

3D-C2FT: Coarse-to-fine Transformer for Multi-view 3D Reconstruction

Leslie Ching Ow Tiong^{1,*} [0000–0003–3786–2117], Dick
 Sigmund^{2,*} [0000–0002–6207–5804], and Andrew Beng Jin
 Teoh^{3,†} [0000–0001–5063–9484]

¹ Computational Science Research Center, Korea Institute of Science and
 Technology, 5, Hwarang-ro 14-gil, Seongbuk-gu, Seoul 02792, Republic of Korea
tiongleslie@kist.re.kr

² AIDOT Inc., 128, Beobwon-ro, Songpa-gu, Seoul 05854, Republic of Korea
dsigmund@aidot.ai

³ School of Electrical and Electronic Engineering, Yonsei University, Seoul 120-749,
 Republic of Korea
bjteoh@yonsei.ac.kr

Abstract. Recently, the transformer model has been successfully employed for the multi-view 3D reconstruction problem. However, challenges remain in designing an attention mechanism to explore the multi-view features and exploit their relations for reinforcing the encoding-decoding modules. This paper proposes a new model, namely 3D coarse-to-fine transformer (3D-C2FT), by introducing a novel coarse-to-fine (C2F) attention mechanism for encoding multi-view features and rectifying defective voxel-based 3D objects. C2F attention mechanism enables the model to learn multi-view information flow and synthesize 3D surface correction in a coarse to fine-grained manner. The proposed model is evaluated by ShapeNet and Multi-view Real-life voxel-based datasets. Experimental results show that 3D-C2FT achieves notable results and outperforms several competing models on these datasets.

Keywords: Multi-view 3D reconstruction · Coarse-to-fine transformer · Multi-scale attention.

1 Introduction

Multi-view 3D reconstruction infers geometry and structure of voxel-based 3D objects from multi-view images. The definition of the voxel is essentially 3D pixels in cube size. Theoretically, voxel representation is the perfect modeling technique for replicating reality [9, 10, 13]. Hence, voxel representation can easily resemble real-world objects into more accurate 3D objects in the virtual environment [14]. Therefore, the multi-view 3D reconstruction plays a crucial role in

*The authors have contributed equally to this work.

†Corresponding author.

robot perception [2, 22], historical artifact [16, 12], dentistry [20, 21], etc. Predicaments in estimating 3D structure of an object include solving an ill-posed inverse problem of 2D images, violation of overlapping views, and unconstrained environment conditions [7]. Here, the ill-posed inverse problem refers to the voxel-based 3D objects estimation problem when disordered or limited sources of the inputs are given as inputs [6]. Accordingly, multi-view 3D object reconstruction remains an active topic in computer vision.

Classical approaches to tackle multi-view 3D reconstruction are to introduce feature extraction techniques to map different reconstruction views [24, 5, 26, 31]. However, these approaches hardly reconstruct a complete 3D object when a small number of images, e.g., 1–4, are presented. With the advances of deep learning (DL), most attempts utilize an encoder-decoder network architecture for feature extraction, fusion, and reconstruction. Existing DL models usually rely on the convolutional neural networks (CNN) or recurrent neural networks (RNN) to fuse multiple deep features encoded from 2D multi-view images [3, 25]. However, the CNN encoder processes each view image independently, and thus, the relations in different views are barely utilized. This leads to difficulty designing a fusion method that can traverse the relationship between views. Although the RNNs can rectify the fusion problem, the model suffers a long processing time [3].

Recently, transformers [27] have gained exponential attention and proved to be a huge success in vision-related problems [4, 18]. In the multi-view 3D reconstruction problem, [28] and [35] propose a unified multi-view encoding, fusion, and decoding framework via a transformer to aggregate features among the patch embeddings that can explore the profound relation between multi-view images and perform the decent reconstruction. However, both studies only consider single-scaled multi-head self-attention (MSA) mechanism to reinforce the view-encoding. Unfortunately, the native single-scaled MSA is not designed to explore the correlation of relevant features between the subsequent layers for object reconstruction.

Considering the limitations above, we propose a new transformer model, namely 3D coarse-to-fine transformer (3D-C2FT). The 3D-C2FT introduces a novel coarse-to-fine (C2F) attention mechanism for encoded features aggregation and decoded object refinement in a multi-scale manner. On top of exploring multi-view images relationship via MSA in the native transformer, the C2F attention mechanism enables aggregating coarse (global) to fine (local) features, which is favored to learn comprehensive and discriminative encoded features. Then, a concatenation operation is used to aggregate multiple sequential C2F features derived from each C2F attention block. In addition, a C2F refiner is devised to rectify the defective decoded 3D objects. Specifically, the refiner leverages a C2F cube attention mechanism to focus on the voxelized object’s attention blocks that benefit coarse-to-fine correction of the 3D surface. In general, C2F attention utilizes the coarse and fine-grained features to be paired and promotes the information flow that helps the 3D objects reconstruction task.

Besides that, we compile a new dataset called Multi-view Real-life dataset for evaluation. However, unlike existing dataset such as ShapeNet [32] that is assembled from synthetic data, our dataset is composed of real-life single or multi-view images taken from internet without ground truth. We propose this dataset to support real-life study for multi-view 3D reconstruction evaluation. The contributions of this paper are summarized as follows:

- A novel multi-scale C2F attention mechanism is outlined for the multi-view 3D reconstruction problem. The proposed C2F attention mechanism learns the correlation of the features in a sequential global to the local manner, on top of exploring multi-view images relation via MSA. Accordingly, coarse-grained features draw attention towards the global object structure, and the fine-grained feature pays attention to each local part of the 3D object.
- A novel transformer model, namely 3D-C2FT, for multi-scale multi-view images encoding, features fusion, and coarse-to-fine decoded objects refinement is proposed.
- Extensive experiments demonstrate better reconstruction results on standard benchmark ShapeNet dataset and challenging real-life dataset than several competing models, even under stringent constraints such as occlusion.

