

IC3D: Image-Conditioned 3D Diffusion for Shape Generation

Cristian Sbrilli Paolo Cudrano Matteo Frosi Matteo Matteucci

Department of Electronics Information and Bioengineering

Politecnico di Milano, Italy

{name.surname}@polimi.it

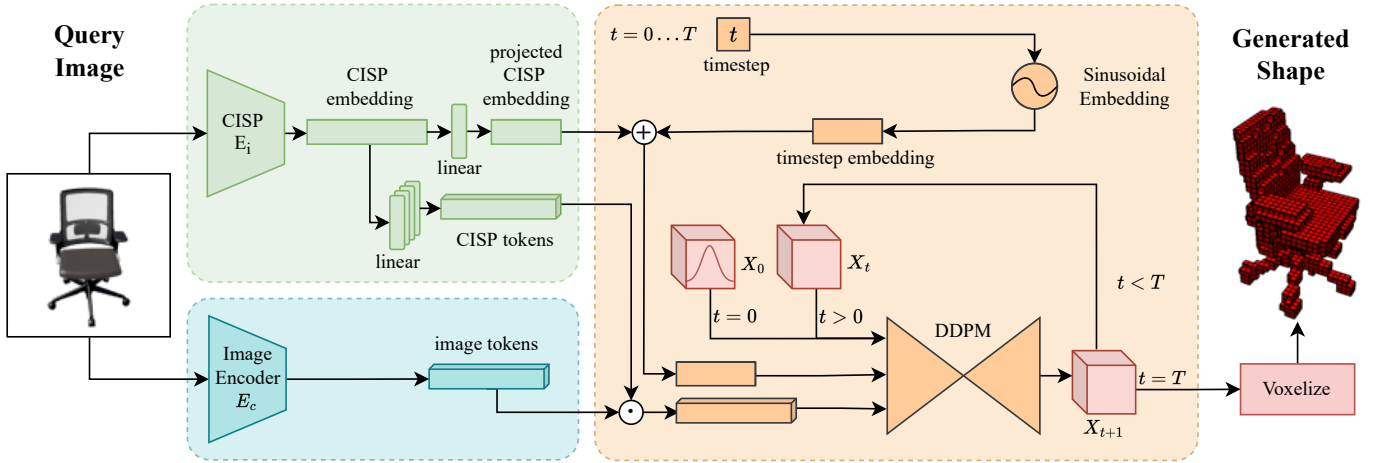


Fig. 1: IC3D, our 3D shape generation pipeline. A joint image-shape embedding space is preemptively learned through our pre-training model CISP (Contrastive Image-Shape Pre-training). Given a query image, we generate its CISP embeddings and use them, together with additional context, to condition a 3D Denoising Diffusion Probabilistic Model (DDPM). In this fashion, IC3D can generate diverse 3D shapes, all maintaining both realism and coherence with the query image./

Abstract—In recent years, Denoising Diffusion Probabilistic Models (DDPMs) have demonstrated exceptional performance in various 2D generative tasks. Following this success, DDPMs have been extended to 3D shape generation, surpassing previous methodologies in this domain. While many of these models are unconditional, some have explored the potential of using guidance from different modalities. In particular, image guidance for 3D generation has been explored through the utilization of CLIP embeddings. However, these embeddings are designed to align images and text, and do not necessarily capture the specific details needed for shape generation. To address this limitation and enhance image-guided 3D DDPMs with augmented 3D understanding, we introduce CISP (Contrastive Image-Shape Pre-training), obtaining a well-structured image-shape joint embedding space. Building upon CISP, we then introduce IC3D, a DDPM that harnesses CISP’s guidance for 3D shape generation from single-view images. This generative diffusion model outperforms existing benchmarks in both quality and diversity of generated 3D shapes. Moreover, despite IC3D’s generative nature, its generated shapes are preferred by human evaluators over a competitive single-view 3D reconstruction model. These properties contribute to a coherent embedding space, enabling latent interpolation and conditioned generation also from out-of-distribution images. We find IC3D able to generate coherent and diverse completions also when presented with occluded views, rendering it applicable in controlled real-world scenarios.

I. INTRODUCTION

Recent years have witnessed astonishing advances in deep generative models, especially in image generation. Denoising Diffusion Probabilistic Models (DDPMs) played a central role in this progress, outperforming previous methods such as variational autoencoders (VAEs) and generative adversarial networks (GANs) in unconditional image synthesis [1] and text-to-image synthesis. Works such as GLIDE [2], DALLE-2 [3], Stable Diffusion [4], and Imagen [5] showed, indeed, how we can effectively condition the generation of images by text prompts.

Transitioning from 2D image generation to 3D shape generation presents challenges, although recent years have witnessed notable progress in this field. Following the initial wave of VAEs and GAN architectures [6], [7], a subsequent wave of models has focused on flow-based and energy-based paradigms [8], [9], [10], showcasing their efficacy in generating high-quality and diverse shapes. Recent studies have not only demonstrated DDPMs’ superiority in shape quality and diversity compared to previous approaches, but have also showcased the potential to condition the generation process using various modalities, such as text prompts [11], shape

latents from autoencoders [12], and single-view images [13].

In particular, the latter—image-to-shape generation—, is pivotal for driving innovation across diverse fields like VR and AR, cultural heritage, medical imaging and diagnosis, and industrial design and manufacturing. Image-conditioned 3D generation has been experimented [13] by leveraging CLIP [14] image embeddings as guidance tokens. However, the inherent alignment of text and images within CLIP embeddings may not optimally suit guiding an image-to-shape diffusion process, potentially leading to loss of fine-grained details crucial for shape generation.

Motivated by this insight, in this work we propose the construction of a CLIP-inspired joint image-shape embedding space, named CISP (Contrastive Image-Shape Pre-training). We show that this embedding space not only exhibits a well-organized structure but also captures the inherent semantic structural characteristics of the represented objects. We exploit this regularity to condition a 3D diffusion model through classifier-free guidance [15], leveraging CISP to steer the generated shapes toward the desired image context. As our aim is to investigate the advantages of image-shape embeddings for 3D generation, we focus on voxels, to avoid additional architectural complexities due to other representations such as point clouds. We find that IC3D generates image-coherent yet diverse shape samples from single-view images, and it achieves cutting edge results for both quality and diversity also in unconditional generation. Thanks to the structural properties acquired via CISP guidance, the shapes modeled by IC3D are preferred by human assessors in terms of quality and coherence also when compared to a single-view 3D reconstruction model, despite IC3D being a generative model. We show the regularity of CISP embeddings through manifold interpolation and OOD probing experiments. Furthermore, we test IC3D’s resilience to occluded views and demonstrate its potential for automated deployment in controlled real-world scenarios.

II. RELATED WORKS

Various strategies have been put forth to develop generative models with the ability to synthesize three-dimensional (3D) objects. Notable methodologies encompass early Variational Autoencoders (VAEs) [16] and Generative Adversarial Networks (GANs) [17] applied to 3D voxelized shapes, as well as the more recent SDF-StyleGAN [18], which constructs a generative model based on an SDF-based implicit representation. The research in this domain has also seen a growing interest in generative techniques like flow-based models [19], Energy-Based Models (EBMs) [20], and Denoising Diffusion Probabilistic Models (DDPMs) [21], [22]. Flow-based models, such as DPF-Net [23], manipulate probability distributions to generate samples through variable transformations. EBMs, exemplified by 3D DescriptorNet [10], optimize energy functions over observed variables and generate new data via Langevin Dynamics [24], [25]. On the other hand, DDPMs train using a denoising objective and subsequently invert the process to generate samples from noise.

Point-Voxel Diffusion (PVD) [26] employs point-voxel CNN [27] to generate point cloud shapes, also highlighting the

inability of training a voxel-based DDPM. Additionally, PVD necessitates separate training for different shape categories due to its unconditional nature. A contrasting approach presented by Luo and Hu [12] involves a point-cloud DDPM conditioned on shape latents extracted from a point cloud autoencoder, enabling a single model to generate diverse object categories. Hui *et al.* [28] employ diffusion on SDFs wavelet coefficients to generate coarse volumes and a refiner network to predict details. Leveraging latent diffusion, LION [13] demonstrates image-conditioning of a 3D generation model using CLIP [14] text or image embeddings. However, although well-structured, CLIP embeddings inherently lack 3D features, and their use in 3D generation might lead to loss of fine-grained structural shape details that are not captured by images and text. On the contrary, joint image-shape embeddings would overcome this limitation.

