tonioni, A.; and Tombari, F. 2021. Legoforner: Transformers for block-by-block multi-view 3d reconstruction. *arXiv preprint arXiv:2106.12102*.
- Yang, B.; Wang, S.; Markham, A.; and Trigoni, N. 2020. Robust attentional aggregation of deep feature sets for multi-view 3D reconstruction. *International Journal of Computer Vision*, 128(1): 53–73.
- Yang, L.; Zhu, Z.; Lin, X.; Nong, J.; and Liang, Y. 2023. Long-Range Grouping Transformer for Multi-View 3D Reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18257–18267.
- Yao, Y.; Luo, Z.; Li, S.; Fang, T.; and Quan, L. 2018. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, 767–783.
- Yao, Y.; Luo, Z.; Li, S.; Shen, T.; Fang, T.; and Quan, L. 2019. Recurrent mvsnet for high-resolution multi-view stereo depth inference. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5525–5534.
- Yoo, P.; Guo, J.; Matsuo, Y.; and Gu, S. S. 2024. Dreamspare: Escaping from plato’s cave with 2d diffusion model given sparse views. *Advances in Neural Information Processing Systems*, 36.
- Zhu, B.; Yang, C.; Dai, J.; Fan, J.; Qin, Y.; and Ye, Y. 2023a. R 2 FD 2: Fast and robust matching of multimodal remote sensing images via repeatable feature detector and rotation-invariant feature descriptor. *IEEE Transactions on Geoscience and Remote Sensing*.
- Zhu, Z.; Yang, L.; Li, N.; Jiang, C.; and Liang, Y. 2023b. Umiformer: Mining the correlations between similar tokens for multi-view 3d reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18226–18235.
- Zhu, Z.; Yang, L.; Lin, X.; Yang, L.; and Liang, Y. 2023c. Garnet: Global-aware multi-view 3d reconstruction network and the cost-performance tradeoff. *Pattern Recognition*, 142: 109674.
- Zou, Z.; Cheng, W.; Cao, Y.-P.; Huang, S.-S.; Shan, Y.; and Zhang, S.-H. 2024. Sparse3d: Distilling multiview-consistent diffusion for object reconstruction from sparse views. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 7900–7908.