This paper is organized as follows: Section 2 reviews the related works of the DL-based multi-view 3D reconstruction. Then, Section 3 presents our proposed method. Next, the experimental setup, results, and ablation study are presented in Section 4. Finally, the conclusions are summarized in the last section.

2 Related Works

This section reviews several relevant works of DL-based multi-view 3D reconstruction. We refer the readers to a more comprehensive survey on this subject [7].

2.1 CNN and RNN-based Models

Early works [3] and [11] utilize modified RNN to fuse information from multi-view images and reconstruct the 3D objects. However, both works have their limitations; for instance, when given the same set of images but in different orders, the model fails to reconstruct the 3D objects consistently. This is due to both models being permutation-variant and relying on the ordered sequence of the input images for feature extraction.

[33], [34] and [36] propose a CNN-based model to address randomly ordered and long sequence forgetting issues of the RNNs. However, these works follow the divide-and-conquer principle by introducing a CNN-based single-view encoder, a single-view decoder, and a fusion model, which work independently. Therefore, the encoder and decoder hardly utilize the relations between different views. Instead, the network relies on the fusion model to integrate the

arbitrary ordered multi-view features for reconstruction. Results suggest that designing a robust fusion method with a CNN-based approach is challenging. Furthermore, these works suffer from the model scaling problem while preserving permutation-invariant capability. For instance, if the input views exceed a specific number, such as 10–16 views, the reconstruction performance level-off, implying the hardness of learning complementary knowledge from a large set of independent inputs.

2.2 Transformer-based Models

Recently, [28] and [35] put forward the transformer-based models that perceive multi-view 3D reconstruction problems as a sequence-to-sequence prediction problem that permits fully exploiting information from 2D images with different views. [28] leverages native multi-head self-attention (MSA) mechanism for feature encoding, fusing, and decoding to generate a 3D volume for each query token. However, this model only relies on the MSA to explore the relation of multi-view images by fostering different representations of each view. Hence, it falls short in analyzing the low to high-level interactions within multi-view images that signify the global structure and local components of the 3D objects.

Yagubbayli et al. [35] proposes another transformer model simultaneously along with [28] known as Legoformer. This model adopts an encoder with pre-norm residual connections [17] and a non-autoregressive decoder that takes advantage of the decomposition factors without explicit supervision. Although this approach is more effective, the model only focuses on learning 3D objects reconstruction parts by parts. Such a strategy is deemed local-to-local attention, which does not benefit low-level interaction to support information flow across local parts of 3D objects.

Considering the limitations of the previous transformer-based models, we are motivated to design a new attention mechanism by introducing a C2F patch attention mechanism in the encoder to extract multi-scale features of the multi-view images. Furthermore, we also put forward a C2F cube attention mechanism in the refiner to rectify the reconstructed object surface in a coarse to fine-grained manner.

3 3D Coarse-to-fine Transformer

3D coarse-to-fine transformer (3D-C2FT) consists of an image embedding module, a 2D-view encoder, a 3D decoder, and a 3D refiner, as illustrated in Fig. 1. Specifically, the encoder accepts either single or multi-view images in the embedding form, which are managed by the image embedding module and then C2F patch attention is used to aggregate the different views of inputs by extracting feature representations for the decoder to reconstruct 3D volume. Finally, the refiner with the C2F cube attention mechanism is to rectify the defective reconstructed 3D volume. The 3D-C2FT network is explained in detail in the following subsections.

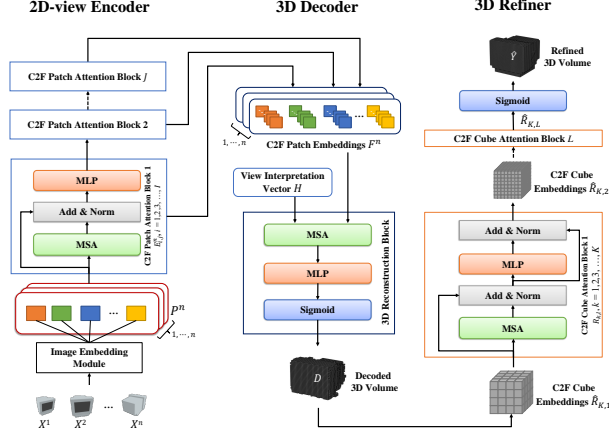


Fig. 1. The architecture of 3D-C2FT. The network comprises an image embedding module - DenseNet121, a 2D-view encoder, a 3D decoder, and a 3D refiner.

3.1 Image Embedding Module

We leverage DenseNet121 [8] as an image embedding module for the proposed 3D-C2FT. For the details, readers are referred to the supplementary materials Section 6.1. Given a set of n view images of a 3D object $X = \{X^1, X^2, \dots, X^n\}$, each view image X^n is fed into the image embedding module to obtain patch embeddings $P^n \in \mathbb{R}^{1 \times d}$, where d is the embedding dimensions. We follow the ViT setting where d is fixed at 768.