Joint image-shape embeddings have been investigated by Li *et al.* [29], establishing a joint image-shape embedding space through a multistep process. It begins by constructing an embedding space for shapes based on shape similarities and then learns to associate images with their corresponding shape embeddings. In contrast, Kuo *et al.* [30] adopt a holistic approach, jointly learning image and 3D CAD shape embeddings. This concept is further extended in [31], learning an image patch-CAD shape mapping. Imagebind [32] shows instead how we can align multiple modalities to images with a contrastive approach. However, they only focus on depth maps, which do not fully capture the 3D scene and may thus lose critical 3D information. On the other hand, we propose to retain full 3D information by extending the contrastive methodology in [14] directly to image-shape domain, with voxels as 3D representation.

III. CONTRASTIVE IMAGE-SHAPE PRE-TRAINING

To condition the 3D shape generation on images, we require a meaningful and well-structured concept space jointly encoding both 2D and 3D information. We obtain this space with CISP (Contrastive Image-Shape Pre-training), which learns to embed images and shapes in a joint space. The CISP pre-training method follows the procedure described in CLIP [14]. In particular, it uses a contrastive approach to learn meaningful representations of images and shapes in a joint space. We define an encoder E_s processing shapes and an encoder E_i processing images, both producing embeddings of size f . Given a batch containing N corresponding (*image*, *shape*) pairs, our training objective is to match each image with its shape. We produce the batches of image and shape embeddings e_i and e_s using E_i and E_s . We then compute a symmetric $N \times N$ similarity matrix between image embeddings (rows) and shape embeddings (columns), using a cosine similarity measure. Our training loss is composed of two cross-entropy terms, measuring both the ability to predict the correct shape given an image and vice-versa. This objective aims at maximizing the similarity of the N matching (*image*, *shape*) pairs and, at the same time, minimizing the similarity of the $N^2 - N$ unmatched pairs. For the implementation details, we refer the reader to [14].

In CLIP, Radford *et al.* [14] adopt a Vision Transformer (ViT) [33] as image encoder. They found it to produce better results and higher computational efficiency compared to CNN architectures when trained on a sufficiently large dataset. As the dataset we use (Sec. V) is limited, however, we employ a Data-efficient image Transformer (DeiT) [34], i.e., DeiT Base (DeiT-B), using 768-dimensional hidden embeddings and 12 layers with 12 attention heads each.

As shape encoder, we extend the ViT model to 3D. We replace the 2D convolutions, mapping an image to patch embeddings, with 3D convolutions mapping a voxel shape to patch embeddings. A learnable token is prepended to the input sequence, processed by the network, and then projected to the desired embedding dimension. This technique allows us to easily extract the image (or shape) CISP embedding, and is inspired by class tokens first used in NLP [35] and later introduced in ViT [33].

IV. SHAPE GENERATION PIPELINE

Figure 1 shows our image-driven shape generation pipeline, which is inspired by 2D text-to-image works [3], [2]. From the input query image we obtain its CISP embedding, carrying information about its 3D shape, and some additional image tokens, encoding information not captured by CISP. We leverage these embeddings to guide our DDPM generative process. As DDPMs are iterative methods, they require information about the current timestep, which is encoded by sinusoidal embeddings [36]. Our DDPM module iteratively refines an input 3D tensor, which is finally transformed into the output shape by binary thresholding. We report a detailed description of DDPMs in the supplementary material.

A. Architecture

We extend the architecture used in GLIDE [2], which is based on the ADM model [1]. Specifically: (1) we replace 2D convolutions with 3D convolutions; (2) we replace the text encoder with an image encoder E_c , which we chose to be a DeiT [34] model; (3) as in [3], we use CISP image embeddings in two ways: first, we project and add them to the timestep embedding; second, in each attention block of the network, we project CISP embeddings into four extra tokens and we concatenate them to the attention context (keys, values); (4) we prepend 8 learnable tokens to E_c and use the corresponding outputs as additional attention context, as with CISP embeddings.

B. Guidance

We leverage CISP to train a single DDPM on all the considered object categories, guiding the generation by using CISP image embeddings. To this end, we train our model with classifier-free guidance [15], which eliminates the need for a separate classifier and obtains similar results compared to classifier guidance [22]. We report the mathematical formulation of classifier-free guidance and guidance scale analysis in the supplementary material.

Input Image

Generated Shapes




Fig. 2: 3D shapes generated by IC3D from the query images on the left. Notice how each shape is diverse and yet it coherently represents structural elements of the query (e.g., flat-topped tail of the airplane, cargo space of the pickup truck, wheels and armrests of the chair.)

V. EXPERIMENTS

We run experiments to evaluate our generative pipeline quantitatively and qualitatively. We compare our model against state-of-the-art 3D generative approaches to quantitatively evaluate our generated shapes, following evaluation approaches and metrics defined by the literature. Furthermore, we evaluate the effectiveness of our guidance approach and the visual quality of generated shapes through a human evaluation experiment. Lastly, we investigate the representation and generalization capabilities of CISP and IC3D through additional experiments and ablations.

Following previous shape generation works [26], [12], [9], [8], [13], we focus on the Airplane, Car, and Chair categories from the ShapeNet [37] test set. Although not used in previous works, we also train our model on the Table and Watercraft categories; see the supplementary material for the results of these two additional categories. Notice that, except for the unconditional experiments, we refer to IC3D as a single model trained on all categories.

A. Shape Generation

The generation of shapes is performed using the inverse diffusion process. In particular, we first generate a pure noise sample of dimension 32^3 , and then run 1000 backward diffusion steps. We obtain a 32^3 float tensor which we voxelize through binary thresholding (threshold at 0.5). For conditional generation, we apply classifier-free guidance as explained in Sec. IV-B, with a guidance factor of 1.5, as we found it produces the best results. Figure 2 shows examples of image-guided generation from our model.

We evaluate the quality and diversity of shapes generated by our model by comparing against other 3D generative works [38], [23], [28], [12], [26], [13] and we report results in Tab. I. Since these models only evaluate unconditional generation, we also configure our pipeline to match this setting for a fair comparison. To achieve unconditional generation with our model, we conduct experiments on each category without guiding images, substituting CISP embeddings and E_c tokens with learned null tokens (only at test time). We

Shape	Model	1-NNA(%)	
		CD	EMD
Airplane	Shape-GF [38]	80.00	76.17
	DPF-Net [23]	75.18	65.55
	[28]	71.69	66.74
	[12]	62.71	67.14
	PVD [26]	73.82	64.81
	LION [13]	67.41	61.23
	Ours	58.93	56.93
Car	Shape-GF [38]	63.20	56.53
	DPF-Net [23]	62.35	54.48
	[28]	-	-
	[12]	-	-
	PVD [26]	54.55	53.83
	LION [13]	53.70	52.34
	Ours	53.20	53.11
Chair	Shape-GF [38]	68.96	65.48
	DPF-Net [23]	62.00	58.53
	[28]	61.47	61.62
	[12]	62.08	64.45
	PVD [26]	56.26	53.32
	LION [13]	53.41	51.14
	Ours	53.30	51.97

TABLE I: Comparison of unconditional generation performance of IC3D against baselines and SoTA 3D generative models. Best results are highlighted in bold. 1-NNA perfect score is 50%. The car category was not evaluated in [12] and [28].

compare mainly on 1-Nearest Neighbor Accuracy (1-NNA), as it was shown to be the most representative metric for 3D generation [9]. We report in the supplementary material results for other common generation metrics, i.e., Minimum Matching Distance (MMD), Coverage (COV) and Shading-image-based FID [18], where IC3D also obtains competitive results. For each metric, we consider both Chamfer distance (CD) and Earth Mover’s Distance (EMD).

1-NNA, presented in [39] and adopted for 3D generation in [9], measures both generated samples diversity and quality by considering the accuracy of a 1-Nearest Neighbor classifier. We define S_g as the set of generated samples and S_r as the set of reference samples (which we take from the test set) with $|S_r| = |S_g|$. Also, let N_x be the nearest neighbor of a sample x , with $N_x \in \{S_g \cup S_r - x\}$. 1-NNA is then defined as

$$1\text{-NNA}(S_g, S_r) = \frac{\sum_{x \in S_g} \mathbb{I}[N_x \in S_g] + \sum_{x \in S_r} \mathbb{I}[N_x \in S_r]}{|S_g| + |S_r|}, \quad (1)$$

with \mathbb{I} being the indicator function. As we are measuring the accuracy of a 1-NN classifier on recognizing generated and reference shape, a value of 50% is the perfect score.