3.2 2D-view Encoder

The encoder of 3D-C2FT receives patch embeddings of view n , P^n as well as its associated positional embeddings. These patches and positional embeddings are sent to a C2F patch attention block that consists of I -layer of MSA [27] and multilayer perceptron (MLP), which is simply a fully connected feed forward network, as shown in Fig. 1. The C2F patch attention block is repeated J times. Thus, the output of the C2F patch attention block, $E_{I,j}^n$ is calculated as follows:

$$\hat{E}_{i,j}^n = \text{MSA}(\text{Norm}(P_{i-1,j}^n)) + P_{i-1,j}^n, \quad (1)$$

$$E_{i,j}^n = \text{MLP}(\text{Norm}(\hat{E}_{i,j}^n)), \quad (2)$$

where $\text{Norm}(\cdot)$ denotes layer normalization, i is the layer index of $E_{i,j}^n$ where $i = 1, \dots, I$ and I is empirically set to 4 in this paper. j is the block index of $E_{i,j}^n$, where $j = 1, \dots, J$.

Each C2F patch attention block is responsible to extract the features i.e., $E_{I,j}^n$ in a coarse to fine-grained manner by shrinking the dimension of $E_{I,j}^n$ in a ratio of two. Thereafter, we fuse each $E_{I,j}^n$ from J blocks by concatenation to

produce a C2F patch embedding, $F^n \in \mathbb{R}^{1 \times [d+d/2+\dots+d/2^{(J-1)}]}$, which captures the aggregation of global and local multi-scale features. Since d is set to 768 in this work and based on our experimental settings and computational complexity concerns, we set $J=3$ as shown in supplementary materials Section 6.1. The C2F patch embedding F^n can be determined as follows:

$$F^n = \text{Norm}([E_{I,1}^n, E_{I,2}^n, \dots, E_{I,J}^n]), \quad (3)$$

where $[\cdot]$ refers to the concatenation operator.

3.3 3D Decoder

As shown in Fig. 1, the 3D decoder learns and explores the global correlation between the multi-view images for reconstruction. The decoder only consists of a single 3D reconstruction block which is crucial for integrating and converting F^n to a 3D volume D with size $32 \times 32 \times 32$. The attention module of the 3D reconstruction block accepts both view interpretation matrix H and F^n to project different views of F^n into the D . Specifically, given F^n and $H \in \mathbb{R}^{g \times [d+d/2+\dots+d/2^{(J-1)}]}$, the decoder can easily aggregate the n -view of F^n together and share across all potential inputs in the network. Here, H is a randomly initialized learnable weight matrix that plays a role in interpreting and addressing F^n , and g denotes the number of cube embeddings that are required to assemble the D . In this paper, we set $g=8 \times 8 \times 8 = 512$. The 3D volume output of the decoder, D , is specified as:

$$D = \sigma(\text{MLP}(\text{MSA}(H, F^n))) \in \mathbb{R}^{32 \times 32 \times 32}, \quad (4)$$

where $\sigma(\cdot)$ is a sigmoid activation function. However, the reconstructed 3D volume, D at this stage is defective (Section 4.3) due to simplicity of the decoder. Therefore, a refiner is needed for further rectification.

3.4 3D Refiner

We propose a C2F cube attention mechanism intending to correct and refine the D surface for the refiner. The refiner is composed of L C2F cube attention blocks where each block consists of K layers of ViT-like attention module (Fig 1). Here, we set $K=6$, which the experiments determine. In addition, $L=2$ is set in this work as the reconstructed volume resolution is a mere $32 \times 32 \times 32$. Therefore, a larger L can be used for high-resolution volume.

Specifically, D from Eq. 4 is transformed to the C2F cube embeddings $\hat{R}_{k,l}$ by partitioning the D to $\hat{R}_{k,1}$ with a size of $8 \times 8 \times 8$ in the first C2F cube attention block. For the second block, $\hat{R}_{k,2}$ size is further reduced by a factor of two; thus $4 \times 4 \times 4$. For each block, the cube attention block output can be computed from:

$$\hat{R}_{k,l} = \text{MSA}(\text{Norm}(\hat{R}_{k-1,l})) + \hat{R}_{k-1,l}, \quad (5)$$

$$R_{k,l} = \text{MLP}(\text{Norm}(\hat{R}_{k,l})) + \hat{R}_{k,l} \quad (6)$$

where k is the layer index of $R_{k,l}$ where $k = 1, 2, \dots, K$ and l is the block index of $R_{k,l}$ where $l = 1, 2, \dots, L$.

By adopting the C2F notion, the proposed refiner leverages the C2F cube attention blocks of different scales of embeddings that can benefit the structure and parts correction of the 3D object. Furthermore, it enables the refiner to iteratively draw attention from coarse to temperate regions by rectifying the spatial surfaces gradually through multi-scale attention.