To make voxel shapes compatible with EMD and CD, we follow the approach from PVD [26], sampling 2048 points from the surface of the generated shape. We also adopt metrics implementations from PVD [26] public code. Our method outperforms or is on par with other models in terms of quality and diversity, achieving higher results on the airplane class and comparable results with LION [13] on cars and chairs. As LION [13] also allows for image guidance, but only evaluates on unconditional generation, we further compare these two models qualitatively in Fig. 3. Thanks to CISP guidance, our




Fig. 3: Thanks to CISP guidance, our model is able to match finer-grained details of the guiding image with respect to LION [13], which uses CLIP. For our model, we show generated voxels and the corresponding sampled point cloud.

model is better able to generate shapes that are coherent with the input image in most details, as we can observe with the plane back engines and the chair circular hole.

We further analyze the quality of our results running a **side-by-side human evaluation**, comparing our model with a SoTA 3D reconstruction approach. We evaluate how well our generated shapes match query images, as well as their realism. Each human evaluator is presented with two unlabeled and randomly ordered 3D shapes, one from our model and one from a reconstruction model. The evaluator is asked to choose which of the two shapes is more realistic. After its choice, the query image is shown, and the evaluator has to choose which shape is more coherent with the displayed image. We chose to only show the query image after the first question is answered to avoid biasing the evaluator when selecting the most realistic shape. Shapes are displayed as GIFs rotating 360° around their volume, allowing the evaluator to see each shape from different perspectives. Following the method from [40], we show each pair of images exactly to 5 independent annotators, allowing us to reduce the variance on each pair by majority voting.

We compare against 3D-RETR [41], the best-performing publicly available model for single-view 3D reconstruction at the time of the experiment. We select 200 random images per category (airplane, car, chair) and use them to condition our model and to predict the shapes using 3D-RETR, obtaining the pair of shapes needed for the evaluation. We publicly release the images, GIFs, and necessary setup for reproducing this evaluation experiment.

The results obtained (Fig. 4) show that our model is preferred by human evaluators in both coherence and realism. In particular, the majority of the evaluators considered our generated shapes to be the most coherent with the query image in 62% of the cases, and the most realistic in 69% of the cases. Examining the per-category scores (reported fully in the supplementary material), chairs are the most preferred category for our model in both coherence (68%) and realism (91%). We speculate that our model achieves better performance because its generative nature forces it to capture the semantics of each structural detail, while for a reconstruction method any voxel has an equivalent meaning, as further discussed in Sec. V-H.




Fig. 4: Human evaluation of coherence and realism of IC3D generated shapes against shapes reconstructed by 3D-RETR [41]. IC3D shapes were considered more coherent with the input image 63% of the times and more realistic 69% of the times. Among the unanimous decisions (5/5 or 0/5 votes), 72% were in favor of IC3D for its coherence and 85% for its realism.




Fig. 5: Generation from occluded views. Our model correctly generates plausible and realistic shapes even for unseen parts.

B. Occluded and partial image views

We study the generative power of IC3D when the conditioning image presents occlusions or partial views. In these situations several completions for the occluded parts are possible. The generative nature of IC3D allows it to produce various shapes, each coherent with the given query image, but presenting diversity in the unseen parts. This is evident in Fig. 5, where only backs of chairs are shown to our model. Despite the missing information, IC3D is capable of (1) identifying the type of object it is presented with (as chairs), and (2) generating chairs with plausible but diverse front parts.

C. Latent space analysis

To gain insight into how CISP embeddings encode meaningful and effective representations, we analyze the learned latent space. We project CISP embeddings by first applying PCA, reducing them to 50 dimensions, and then using t-SNE to project into the 2D plane. Figure 6 shows the projected shape embeddings of the training set, in which we observe that the categories are well separated. The figure also shows example images depicting how shapes are organized in the

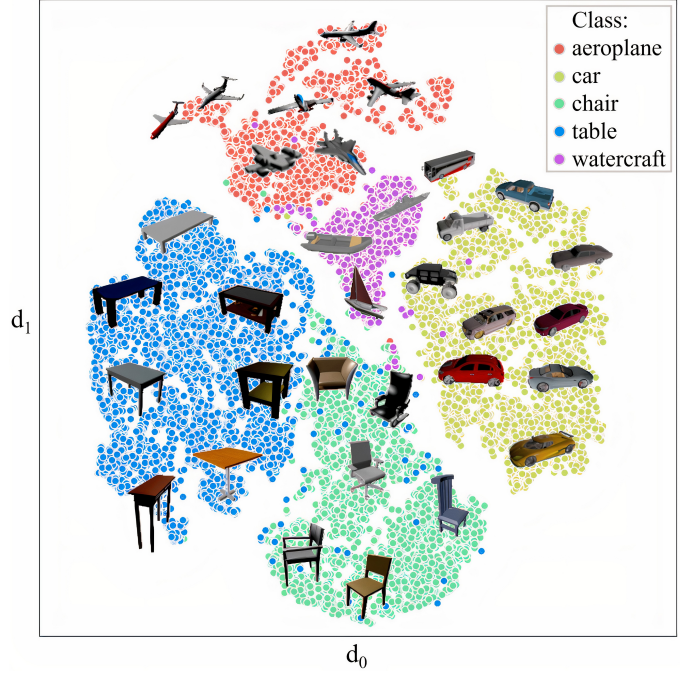


Fig. 6: Projection of the CISP shape embeddings into a 2D space, displaying instances of the encoded images for each considered class at their respective locations.

embedding space. We observe how the model captures details and subcategories, allowing it to better structure the space. For example, notice how the height of the tables increases as d_1 decreases or how shelves begin to appear as d_0 increases. At the same time, for cars, notice how sports cars fade into city cars and then into progressively larger cars as d_1 increases. We also see inter-class interactions, e.g., the small cloud of buses lies close to longest watercrafts since they share a similar form factor when voxelized. This well-structured joint space makes CISP the core of our guidance method, allowing us to generate shapes coherent with input images.

D. Sketch-to-Shape generation

To additionally prove the effectiveness and generalization capabilities of the CISP model, we test the generation of 3D shapes by guiding our DDPM with CISP embeddings of hand-drawn sketches. Examples of sketch-to-shape generation are shown in Fig. 7. CISP embeddings correctly encapsulate shape categories and even more specific intraclass details from the hand-drawn sketches.

E. Interpolations

We show that it is possible to traverse the concept space between two objects by interpolating between their CISP embeddings. As in [3], we do so by moving between the two considered CISP embeddings with Slerp (spherical linear interpolation). We found Slerp to work much better than standard linear interpolation, which failed to effectively blend the two input shapes. Figure 8 shows some examples of interpolated shapes obtained with our model. We show that both intra- and inter-class interpolations are effective, further




Fig. 7: The generalization power of CISP allows us to generate coherent shapes even when feeding IC3D with a hand-drawn sketch of the desired object. Notice how the generated samples accurately capture characteristic aspects of the drawn items, such as the type of chair bottom (wheels, single leg, four legs), while still providing a degree of diversity.

proving the relevance of the knowledge captured by the CISP embedding space. Notice in the first row how the wheels and armrests are added to the original chair structure before turning it into an office chair. Also notice how, in an inter-class scenario (bottom row), the chair is first filled, then the base is removed, and finally the legs are added, completing its transformation into a table.

F. Ablation study

To better understand the contribution of different elements in our model, we perform an ablation study on each of them and report the main results in Tab. II. We evaluate the generation performance of IC3D when E_c is completely removed and when guidance tokens from CISP and/or E_c are substituted, at inference time, with a null token. For configurations using guidance images (when not using null tokens), we use images from a different set than S_r . Unlike text-to-image models [3], where the extra text encoder has a negligible role, our model suffers a significant drop in performance without E_c . Specifically, in this case, the generated shapes retain their

Model guidance		1-NNA(%)		COV(%)↑	
CISP	E_c	CD	EMD	CD	EMD
Yes	No	68.71	62.37	51.49	52.35
Null token	Null token	55.14	54.00	51.99	52.70
Yes	Null token	55.21	52.90	51.14	52.59
Yes	Yes	54.55	52.43	50.72	52.33

TABLE II: We analyze the impact of different guidance elements in IC3D, testing whether removing completely a guidance encoder or substituting it with the null token affects the generative results. Scores are averaged over the three considered categories. We notice that while coverage remains similar in all cases, 1-NNA score degrades when E_c is removed.

diversity, but their visual quality deteriorates considerably. On the other hand, using null tokens has a minor effect on generation quality, while achieving the highest coverage scores. This is expected, as the model is freely generating shapes. These results indicate that CISP and E_c not only offer effective guidance, but also help to learn to generate higher quality samples during training. Moreover, we notice that the guidance from CISP and E_c enables a tradeoff between quality (1-NNA) and diversity (COV).