3.5 Loss Function

For training, we adopt a combination of mean-square-error (\mathcal{L}_{MSE}) loss and 3D structural similarity ($\mathcal{L}_{\text{3D-SSIM}}$) loss as a loss function ($\mathcal{L}_{\text{total}}$) to capture better error information about the quality degradation of the 3D reconstruction. The motivation of adopting $\mathcal{L}_{\text{3D-SSIM}}$ loss [30] is for 3D-C2FT to learn and produce visually pleasing 3D objects by quantifying volume structural differences between a reference and a reconstructed object. Furthermore, the \mathcal{L}_{MSE} is used to evaluate the voxel loss that measures the L2 differences between the ground-truth (GT) and reconstructed objects. Thus, $\mathcal{L}_{\text{total}}$ is given as follows:

$$\mathcal{L}_{\text{total}}(Y, \hat{Y}) = \mathcal{L}_{\text{MSE}}(Y, \hat{Y}) + \mathcal{L}_{\text{3D-SSIM}}(Y, \hat{Y}), \quad (7)$$

$$\mathcal{L}_{\text{MSE}}(Y, \hat{Y}) = \frac{1}{M^3} \sum_{x=0}^M \sum_{y=0}^M \sum_{z=0}^M (Y_{x,y,z} - \hat{Y}_{x,y,z})^2, \quad (8)$$

$$\mathcal{L}_{\text{3D-SSIM}}(Y, \hat{Y}) = 1 - \frac{(2\mu_Y \mu_{\hat{Y}} + c_1)/(2\sigma_{Y\hat{Y}} + c_2)}{(\mu_Y^2 + \mu_{\hat{Y}}^2 + c_1)(\sigma_Y^2 + \sigma_{\hat{Y}}^2 + c_2)}, \quad (9)$$

where Y is the GT of 3D volume and \hat{Y} is reconstructed 3D volume. Note that, in Eq. 8, M refers to the dimension of 3D volume; x , y and z are defined as the indexes of voxel coordinates, respectively. In Eq. 9, μ_Y and $\mu_{\hat{Y}}$ denote the mean of the voxels; σ_Y^2 and $\sigma_{\hat{Y}}^2$ are the variance of the voxels, respectively. In addition, $\sigma_{Y\hat{Y}}$ denotes the covariance between Y and \hat{Y} . We set $c_1 = 0.01$ and $c_2 = 0.03$ to avoid instability when the mean and variances are close to zero.

4 Experiments

4.1 Evaluation Protocol and Implementation Details

Dataset To have a fair comparison, we follow the protocols specified by [3] and [35] evaluate the proposed model. A subset of the ShapeNet dataset, which comprises 43,783 objects from 13 categories, is adopted. Each object is rendered from 24 different views. We resize the original images from 137×137 to 224×224 , and the uniform background color is applied before passing them to the network. For each 3D object category, three subsets with a ratio of 70:10:20 for training, validation, and testing are partitioned randomly. Note that all the 3D object samples are provided with a low-resolution voxel representation of $32 \times 32 \times 32$.

Evaluation Metrics Intersection over Union (IoU) and F-score are adopted to evaluate the reconstruction performance of the proposed model. A higher IoU score implies a better reconstruction. A higher F-score value suggests a better reconstruction. Therefore, the F-score@1% is adopted in this paper.

Implementation The proposed model is trained with the batch size 32 with view input images of size 224×224 and the dimension of 3D volume is set to $32 \times 32 \times 32$. For multi-view training, the number of input views is fixed to 8, which is the best obtained from the experiments. The full-fledged results are given in supplementary materials Section 6.1.

During training, the views are randomly sampled out of 24 views at each iteration, and the network is trained by an SGD optimizer with an initial learning rate of 0.01. The learning rate decay is used, and it is subsequently reduced by 10^{-1} for every 500 epochs. The minimum learning rate is defined as 1.0×10^{-4} . Our network is implemented by using the PyTorch toolkit [19], and it is performed by an NVidia Tesla V100. The source code is provided at GitHub[‡].

4.2 Result

Multi-view 3D Reconstruction The performance comparisons of 3D-C2FT against several benchmark models, namely 3D-R2N2 (RNN-based) [3], AttSets (RNN-based) [36], Pix2Vox++/F (CNN-based) [34], Pix2Vox++/A (CNN-based) [34], LegoFormer (transformer-based) [35], 3D-RETR (transformer-based) [23], VoIT+ (transformer-based) [28], and EVoIT (transformer-based) [28] are presented in this subsection. To have fair comparisons, we re-implemented several models such as AttSets, Pix2Vox++/F, Pix2Vox++/A, and 3D-RETR that were provided by the respective authors in order to follow the same experimental settings and protocols. Table 1 shows reconstruction performance over a different number of views in terms of IoU and F-score on the ShapeNet dataset.

Compared with RNN-based and CNN-based models, it can be observed that 3D-C2FT achieves the highest IoU and F-score with all views as tabulated in Table 1. As an illustration, we demonstrate several reconstruction instances from the ShapeNet dataset as shown in Fig. 2 and supplementary material Video01. As can be seen in Fig. 2, the objects reconstructed by the 3D-C2FT have complete and smoother surfaces, even compared with the LegoFormer. These results suggest that C2F patch and cube attention mechanisms play a crucial role in multi-view 3D object reconstruction.

Among transformer-based models (see Table 1), 3D-C2FT outperforms LegoFormer, VoIT+, EVoIT, and 3D-RETR significantly in almost all views in terms of IoU score and F-score, even under extreme cases where one view image is used as input. Note VoIT+ and EVoIT are published without code availability; thus, we only report the results for 4, 8, 12, and 20 views inputs [28]. Nevertheless, the results reveal the distinctive advantages of the C2F patch and cube attention

[‡]Source Code URL: <https://github.com/tiongleslie/3D-C2FT/>

Table 1. Performance comparisons of single and multi-view 3D reconstruction on ShapeNet with IoU and F-score. The best score for each view is written in bold. * refers to re-implemented models with the same settings of this study. The ‘-’ results of VoIT+ and EVoIT are not provided in [28] and source code is not publicly available.

Model	Number of views								
	1	2	3	4	5	8	12	18	20
<i>Metric: IoU</i>									
RNN & CNN-based									
3D-R2N2 (2016)	0.560	0.603	0.617	0.625	0.634	0.635	0.636	0.636	0.636
AttSets* (2020)	0.607	0.634	0.639	0.658	0.669	0.674	0.677	0.679	0.681
Pix2Vox++/F* (2019)	0.585	0.636	0.658	0.669	0.677	0.686	0.690	0.693	0.694
Pix2Vox++/A* (2020)	0.623	0.674	0.690	0.695	0.699	0.704	0.707	0.709	0.710
Transformer-based									
LegoFormer (2021)	0.519	0.644	0.679	0.694	0.703	0.713	0.717	0.719	0.721
VoIT+ (2021)	-	-	-	0.695	-	0.707	0.714	-	0.715
EVoIT (2021)	-	-	-	0.609	-	0.698	0.720	-	0.738
3D-RETR* (2021)	0.608	0.661	0.672	0.679	0.682	0.685	0.687	0.688	0.689
3D-C2FT	0.629	0.678	0.695	0.702	0.708	0.716	0.720	0.723	0.725
<i>Metric: F-score</i>									
RNN & CNN-based									
3D-R2N2 (2016)	0.351	0.368	0.372	0.378	0.382	0.383	0.382	0.382	0.383
AttSets* (2020)	0.358	0.360	0.370	0.379	0.384	0.386	0.389	0.403	0.408
Pix2Vox++/F* (2019)	0.341	0.367	0.388	0.398	0.405	0.417	0.429	0.430	0.432
Pix2Vox++/A* (2020)	0.365	0.419	0.435	0.443	0.447	0.452	0.457	0.460	0.461
Transformer-based									
LegoFormer (2021)	0.282	0.392	0.428	0.444	0.453	0.464	0.470	0.472	0.473
VoIT+ (2021)	-	-	-	0.451	-	0.464	0.469	-	0.474
EVoIT (2021)	-	-	-	0.358	-	0.448	0.475	-	0.497
3D-RETR* (2021)	0.355	0.412	0.425	0.432	0.436	0.440	0.442	0.443	0.444
3D-C2FT	0.371	0.424	0.443	0.452	0.458	0.468	0.476	0.477	0.479

over competing transformer models that utilize a plain single-scaled MSA mechanism in both encoder and decoder. However, 3D-C2FT underperforms EVoIT for 20 views. We speculate that the issue is associated with the attention layer in the decoder being less effective for a large number of input views.

Next, we also perform single-view 3D reconstruction against several benchmark models, namely 3D-R2N2, OGN (CNN-based) [25], Pix2Mesh (CNN-based) [29], OccNet (CNN-based) [15], Pix2Vox++F, Pix2Vox++A, LegoFormer, and 3D-RETR. For reference, we follow the protocols specified by [3] and [15] in this experiment. As shown in Table 2, 3D-C2FT attains the highest IoU value and F-score for most of the categories. By comparison, the C2FT attention mechanism clearly achieves better aggregation performance even with single-view images.

Performance Comparisons with Occlusion Images In this subsection, the reconstruction performance of 3D-C2FT and the competing models is evaluated with occluded images. Here, we use the testing set from the subset of the ShapeNet dataset by following the same protocols in the previous section. We add occlusion boxes with different sizes to the odd number ordered lists of 2D view images: 20×20 , 25×25 , 30×30 , 35×35 and 40×40 , which impeded the essential parts of the images intentionally. Readers are referred to the supplementary materials Section 6.2 for details.




Fig. 2. 3D object reconstruction using 8 views (only 1 are shown) for specific categories: *car*, *chair*, *lamp*, *rifle*, *table*, *display*, and *watercraft*.

Table 2. Performance comparisons of single-view 3D reconstruction results for each category on the ShapeNet dataset. * refers to re-implemented models with the same settings in this study. The best score for each category is written in bold.

Category	3D-R2N2	OCN	Pix2Mesh	OccNet	Pix2Vox++/F*	Pix2Vox++/A*	3D-RETR*	3D-C2FT
Airplane	0.513	0.587	0.420	0.571	0.533	0.652	0.647	0.640
Bench	0.421	0.481	0.323	0.485	0.454	0.547	0.547	0.549
Cabinet	0.716	0.729	0.664	0.733	0.726	0.760	0.712	0.761
Car	0.798	0.816	0.552	0.737	0.801	0.831	0.818	0.832
Chair	0.466	0.483	0.396	0.501	0.482	0.522	0.511	0.525
Display	0.468	0.502	0.490	0.471	0.463	0.499	0.434	0.489
Lamp	0.381	0.398	0.323	0.371	0.427	0.445	0.436	0.468
Loudspeaker	0.662	0.637	0.599	0.647	0.690	0.695	0.632	0.697
Rifle	0.544	0.593	0.474	0.402	0.528	0.583	0.567	0.583
Sofa	0.628	0.646	0.613	0.680	0.637	0.695	0.672	0.668
Table	0.513	0.536	0.395	0.506	0.541	0.568	0.536	0.571
Telephone	0.661	0.702	0.661	0.720	0.702	0.720	0.624	0.720
Watercraft	0.513	0.632	0.397	0.530	0.514	0.570	0.547	0.558
Overall	0.560	0.596	0.480	0.571	0.585	0.623	0.607	0.629

As shown in Table 3, the proposed 3D-C2FT achieves the best IoU scores for all sizes of occlusion boxes. Among the benchmark models, LegoFormer performs second-best over various occlusion boxes. Unlike the LegoFormer that draws attention to the fine-grained features only to explore the correlation between multi-view images, 3D-C2FT focuses on both coarse and fine-grained features so that they can assist the model in exploiting global and local interactions within the occluded regions, which make the model more robust against predicting unknown elements in the multi-view images.