G. Generation from in-the-wild images

Although our model is trained only on synthetic data (ShapeNet), we study its ability to generate shapes from in-the-wild images. To perform such analysis in an exhaustive way, we provide IC3D with images coming from three sources, each more in-the-wild: (1) an online chair catalog, (2) a known real-world image-shape dataset (Pix3D [48]), and (3) pictures of an office chair captured by the authors with a smartphone. As the samples in our training dataset, ShapeNet, are rendered only on white backgrounds, the input images to our model are pre-processed with a simple and completely automatic background removal tool. Figure 9 highlights how our model is capable of generating coherent 3D shapes even in these situations, demonstrating potential for automatic use in controlled environments, such as industrial settings. As already seen in Fig. 5, also in this scenario IC3D introduces a level of coherent variability to the generated shapes when occlusion and partial image views are given, making it even more appealing for applications in graphics and design.

H. 3D Reconstruction abilities

We investigate the ability of our model to perform single-view 3D shape reconstruction. We compare our generated shapes against several baselines [42], [43], [44], [45], [46] and SoTA works [41], [47], evaluating their Intersection over Union (IoU) and F-score. Since IC3D is a generative model, we display the maximum scores obtained when sampling an increasing amount of shapes. Results, presented in Tab. III, show that our generative model achieves a satisfactory performance, on par with all baselines, even when used for 3D reconstruction, although this was not our target. Figure 10 shows

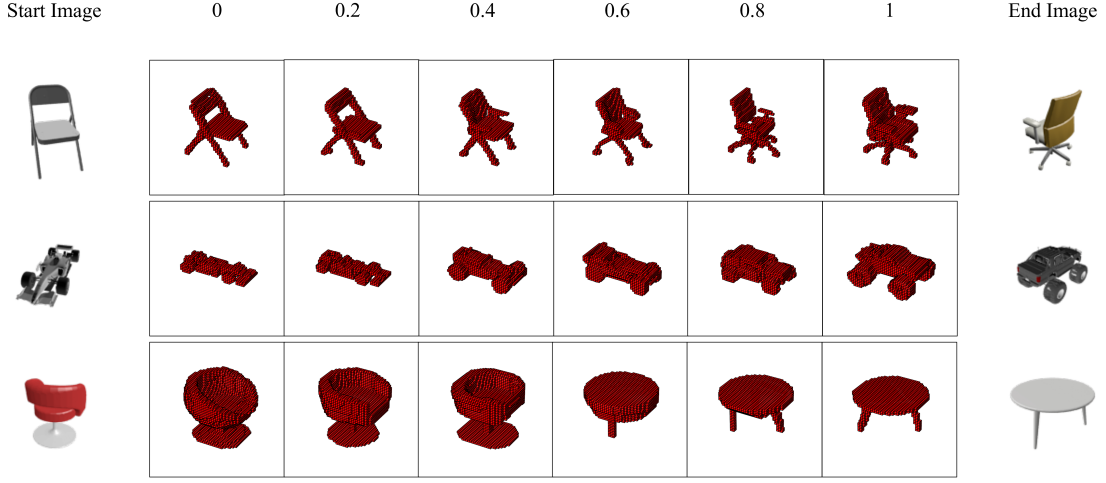


Fig. 8: The regularity of our CISP concept space allows IC3D to generate consistent shapes even when interpolating between two unrelated embeddings. Notice how the intermediate shapes mutate smoothly while maintaining perfect structural realism throughout the transformation. At each interpolation step (from 0.0 to 1.0), only a few, crisp structural details are changed, behavior that we observe indifferently in intra-class (rows 1–2) and inter-class (row 3) interpolations.

	3D-R2N2 [42]	OGN [43]	Pixel2Mesh [44]	IM-Net [45]	Pix2Vox++/F [46]	3D-RETR [41]	TMVNet [47]	Ours (1)	Ours (5)	Ours (10)	Ours (15)
IoU	0,592	0,633	0,554	0,701	0,664	0,719	0,761	0,579	0,633	0,649	0,658
F-Score	0,377	0,409	0,416	0,445	0,485	0,492	0,572	0,363	0,402	0,414	0,421

TABLE III: Reconstruction performance of IC3D against 3D reconstruction models. To evaluate our reconstruction performance, we sample n shapes (n in brackets) and report the maximum reconstruction metric obtained. Being IC3D a generative model, its reconstruction score increases as we sample more shapes.




Fig. 9: Generation from smartphone photo in a controlled environment and catalogue photo (first row) and from pix3d images (second row). Backgrounds are removed by an automatic tool before feeding them to IC3D, which is able to generate coherent and realistic shapes in all these settings.

how our shapes are more structurally correct and thus realistic, although this is often not reflected in the reconstruction metrics (which focus only on voxel correspondence).

VI. DISCUSSION AND LIMITATIONS

The experiments presented in the previous sections highlight the ability of our pipeline to produce structurally realistic and diverse shapes, correctly capturing the input data distribution. They also prove the applicability of DDPMs for the image-driven generation of voxel-based shapes. To this end, we presented CISP, a model exploiting contrastive training to learn a joint embedding space for images and shapes.




Fig. 10: We investigate the application of IC3D to 3D reconstruction, comparing its generated shapes to the reconstructions of 3D-RETR [41]. Being optimized on mere reconstruction losses, 3D-RETR achieves higher reconstruction scores (IoU); however, it fails to capture small details, such as the thin legs of a tall chair, which are instead well modeled by IC3D.

We leveraged its well-structured embedding space to guide IC3D generation toward query images. Our quantitative results showed the effectiveness of IC3D in generating high-quality shapes. Our human evaluation survey also confirmed that IC3D achieves higher structural realism than a competitive 3D reconstruction model while maintaining consistency with

the query image. We speculate that the generative nature of our model forces it to learn the structural semantics of the training data distribution, to be able to generate diverse but coherent samples. Instead, most reconstruction methods focus only on minimizing a reconstruction metric, without giving any importance to the realism, integrity, and structural correctness of the produced shape. For this reason, given for example the task of reconstructing a chair with thin legs (Fig. 10), if a reconstruction model were to omit an entire leg, its reconstruction score would not be significantly impacted, but the chair would certainly fall over. The generative nature of our model and the usage of a well-structured joint image-shape embedding space, instead, enforces the inductive biases necessary to learn the importance of producing stable chair legs.

The main limitation of IC3D resides in its low sampling speed, an issue that is well-known for simple DDPM models as the one we employ. Several techniques were recently presented to obtain a substantial speed-up [49], [50], [51], [52], and further research could be invested into integrating them within our pipeline. In order to explore the capabilities of CISP-guided DDPMs for image-conditioned shape generation, we developed IC3D focusing on the simplest 3D representation: voxels. This allowed us to decouple the complexities due to the representation (e.g., point clouds [53], [27]) from the capabilities of the model. We are aware, however, of the limitations of voxels, in particular regarding their high scaling cost. Future extensions will expand our work to produce more complex 3D data, such as point cloud or implicit functions.

Experiments in Secs. V-D and V-E highlight the importance of the CISP embedding space for producing high quality realistic results. In its current form, CISP was trained on ShapeNet, a dataset well-known but limited in size. Training CISP on larger datasets would allow for a stronger generalization power, also improving its effectiveness as a zero-shot model. Furthermore, the extension of our single-view pipeline to multi-view guidance through CISP embeddings represents an exciting research direction for future work.