Table 3. Performance comparisons of 12-view 3D reconstruction on ShapeNet using different sizes of occlusion box. The best score for each size is highlighted in bold.

Model	Occlusion Box				
	20×20	25×25	30×30	35×35	40×40
<i>Metric: IoU</i>					
Pix2Vox++/F	0.672	0.661	0.651	0.636	0.621
Pix2Vox++/A	0.698	0.695	0.690	0.685	0.678
LegoFormer	0.701	0.699	0.695	0.691	0.686
3D-RETR	0.667	0.651	0.623	0.603	0.584
3D-C2FT	0.717	0.715	0.712	0.706	0.698
<i>Metric: F-score</i>					
Pix2Vox++/F	0.416	0.410	0.405	0.396	0.388
Pix2Vox++/A	0.451	0.449	0.446	0.442	0.437
LegoFormer	0.453	0.451	0.448	0.443	0.439
3D-RETR	0.414	0.392	0.361	0.340	0.319
3D-C2FT	0.471	0.468	0.465	0.459	0.451

Experiment on Multi-view Real-life Dataset This subsection presents the performance comparisons on Multi-view Real-life Dataset. All models were only trained with the ShapeNet dataset. To create our dataset, we randomly searched the images by following the 3D object categories of the ShapeNet database from the Google image search engine. After that, all the images were manually verified to ensure they were correctly labeled. However, collecting many different input views for real images is difficult. Therefore, we manually combined the same samples with similar images from different views. The dataset is designed as a testing set that only contains 60 samples across 13 categories with three cases:

- *Case I*: samples with at least 1 to 5 views;
- *Case II*: samples with at least 6 to 11 views;
- *Case III*: samples with at least 12 or above views;

The experiment is designed to verify the 3D objects based on human perception without the GT of 3D volume. In Fig. 3 and supplementary materials Section 6.3, we show the qualitative results for single and multi-view 3D reconstruction. It is observed that 3D-C2FT performs substantially better than the benchmark models in terms of the refined 3D volume and surface quality even in *Case I*. Interestingly, most benchmark models fail to reconstruct the 3D objects, or the surface is not accurately generated. This is likely because the benchmark models lack in utilizing multi-scale feature extraction mechanisms, which causes these models to fail to capture relevant components more richly.

Note that [35] advocates a part-by-part composition notion by means of fine-to-fine attention, which is opposed to our global-to-local attention notion. As depicted in Fig. 3, LegoFormer underperforms the proposed model. This could be due to LegoFormer does not utilize 3D structure information essential in interpreting global interaction between semantic components. As a result, the model ignores the importance of the structure and orientation information of the 3D objects in guiding the model to perform reconstruction in challenging environments. In summary, the advantage of the C2F attention mechanism is that it fully utilizes the coarse and fine-grained features to be paired and represents the specific information flow that benefits the 3D object reconstruction task.




Fig. 3. 3D object reconstruction for *Case I*, *Case II* and *Case III* with Multi-view Real-life dataset

Table 4. Comparisons on parameter sizes of competing models

Model	Param. (M)	Backbone	Speed per View (sec)
3D-R2N2	36.0	-	-
AttSets	53.1	-	0.011
Pix2Vox++/F	114.2	ResNet50	0.010
Pix2Vox++/A	96.3	ResNet50	0.011
LegoFormer	168.4	VGG16	0.013
3D-RETR	168.0	RETR	0.012
3D-C2FT	90.1	DenseNet121	0.013

Computational Cost Table 4 tabulates the number of parameters in various benchmark models. Note that the backbone refers to the pre-trained image embedding module. 3D-C2FT has a smaller number of parameters than most of the competing models. Although 3D-R2N2 has significantly smaller parameter sizes, the reconstruction performances are below par, as revealed in Table 1. Besides, we also compare reconstruction time (speed per view) in Table 4. As can be seen in this table, 3D-C2FT has a well-balanced of the best balance of speed (see Table 4) and performance (IoU score, see Table 1).

4.3 Ablation Study

In this section, all the experiments are conducted with the ShapeNet dataset described in Section 4.1.

Visualization on C2F Attention To better understand the benefit of C2F attention, we present visualization results of three C2F patch attention blocks separately, labeled as $C2F_1$, $C2F_2$ and $C2F_3$ with attention rollout [1]. Fig. 4 depicts the attention maps of each stage of the C2F patch attention block, which interpret and focus semantically relevant regions that result in a coarse to fine-grained manner.




Fig. 4. Visualization of attention maps for each C2F patch attention block.

As shown in Fig. 4, the $C2F_1$ attention is drawn to the central parts of the given images, which represent coarse or global features of the 3D objects. Similarly, the $C2F_2$ shows attention is drawn toward the edges of the objects, which can assist the model to distinguish the view orientations for reconstruction. The final block ($C2F_3$) suggests that the attention is drawn to the specific parts of the objects, which can be considered as fine-grained features. The results indicate that the proposed coarse-to-fine mechanism pays specific attention to the relevant components in the 3D objects, which is essential for the decoder to predict and reconstruct accurate 3D volume.

Reconstruction with and without Refiner In Fig. 5a and 5b, we evaluate the influence of the refiner on 3D reconstruction results using 3D-C2FT and 3D-C2FT without refiner (3D-C2FT-WR). We observe that 3D-C2FT can significantly achieve the best IoU score and F-score in all views. We also show several qualitative 3D reconstruction results in Fig. 5c. Without refiner, the decoder generates defective 3D surfaces due to a lack of drawing attention to rectifying wrongly reconstructed parts within the 3D objects. Therefore, our analysis indicates that the refiner plays a crucial role in improving the 3D surface quality and can remove the noise of the reconstructed 3D volume from the decoder.

Loss Functions We also show the performance of 3D-C2FT trained with different loss functions, namely \mathcal{L}_{total} , \mathcal{L}_{MSE} and $\mathcal{L}_{3D-SSIM}$, as shown in Table 5. It is noticed that 3D-C2FT trained with \mathcal{L}_{total} achieves the best IoU and F-score compared to either \mathcal{L}_{MSE} or $\mathcal{L}_{3D-SSIM}$ alone. \mathcal{L}_{MSE} minimizes the voxel grid error between the predicted and GT of 3D objects. In contrast, $\mathcal{L}_{3D-SSIM}$ reduces the structural difference, which benefits in improving the surface of predicted 3D objects. As a result, both loss functions drive the model training towards more reliable and better surface quality reconstruction.

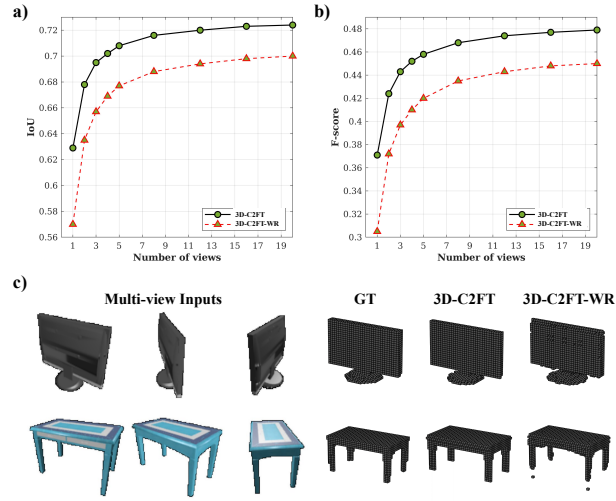


Fig. 5. Performance comparisons of multi-view 3D reconstruction results on 3D-C2FT and 3D-C2FT-WR. a) IoU score; b) F-score; c) 3D reconstruction results using 8 views (only 3 are shown) between 3D-C2FT and 3D-C2FT-WR.

Table 5. Performance comparisons of multi-view 3D reconstruction results on 3D-C2FT using different loss functions. The best score for each view is written in bold.

Loss function	Number of views								
	1	2	3	4	5	8	12	18	20
<i>Metric: IoU</i>									
\mathcal{L}_{MSE}	0.623	0.673	0.683	0.696	0.701	0.708	0.713	0.716	0.716
$\mathcal{L}_{\text{3D-SSIM}}$	0.617	0.668	0.683	0.691	0.696	0.705	0.709	0.712	0.712
$\mathcal{L}_{\text{total}}$	0.629	0.678	0.695	0.702	0.708	0.716	0.720	0.723	0.724
<i>Metric: F-score</i>									
\mathcal{L}_{MSE}	0.369	0.421	0.437	0.451	0.457	0.466	0.471	0.474	0.476
$\mathcal{L}_{\text{3D-SSIM}}$	0.365	0.418	0.437	0.446	0.452	0.462	0.467	0.470	0.472
$\mathcal{L}_{\text{total}}$	0.371	0.424	0.443	0.452	0.458	0.468	0.475	0.477	0.479

5 Conclusion

This paper proposed a multi-view 3D reconstruction model that employs a C2F patch attention mechanism in the 2D encoder, a 3D decoder, and a C2F cube attention mechanism in the refiner. Our experiments showed that 3D-C2FT could achieve significant results compared to the several competing models on the ShapeNet dataset. Further study with the Multi-view Real-life dataset showed that 3D-C2FT is far more robust than other models. For future works, we will consider improving the attention mechanism and the scaling of 3D-C2FT. Finally, we will explore other mechanisms to enhance the 3D reconstruction task, which we hope will facilitate the practical usage of robotics, historical artifacts, and other related domains.