REFERENCES

- [1] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in Neural Information Processing Systems*, vol. 34, pp. 8780–8794, 2021. 1, 3
- [2] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, "Glide: Towards photorealistic image generation and editing with text-guided diffusion models," 2021. [Online]. Available: <https://arxiv.org/abs/2112.10741> 1, 3
- [3] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," 2022. [Online]. Available: <https://arxiv.org/abs/2204.06125> 1, 3, 5, 6
- [4] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 684–10 695. 1
- [5] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes *et al.*, "Photorealistic text-to-image diffusion models with deep language understanding," *arXiv preprint arXiv:2205.11487*, 2022. 1
- [6] A. Brock, T. Lim, J. M. Ritchie, and N. Weston, "Generative and discriminative voxel modeling with convolutional neural networks," 2016. [Online]. Available: <https://arxiv.org/abs/1608.04236> 1
- [7] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, "Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016. [Online]. Available: <https://proceedings.neurips.cc/paper/2016/file/44f683a84163b3523afe57c2e008bc8c-Paper.pdf> 1
- [8] H. Kim, H. Lee, W. H. Kang, J. Y. Lee, and N. S. Kim, "Softflow: Probabilistic framework for normalizing flow on manifolds," *Advances in Neural Information Processing Systems*, vol. 33, pp. 16 388–16 397, 2020. 1, 3
- [9] G. Yang, X. Huang, Z. Hao, M.-Y. Liu, S. Belongie, and B. Hariharan, "Pointflow: 3d point cloud generation with continuous normalizing flows," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4541–4550. 1, 3, 4
- [10] J. Xie, Z. Zheng, R. Gao, W. Wang, Z. Song-Chun, and Y. N. Wu, "Learning descriptor networks for 3d shape synthesis and analysis," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [11] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall, "Dreamfusion: Text-to-3d using 2d diffusion," 2022. [Online]. Available: <https://arxiv.org/abs/2209.14988> 1
- [12] S. Luo and W. Hu, "Diffusion probabilistic models for 3d point cloud generation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2837–2845. 2, 3, 4
- [13] A. Vahdat, F. Williams, Z. Gojcic, O. Litany, S. Fidler, K. Kreis *et al.*, "Lion: Latent point diffusion models for 3d shape generation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 10 021–10 039, 2022. 2, 3, 4
- [14] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," 2021. [Online]. Available: <https://arxiv.org/abs/2103.00020> 2, 3
- [15] J. Ho and T. Salimans, "Classifier-free diffusion guidance," *arXiv preprint arXiv:2207.12598*, 2022. 2, 3
- [16] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. 2
- [17] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014. [Online]. Available: <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf> 2
- [18] X.-Y. Zheng, Y. Liu, P.-S. Wang, and X. Tong, "Sdf-stylegan: Implicit sdf-based stylegan for 3d shape generation," in *Comput. Graph. Forum (SGP)*, Oct. 2022. 2, 4
- [19] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *International conference on machine learning*. PMLR, 2015, pp. 1530–1538. 2
- [20] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. Huang, "A tutorial on energy-based learning," *Predicting structured data*, vol. 1, no. 0, 2006. 2
- [21] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265. 2
- [22] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 6840–6851. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/4c5bfcfec8584af0d967f1ab10179ca4b-Paper.pdf> 2, 3
- [23] R. Klokov, E. Boyer, and J. Verbeek, "Discrete point flow networks for efficient point cloud generation," in *European Conference on Computer Vision*. Springer, 2020, pp. 694–710. 2, 3, 4
- [24] Y. Du and I. Mordatch, "Implicit generation and modeling with energy based models," *Advances in Neural Information Processing Systems*, vol. 32, 2019. 2
- [25] E. Nijkamp, M. Hill, S.-C. Zhu, and Y. N. Wu, "Learning non-convergent non-persistent short-run mcmc toward energy-based model," *Advances in Neural Information Processing Systems*, vol. 32, 2019. 2

- [26] L. Zhou, Y. Du, and J. Wu, “3d shape generation and completion through point-voxel diffusion,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5826–5835. 2, 3, 4
- [27] Z. Liu, H. Tang, Y. Lin, and S. Han, “Point-voxel cnn for efficient 3d deep learning,” *Advances in Neural Information Processing Systems*, vol. 32, 2019. 2, 8
- [28] K.-H. Hui, R. Li, J. Hu, and C.-W. Fu, “Neural wavelet-domain diffusion for 3d shape generation,” in *SIGGRAPH Asia 2022 Conference Papers*, ser. SA ’22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: <https://doi.org/10.1145/3550469.3555394> 2, 3, 4
- [29] Y. Li, H. Su, C. R. Qi, N. Fish, D. Cohen-Or, and L. J. Guibas, “Joint embeddings of shapes and images via cnn image purification,” *ACM transactions on graphics (TOG)*, vol. 34, no. 6, pp. 1–12, 2015. 2
- [30] W. Kuo, A. Angelova, T.-Y. Lin, and A. Dai, “Mask2cad: 3d shape prediction by learning to segment and retrieve,” in *European Conference on Computer Vision*. Springer, 2020, pp. 260–277. 2
- [31] —, “Patch2cad: Patchwise embedding learning for in-the-wild shape retrieval from a single image,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 589–12 599. 2
- [32] R. Girdhar, A. El-Nouby, Z. Liu, M. Singh, K. V. Alwala, A. Joulin, and I. Misra, “Imagebind: One embedding space to bind them all,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 15 180–15 190. 2
- [33] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020. 3
- [34] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, “Training data-efficient image transformers & distillation through attention,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 10 347–10 357. 3
- [35] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018. 3
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017. 3
- [37] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su *et al.*, “Shapenet: An information-rich 3d model repository,” *arXiv preprint arXiv:1512.03012*, 2015. 3
- [38] R. Cai, G. Yang, H. Averbuch-Elor, Z. Hao, S. Belongie, N. Snavely, and B. Hariharan, “Learning gradient fields for shape generation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 3, 4
- [39] D. Lopez-Paz and M. Oquab, “Revisiting classifier two-sample tests,” *arXiv preprint arXiv:1610.06545*, 2016. 4
- [40] J. Y. Koh, J. Baldridge, H. Lee, and Y. Yang, “Text-to-image generation grounded by fine-grained user attention,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 237–246. 4
- [41] Z. M. Zai Shi, Y. Xing, Y. Ma, and R. Wattenhofer, “3d-retr: End-to-end single and multi-view 3d reconstruction with transformers,” in *BMVC*, 2021. 4, 5, 6, 7
- [42] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese, “3d-r2n2: A unified approach for single and multi-view 3d object reconstruction,” in *European conference on computer vision*. Springer, 2016, pp. 628–644. 6, 7
- [43] M. Tatarchenko, A. Dosovitskiy, and T. Brox, “Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2088–2096. 6, 7
- [44] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang, “Pixel2mesh: Generating 3d mesh models from single rgb images,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 52–67. 6, 7
- [45] Z. Chen and H. Zhang, “Learning implicit fields for generative shape modeling,” *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 6, 7
- [46] H. Xie, H. Yao, S. Zhang, S. Zhou, and W. Sun, “Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images,” *International Journal of Computer Vision*, vol. 128, no. 12, pp. 2919–2935, 2020. 6, 7
- [47] K. Peng, R. Islam, J. Quarles, and K. Desai, “Tmvnet: Using transformers for multi-view voxel-based 3d reconstruction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 222–230. 6, 7
- [48] X. Sun, J. Wu, X. Zhang, Z. Zhang, C. Zhang, T. Xue, J. B. Tenenbaum, and W. T. Freeman, “Pix3d: Dataset and methods for single-image 3d shape modeling,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [49] A. Q. Nichol and P. Dhariwal, “Improved denoising diffusion probabilistic models,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 8162–8171. 8
- [50] D. Watson, J. Ho, M. Norouzi, and W. Chan, “Learning to efficiently sample from diffusion probabilistic models,” *arXiv preprint arXiv:2106.03802*, 2021. 8
- [51] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” *arXiv preprint arXiv:2010.02502*, 2020. 8
- [52] A. Vahdat, K. Kreis, and J. Kautz, “Score-based generative modeling in latent space,” in *Neural Information Processing Systems (NeurIPS)*, 2021. 8
- [53] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660. 8