References

1. Abnar, S., Zuidema, W.: Quantifying attention flow in transformers. arXiv e-prints (2020), <https://arxiv.org/abs/2005.00928>
2. Burchfiel, B., Konidaris, G.: Bayesian eigenobjects: A unified framework for 3D robot perception. In: *Robotics: Science and Systems*. vol. 13 (2017)
3. Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3D-R2N2: A unified approach for single and multi-view 3D object reconstruction. In: *European Conference on Computer Vision (ECCV)*. pp. 628–644 (2016)
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations (ICLR)* (2021)
5. Gao, Y., Luo, J., Qiu, H., Wu, B.: Survey of structure from motion. In: *Proceedings of 2014 International Conference on Cloud Computing and Internet of Things*. pp. 72–76 (2014)
6. Groen, I.I.A., Baker, C.I.: Previews scenes in the human brain: Comparing 2D versus 3D representations. *Neuron* **101**(1), 8–10 (2019)
7. Han, X.F., Laga, H., Bennamoun, M.: Image-based 3D object reconstruction: State-of-the-art and trends in the deep learning era. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(5), 1578–1604 (2021)
8. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4700–4708 (2017)
9. Jabłoński, S., Martyn, T.: Real-time voxel rendering algorithm based on screen space billboard voxel buffer with sparse lookup textures. In: *24th Conference on Computer Graphics, Visualization and Computer Vision*. pp. 27–36 (2016)
10. Kanzler, M., Rautenhaus, M., Westermann, R.: A voxel-based rendering pipeline for large 3d line sets. *IEEE Transactions on Visualization and Computer Graphics* **25**(7), 2378–2391 (2019)
11. Kar, A., Häne, C., Malik, J.: Learning a multi-view stereo machine. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS)*. p. 364–375. Curran Associates, Inc. (2017)
12. Kargas, A., Loumos, G., Varoutas, D.: Using different ways of 3D reconstruction of historical cities for gaming purposes: The case study of Nafplio. *Heritage* **2**(3), 1799–1811 (2019)
13. Kniaz, V.V., Knyaz, V.A., Remondino, F., Bordodymov, A., Moshkantsev, P.: Image-to-voxel model translation for 3D scene reconstruction and segmentation. In: *European Conference on Computer Vision (ECCV)*. pp. 105–124 (2020)
14. Malik, J., Abdelaziz, I., Elhayek, A., Shimada, S., Ali, S.A., Golyanik, V., Theobalt, C., Stricker, D.: Handvoxnet: Deep voxel-based network for 3d hand shape and pose estimation from a single depth map. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7111–7120 (2020)
15. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
16. Nabil, M., Saleh, F.: 3D reconstruction from images for museum artefacts: A comparative study. In: *International Conference on Virtual Systems and Multimedia (VSMM)*. pp. 257–260. IEEE (2014)

17. Nguyen, T.Q., Salazar, J.: Transformers without tears: Improving the normalization of self-attention. In: Proceedings of the 16th International Conference on Spoken Language Translation. Hong Kong (2019)
18. Park, N., Kim, S.: How do vision transformers work? In: International Conference on Learning Representations (ICLR) (2022)
19. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS). pp. 8024–8035 (2019)
20. Păvăloiu, I.B., Vasilăţeanu, A., Goga, N., Marin, I., Ilie, C., Ungar, A., Pătraşcu, I.: 3D dental reconstruction from CBCT data. In: International Symposium on Fundamentals of Electrical Engineering (ISFEE). pp. 4–9 (2014)
21. Roointan, S., Tavakolian, P., Sivagurunathan, K.S., Floryan, M., Mandelis, A., Abrams, S.H.: 3D dental subsurface imaging using enhanced truncated correlation-photothermal coherence tomography. *Scientific Reports* **9**(1), 1–12 (2019)
22. Shi, Q., Li, C., Wang, C., Luo, H., Huang, Q., Fukuda, T.: Design and implementation of an omnidirectional vision system for robot perception. *Mechatronics* **41**, 58–66 (2017)
23. Shi, Z., Meng, Z., Xing, Y., Ma, Y., Wattenhofer, R.: 3D-RETR: End-to-end single and multi-view 3D reconstruction with transformers. In: British Machine Vision Conference (BMVC). pp. 1–14 (2021)
24. Silveira, G., Malis, E., Rives, P.: An efficient direct approach to visual SLAM. *IEEE Transactions on Robotics* **24**(5), 969–979 (2008)
25. Tatarchenko, M., Dosovitskiy, A., Brox, T.: Octree generating networks: Efficient convolutional architectures for high-resolution 3D outputs. In: IEEE International Conference on Computer Vision (ICCV). pp. 2088–2096 (2017)
26. Tron, R., Vidal, R.: Distributed 3-D localization of camera sensor networks from 2-D image Measurements. *IEEE Transactions on Automatic Control* **59**(12), 3325–3340 (2014)
27. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS). vol. 30, p. 6000–6010 (2017)
28. Wang, D., Cui, X., Chen, X., Zou, Z., Shi, T., Salcudean, S., Wang, Z.J., Ward, R.: Multi-view 3D reconstruction with transformer. In: International Conference on Computer Vision (ICCV). pp. 5722–5731 (2021)
29. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2Mesh: Generating 3D mesh models from single RGB images. In: European Conference on Computer Vision (ECCV) (2018)
30. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004)
31. Wilson, K., Snavely, N.: Robust global translations with 1DSfM. In: European Conference on Computer Vision (ECCV). pp. 61–75 (2014)
32. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3D ShapeNets: A deep representation for volumetric shapes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1912–1920 (2015)

33. Xie, H., Yao, H., Sun, X., Zhou, S., Zhang, S.: Pix2Vox: Context-aware 3D reconstruction from single and multi-view images. In: IEEE International Conference on Computer Vision (ICCV). pp. 2690–2698 (2019)
34. Xie, H., Yao, H., Zhang, S., Zhou, S., Sun, W.: Pix2Vox++: Multi-scale context-aware 3D object reconstruction from single and multiple images. *International Journal of Computer Vision* **128**(12), 2919–2935 (2020)
35. Yagubbayli, F., Tonioni, A., Tombari, F.: LegoFormer: Transformers for block-by-block multi-view 3D reconstruction. *arXiv e-prints* (2021), <http://arxiv.org/abs/2106.12102>
36. Yang, B., Wang, S., Markham, A., Trigoni, N.: Robust attentional aggregation of deep feature sets for multi-view 3D reconstruction. *International Journal of Computer Vision* **128**(1), 53–73 (2020)