Supplementary Material

A. DIFFUSION MODELS

Denoising Diffusion Probabilistic models [1], [2] are latent variable models consisting of two processes: the forward noising process q and the backward denoising process p . We define x_0 as the (clean) input data and x_1, \dots, x_T as the latent variables sampled at each noising step, each with the same dimensionality as x_0 . The forward process is a fixed Markov chain gradually adding noise to the data with a variance schedule β_1, \dots, β_T . Each forward step is then a Gaussian transition following $q(x_t|x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}; \beta_t\mathbf{I})$. With a well-designed variance schedule and a sufficiently large T , we get that x_T converges to an isotropic Gaussian distribution. Also the backward process is parametrized by a Markov chain with Gaussian transitions, but it is learned instead, and the initial state is sampled from $x_T \sim \mathcal{N}(x_T; 0; \mathbf{I})$. Backward transitions follow $p(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t); \Sigma_\theta(x_t, t))$, where θ are the model parameters. A crucial property of this formulation is that the defined forward process allows us to directly sample from any given timestep. In fact, defining $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=0}^t \alpha_s$ we can write the marginal at timestep t as

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0; (1 - \bar{\alpha}_t)\mathbf{I}), \quad (1)$$

which allows us to sample x_t as

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, 1). \quad (2)$$

B. TRAINING APPROACH AND SAMPLE GENERATION

As q and p form a Variational Auto-Encoder [3], we can write a variational lower bound (VLB) [3] and use it as objective for the model. However, Ho *et al.* [1] propose a reparametrization of $\mu_\theta(x_t, t)$ allowing a simpler objective and better sample quality with respect to optimizing the VLB. Instead of predicting $\mu_\theta(x_t, t)$, they propose to predict the noise ϵ , expressing $\mu_\theta(x_t, t)$ as

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}}(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_\theta(x_t, t)). \quad (3)$$

This formulation allows for a simpler objective function, which is a reweighted form of the variational lower bound:

$$L_{simple} = E_{t, x_0, \epsilon}[\|\epsilon - \epsilon_\theta(x_t, t)\|^2]. \quad (4)$$

As this formulation gives no learning signal for the variances $\Sigma_\theta(x_t, t)$, Ho *et al.* [1] propose to fix $\Sigma_\theta(x_t, t)$ to time-dependent constants $\sigma_t^2 = \beta_t$. We can then train the model by randomly sampling a timestep t , applying Eq. (1) to get a noised sample x_t , use (x_t, t) as input to our model to predict $\epsilon_\theta(x_t, t)$ and optimize the loss in Eq. (4).

To sample from the trained model, we first generate a sample $x_T \sim \mathcal{N}(0, 1)$ and we use the trained network to predict $\epsilon_\theta(x_t, t)$, thus obtaining $\mu_\theta(x_t, t)$ from Eq. (3). We can now sample $x_{t-1} \sim \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t); \sigma_t^2\mathbf{I})$. This process is repeated recursively, eventually obtaining x_0 , the generated output of our pipeline.

C. GUIDANCE

We apply classifier-free guidance [4] to guide our model. To this end, we jointly train a conditional and an unconditional model, then we make predictions by combining their score estimates to step toward the guidance direction. To jointly train a conditional and an unconditional model, we replace with probability p the input conditioning with a null token \emptyset . At inference time, we combine the conditional and unconditional predictions as

$$\hat{\epsilon}_\theta(x_t, t|y) = \epsilon_\theta(x_t, t|\emptyset) + w \cdot (\epsilon_\theta(x_t, t|y) - \epsilon_\theta(x_t, t|\emptyset)), \quad (5)$$

where y is the guidance token and $w \geq 1$ is the guidance scale. We apply classifier-free guidance with $p = 0.1$ on image tokens and CISP embeddings independently.

A. Guidance scale

In the original paper [4] and in GLIDE [5], both dealing with text-to-image generation, it was shown that the guidance scale balances between quality and diversity: higher guidance scales produce samples of higher quality and more coherent to the conditioning tokens, thus reducing the sample diversity. However, in [6] the authors find the guidance scale to only affect photorealism (e.g., lighting, transparency) and not diversity, thus obtaining images correctly matching the input text across different guidance scales. We also find that the diversity of our generated shapes is independent of the guidance scale, although in our case 1-NNA and visual realism also remain similar. This could be due to the nature of our 3D representation. Indeed, voxels only represent the

Shape	Model	1-NNA(%)	
		CD	EMD
Aeroplane	Unguided	62.78	61.13
	Class-guided	63.23	61.95
	CISP	75.12	66.94
	Null token + null token	58.93	56.93
	CISP + null token	58.11	54.47
	CISP + E_c (IC3D)	57.64	53.89
Car	Unguided	58.17	57.01
	Class-guided	57.25	56.37
	CISP	62.23	58.69
	Null token + null token	53.20	53.11
	CISP + null token	53.35	52.04
	CISP + E_c (IC3D)	52.44	51.68
Chair	Unguided	70.34	71.39
	Class-guided	68.98	68.98
	CISP	68.77	61.49
	Null token + null token	53.30	51.97
	CISP + null token	54.18	52.20
	CISP + E_c (IC3D)	53.58	51.73
Average	Unguided	63.76	63.18
	Class-guided	63.15	62.43
	CISP	68.71	62.37
	Null token + null token	55.14	54.00
	CISP + null token	55.21	52.90
	CISP + E_c (IC3D)	54.55	52.43

TABLE I: Extended ablation study on CISP and E_c . Best results are highlighted in bold.

shape of an object and do not consider complex aspects such as textures and light effects. Thus, if low guidance scales already obtain coherent and realistic shapes, we do not expect that more relevant details can be added when increasing the guidance scale.

B. Guidance tokens ablation

Our approach leverages both CISP and E_c tokens to produce shapes that are realistic and coherent to the input image. As discussed in the main paper (Sec. 5.6), training a model with only E_c tokens is not effective in guiding the generation toward the input image. Differently, even though the introduction of CISP tokens enables an effective guidance, the joint use of E_c and CISP tokens provides even better results, allowing superior modeling of low-level details and symmetries.

To further investigate the role of our guidance elements, we extend here our ablation by testing our model in 3 additional settings. First, we analyze the impact of training our DDPM in an unguided and class-guided setting, as to consider these two setups as baselines for our image-guided method. In the unguided setting, we train multiple models to generate shapes without any guidance input. The class-guided scenario, instead, focuses on a single model, which is guided by a class

token to generate shapes of different categories. The results, reported in Tab. I, show that unguided and class-guided models obtain scores comparable to those of CISP-only guidance. CISP-only guidance, however, allows us to drive the generation through an input image instead of a class token, resulting in a more fine control of the generated shapes. As reported in the main paper, the addition of E_c at training time allows us to obtain much higher generation quality. Instead, at generation time, replacing its embeddings with the learned null token at generation time does not significantly impact the results.

D. ADDITIONAL GENERATION METRICS

Following PVD [7] and PointFlow [8], we also measure the results of our model on Minimum Matching Distance (MMD) and Coverage (COV). We define S_g as the set of generated samples and S_r as the set of reference samples (taken from the test set), with $|S_r| = |S_g|$. Minimum matching distance measures the distance from each sample in S_r to its nearest neighbor in S_g . It is defined as

$$\text{MMD}(S_g, S_r) = \frac{1}{|S_r|} \sum_{r \in S_r} \min_{g \in S_g} d(g, r), \quad (6)$$

where $d(\cdot)$ is a distance metric. As done in the literature, we consider both Chamfer Distance and Earth Mover's Distance. While MMD measures the quality of the generated samples, it does not take into account their diversity. A measure of their diversity is given by COV, widely used to complement MMD.

COV measures the percentage of samples in S_r that are matched with at least one of the samples in S_g . A match is defined as the nearest neighbor according to a distance function d . COV can then be defined as

$$\text{COV}(S_g, S_r) = \frac{|\arg \min_{r \in S_r} d(g, r) \mid g \in S_g|}{|S_r|}. \quad (7)$$

Tab. II shows MMD and COV results of our model compared to previous works. Our model is comparable to the others in MMD, which slightly suffers from the conversion from voxel to point cloud; on the other hand, it obtains significantly higher results in COV in nearly all cases. In these two cases LION [12] performs slightly better, however our model is the second best and still competitive to it.

E. GENERATION METRICS FOR ADDITIONAL CATEGORIES

We report in Tab. III our results on the table and watercraft category, on which generative models in the literature do not compare.

Shape	Model	MMD↓		COV(%)↑	
		CD	EMD	CD	EMD
Airplane	Shape-GF[9]	2.703	0.659	40.74	40.49
	DPF-Net[10]	0.264	0.409	46.17	48.89
	PointFlow[8]	0.224	0.390	47.90	46.41
	[11]	0.221	0.385	49.02	44.98
	PVD[7]	0.228	0.380	48.88	52.09
	LION[12]	0.219	0.372	47.16	49.63
	Ours	0.224	0.386	54.03	54.13
Car	Shape-GF[9]	9.232	0.756	49.43	50.28
	DPF-Net[10]	1.129	0.853	45.74	49.43
	PointFlow[8]	0.901	0.807	46.88	50.00
	[11]	-	-	-	-
	PVD[7]	1.077	0.794	41.19	50.56
	LION[12]	0.913	0.752	50.00	56.53
	Ours	1.083	0.812	50.90	54.93
Chair	Shape-GF[9]	2.889	1.702	46.67	48.03
	DPF-Net[10]	2.536	1.632	44.71	48.79
	PointFlow[8]	2.409	1.595	42.90	50.00
	[11]	2.415	1.564	49.66	50.22
	PVD[7]	2.622	1.556	49.84	50.60
	LION[12]	2.640	1.550	48.94	52.11
	Ours	2.492	1.573	51.08	51.44

TABLE II: Comparison of our model on generation metric against other works. The car category was not part of the evaluation of [11]. Best results are highlighted in bold.

Shape	1-NNA(%)		MMD↓		COV(%)↑	
	CD	EMD	CD	EMD	CD	EMD
Table	55.63	53.89	2.326	1.573	53.85	54.10
Watercraft	56.95	54.03	1.984	1.453	53.57	52.50

TABLE III: Generation performance of IC3D on additional object categories.

F. CONSIDERATIONS AGAINST STANDARD RECONSTRUCTION METRICS

As discussed in the main paper, metrics such as IoU and F-score fail to capture properties such as realism, integrity and structural correctness of 3D shapes. Indeed, they only measure the reconstruction quality by considering the number of correctly predicted voxels values. As a consequence, models trained to optimize these quantities can generate structurally wrong shapes, such as tables and chairs with discontinuous or missing legs, or cars with flat wheels. Indeed, such errors are caused by only a few wrongly predicted voxels, thus having a small impact on metrics such as IoU and F-score. Moreover, IoU is sensitive to the slightest shift or scale of the predicted shape with respect to the ground truth. This is problematic for our model, since, being generative, it may produce a conceptually correct shape but place it slightly scaled or shifted compared to the ground truth. Examples of these scenarios are highlighted in Fig. 1.




Fig. 1: Even if our model produces coherent and realistic shapes, IoU favors the shapes reconstructed by 3D-RETR [13], which may contain structural integrity problems. Although this image is already reported in the main paper (Fig. 10), we report it here for ease of discussion.




Fig. 2: Human evaluation interface. The query image is shown only after the first question has been answered. The second question is unlocked together with the image.

Although our model produces realistic and structurally intact shapes, they may exhibit different details and be slightly scaled or shifted. This is due to the probabilistic approach and heavily impacts our overall score. These considerations led us to perform a human evaluation, reported in the main paper, in order to evaluate both our coherence to the query image and the realism of our shapes from a human perspective.

Category	Votes for IC3D (%)						
	0/5	1/5	2/5	3/5	4/5	5/5	3/5 or higher
Aeroplane	6.00	16.50	16.00	24.50	19.50	17.50	61.50
Car	12.50	9.50	19.00	23.50	19.50	16.00	59.00
Chair	7.50	11.00	13.00	20.00	16.50	32.00	68.50
Overall	8.67	12.33	16.00	22.67	18.50	21.83	63.00

TABLE IV: **Coherence** per-class human evaluation results. Our model obtains the majority vote (in bold) on all object categories.

Category	Votes for IC3D (%)						
	0/5	1/5	2/5	3/5	4/5	5/5	3/5 or higher
Aeroplane	3.50	12.50	19.00	16.50	21.00	27.50	65.00
Car	9.50%	18.50	20.50	18.50	19.00	14.00	52.00
Chair	4.00	4.00	1.50	9.50	20.00	61.00	91.00
Overall	5.67	11.67	13.67	14.83	20.00	34.17	69.00

TABLE V: **Realism** per-class human evaluation results. Our model obtains the majority vote (in bold) on all object categories.

	3D-R2N2 [14]	OGN [15]	Pixel2Mesh [16]	IM-Net [17]	Pix2Vox++/F [18]	3D-RETR [13]	TMVNet [19]	Ours (1)	Ours (5)	Ours (10)	Ours (15)
Aeroplane	0.512	0.587	0.508	0.702	0.607	0.704	0.691	0.540	0.600	0.620	0.630
Car	0.798	0.828	0.670	0.756	0.841	0.861	0.870	0.790	0.824	0.833	0.838
Chair	0.466	0.483	0.484	0.644	0.548	0.592	0.721	0.407	0.476	0.494	0.506

TABLE VI: Per-class IoU single-view 3D reconstruction results.

	3D-R2N2 [14]	OGN [15]	Pixel2Mesh [16]	IM-Net [17]	Pix2Vox++/F [18]	3D-RETR [13]	TMVNet [19]	Ours (1)	Ours (5)	Ours (10)	Ours (15)
Aeroplane	0.412	0.487	0.376	0.589	0.583	0.612	0.594	0.492	0.543	0.561	0.570
Car	0.481	0.514	0.486	0.304	0.564	0.511	0.602	0.394	0.425	0.433	0.439
Chair	0.238	0.226	0.386	0.442	0.309	0.352	0.520	0.203	0.238	0.248	0.255

TABLE VII: Per-class F-score single-view 3D reconstruction results.

G. HUMAN EVALUATION ADDITIONAL MATERIAL

Interface details We report additional details on the human evaluation procedure. Fig. 2 shows the interface presented to our human evaluators. To ensure that the evaluators would not be biased by our interface, we adopted the following precautions: (1) the shapes generated by our model and by 3D-RETR are placed to the left or to the right randomly. (2) the GIFs are synchronized to allow for a direct comparison of details. (3) the query image is centered to avoid a closeness bias. Throughout our evaluation campaign, we recorded only the anonymized answers of our evaluators. We did not record any personal information, and we made this clear to each participant.

Per-class results We report the per-class human evaluation results in Tabs. IV and V. Human evaluators prefer shapes generated by our model for both coherence and

realism in all categories. Notice how the scores obtained by our model for the chair class are the highest among all categories. Our model achieves the majority of the votes for realism in 91% of the tested shapes. For cars, instead, we register the lower margin, while still being preferred by the evaluators. We think this is because cars are characterized by less complex shapes, thus both models produce very similar results, forcing the evaluators to chose almost randomly between the two. Single-view 3D reconstruction results (Tabs. VI and VII) also support this hypothesis, as, for most models, cars are the best reconstructed class, while chairs are the worst. In fact, among the three tested categories, chairs display the most complex structure, as they often present thin elements. For example, legs and armrests are often modeled incompletely or incorrectly by reconstruction pipelines. Our model IC3D, instead, thanks to its generative nature, is

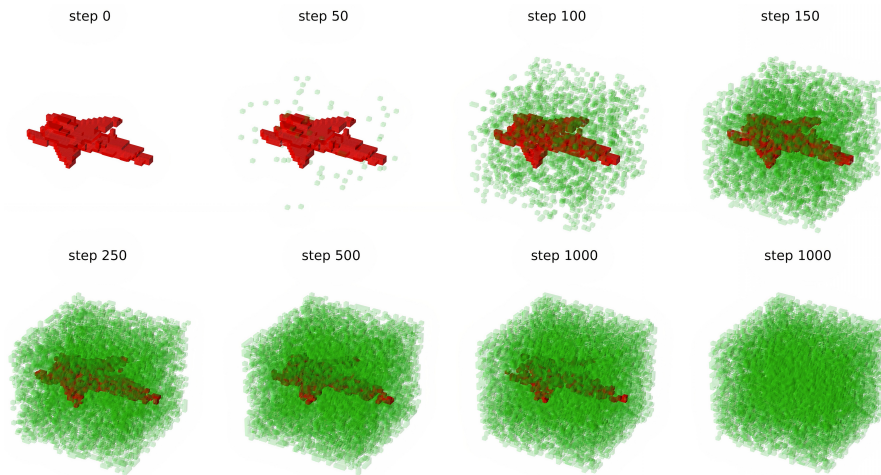


Fig. 3: Voxel diffusion process. At each diffusion step, noise is added to the volumetric data, and samples are thresholded at 0.5 to visualize binary data. Originally filled values are shown in red for comparison. Step 1000 is displayed also without red highlighting, showing how the original shape is completely lost.

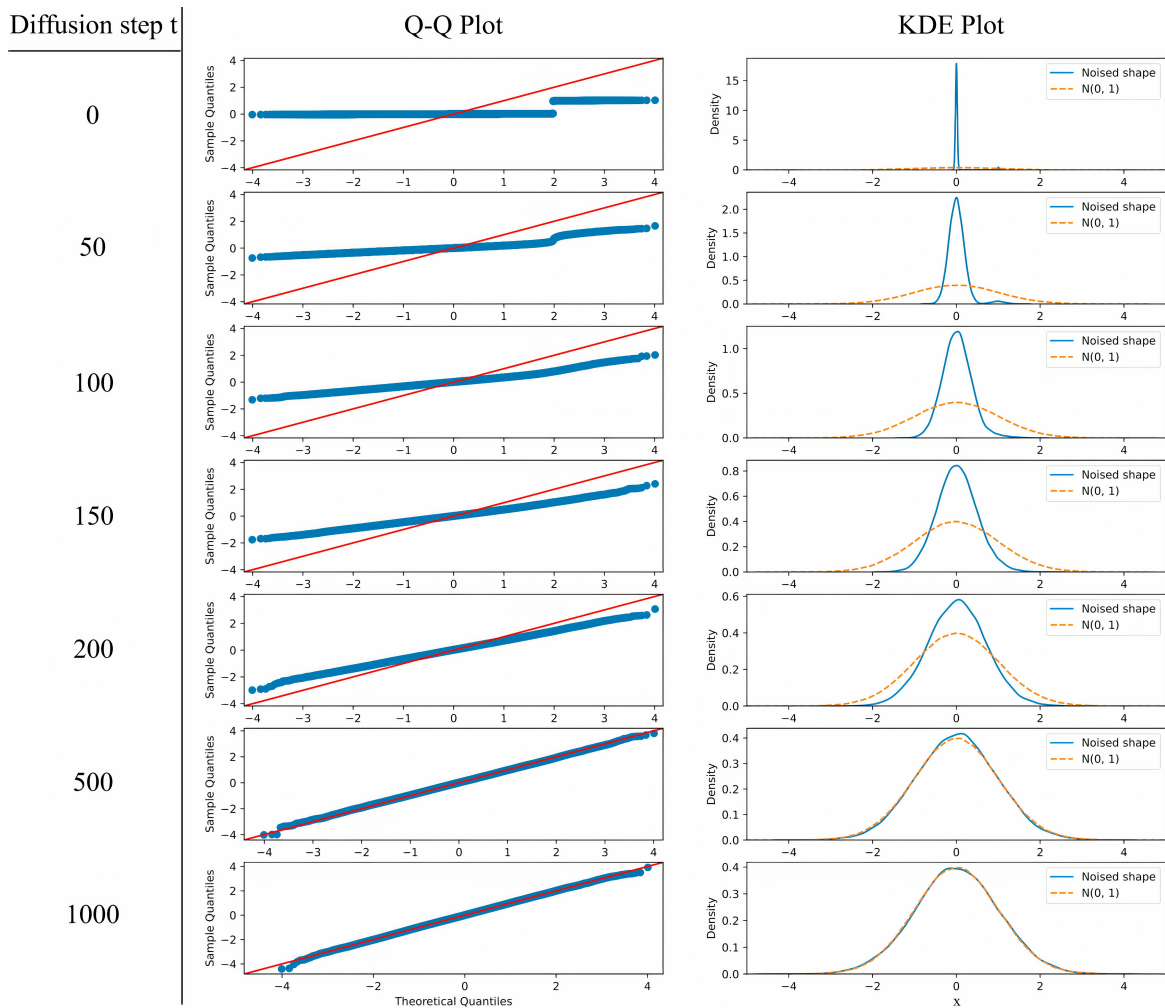


Fig. 4: Distribution of values throughout the forward process. We observe that the diffusion process progressively transforms the initial distribution into a standard Normal distribution.

able to correctly produce structurally realistic samples.

H. ANALYZING VOXEL DIFFUSION

As voxel DDPMs have never been deployed before, we study the effect of the diffusion process with Gaussian transitions on voxels. In fact, being voxels binary data, we may ask if a Gaussian process is effectively able to gradually destroy the information in a sample, which is crucial for the reverse step.

We study the distribution of volume values across different steps in the forward diffusion process, and show the results in Figs. 3 and 4.

Specifically, Fig. 3 visually shows the effect of the forward process on a binarized shape. We highlight in red voxels with value 1 belonging to the original shape and in green voxels with value 1 which are not from the original shape. We can observe both how originally filled voxels are erased by the process and how new voxels are filled. The result of the last diffusion step is displayed with and without highlighting, to show that information about the original shape is lost.

In Fig. 4, we analyze the distribution of the data, comparing it to a standard Normal distribution at different timesteps by computing Quantile-Quantile (QQ) and Kernel Density Estimation (KDE) plots. QQ plots assess the plausibility that a set of data is coming from a theoretical distribution (in our case, a standard Normal distribution) by plotting their sets of quantiles against one another. Data coming from the theoretical distribution give origin to a straight line, as in our case. Instead, KDE plots allow us to visualize the distribution of our data by estimating its density function. Fig. 4 shows QQ and KDE at different timesteps, proving how the sample distribution progressively approximates a standard Normal distribution.

REFERENCES

- [1] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 6840–6851. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/4c5bfec8584af0d967f1ab10179ca4b-Paper.pdf> 10
- [2] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265. 10
- [3] D. P. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. 10
- [4] J. Ho and T. Salimans, “Classifier-free diffusion guidance,” *arXiv preprint arXiv:2207.12598*, 2022. 10
- [5] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, “Glide: Towards photorealistic image generation and editing with text-guided diffusion models,” 2021. [Online]. Available: <https://arxiv.org/abs/2112.10741> 10
- [6] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with clip latents,” 2022. [Online]. Available: <https://arxiv.org/abs/2204.06125> 10
- [7] L. Zhou, Y. Du, and J. Wu, “3d shape generation and completion through point-voxel diffusion,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5826–5835. 11, 12
- [8] G. Yang, X. Huang, Z. Hao, M.-Y. Liu, S. Belongie, and B. Hariharan, “Pointflow: 3d point cloud generation with continuous normalizing flows,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4541–4550. 11, 12
- [9] R. Cai, G. Yang, H. Averbuch-Elor, Z. Hao, S. Belongie, N. Snavely, and B. Hariharan, “Learning gradient fields for shape generation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 12
- [10] R. Klokov, E. Boyer, and J. Verbeek, “Discrete point flow networks for efficient point cloud generation,” in *European Conference on Computer Vision*. Springer, 2020, pp. 694–710. 12
- [11] S. Luo and W. Hu, “Diffusion probabilistic models for 3d point cloud generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2837–2845. 12
- [12] A. Vahdat, F. Williams, Z. Gojcic, O. Litany, S. Fidler, K. Kreis *et al.*, “Lion: Latent point diffusion models for 3d shape generation,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 10 021–10 039, 2022. 11, 12
- [13] Z. M. Zai Shi, Y. Xing, Y. Ma, and R. Wattenhofer, “3d-retr: End-to-end single and multi-view 3d reconstruction with transformers,” in *BMVC*, 2021. 12, 13
- [14] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese, “3d-r2n2: A unified approach for single and multi-view 3d object reconstruction,” in *European conference on computer vision*. Springer, 2016, pp. 628–644. 13
- [15] M. Tatarchenko, A. Dosovitskiy, and T. Brox, “Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2088–2096. 13
- [16] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang, “Pixel2mesh: Generating 3d mesh models from single rgb images,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 52–67. 13
- [17] Z. Chen and H. Zhang, “Learning implicit fields for generative shape modeling,” *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 13
- [18] H. Xie, H. Yao, S. Zhang, S. Zhou, and W. Sun, “Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images,” *International Journal of Computer Vision*, vol. 128, no. 12, pp. 2919–2935, 2020. 13
- [19] K. Peng, R. Islam, J. Quarles, and K. Desai, “Tmvnet: Using transformers for multi-view voxel-based 3d reconstruction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 222–230. 13