

DOCTR: Disentangled Object-Centric Transformer for Point Scene Understanding

Xiaoxuan Yu^{1*}, Hao Wang¹, Weiming Li¹, Qiang Wang¹, Soonyong Cho², Younghun Sung²

¹Samsung Research China – Beijing

²Samsung Advanced Institute of Technology

xiaoxuan1.yu@samsung.com, hao1.wang@samsung.com, weiming.li@samsung.com, qiang.w@samsung.com, soonyong.cho@samsung.com, younghun.sung@samsung.com

Abstract

Point scene understanding is a challenging task to process real-world scene point cloud, which aims at segmenting each object, estimating its pose, and reconstructing its mesh simultaneously. Recent state-of-the-art method first segments each object and then processes them independently with multiple stages for the different sub-tasks. This leads to a complex pipeline to optimize and makes it hard to leverage the relationship constraints between multiple objects. In this work, we propose a novel Disentangled Object-Centric TRansformer (DOCTR) that explores object-centric representation to facilitate learning with multiple objects for the multiple sub-tasks in a unified manner. Each object is represented as a query, and a Transformer decoder is adapted to iteratively optimize all the queries involving their relationship. In particular, we introduce a semantic-geometry disentangled query (SGDQ) design that enables the query features to attend separately to semantic information and geometric information relevant to the corresponding sub-tasks. A hybrid bipartite matching module is employed to well use the supervisions from all the sub-tasks during training. Qualitative and quantitative experimental results demonstrate that our method achieves state-of-the-art performance on the challenging ScanNet dataset. Code is available at <https://github.com/SAITPublic/DOCTR>.

Introduction

Understanding 3D scene is important for various spatial modeling and interaction applications such as augmented reality (AR), autonomous driving, and robotics. With the progress of 3D deep learning, recent real-world applications seek more comprehensive understanding of rich and detailed attributes for all objects of interests in a scene. This paper addresses a recent task of point scene understanding (Nie et al. 2021; Tang et al. 2022) that simultaneously involves several sub-tasks including object classification, instance segmentation, pose estimation, and object mesh reconstruction. The input is point cloud of real-world scene that is obtained by 3D scanning or reconstruction. Such data often contains noisy and missing portions due to occlusions or sensor limitations, which make the task even more challenging.

*Corresponding author




Figure 1: Point Scene Understanding. With an incomplete point cloud scene as input, our method learns to segment each object instance and reconstruct its complete mesh. Compared to the recent DIMR (Tang et al. 2022), our proposed DOCTR has a compact pipeline and achieves more accurate results in challenging cases such as when multiple objects are in close proximity.

For the point scene understanding task, an early method RfD-Net(Nie et al. 2021) is among the first to learn object meshes at semantic-instance level directly from points. RfD-Net proposes a reconstruction-from-detection framework that enables identification and reconstruction of object meshes at a high resolution. It confirms that object recognition and reconstruction are mutually reinforcing tasks. Recently, DIMR (Tang et al. 2022) finds that the reconstruction-from-detection pipeline tends to fail when reconstructing high fidelity objects. DIMR shows that replacing the detection backbone in RfD-Net by a segmentation backbone and disentangling instance completion and mesh generation contribute to improve object reconstruction quality. These designs empower DIMR to achieve state-of-the-art (SOTA) performance. However, DIMR still relies on geometric clustering for object segmentation that requires tuning of hyper-parameters such as the radius for clustering. Meanwhile, after object segmentation, DIMR treats each object instance independently to estimate its pose, shape latent code, and reconstruct its mesh. DIMR’s pipeline is complex to jointly optimize the segmentation and other sub-tasks. Further, the separate processing for each object makes it hard to leverage the relationship between multiple objects. In test with real-world scenes, DIMR is prone to artifacts for

multiple nearby objects as shown in Figure 1.

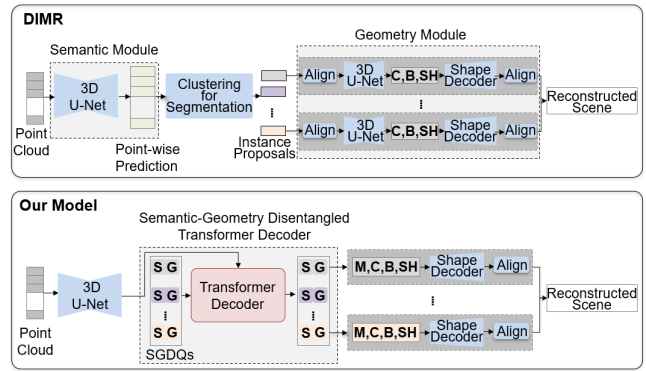
To solve the above issues, our work is inspired by object-centric learning (Carion et al. 2020; Locatello et al. 2020), which proposes to learn efficient representations of a scene by decomposing it into objects, thus is natural to explore the object relationship. Recently, object-centric learning methods are introduced into 3D point cloud segmentation task (Schult et al. 2023; Sun et al. 2023) and show promising performance. In this paper, we attempt to introduce an object-centric learning based method for the point scene understanding task. Specifically, each object is represented as a query, and a Transformer decoder is adapted to iteratively optimize all the queries involving their relationship.

However, applying an object-centric framework to the point scene understanding task is not straight forward. In existing methods such as RfD-Net and DIMR, after obtaining object proposals, each object’s point cloud is transformed to object’s canonical coordinate frame using the initial pose estimation result. This 3D alignment is important to improve subsequent geometric sub-tasks such as object pose refinement, shape completion, and mesh reconstruction. However, object-centric Transformer uses a unified representation for each object, which is not meaningful to perform any 3D coordinate frame alignment. This makes it difficult to learn object geometric information such as pose and shape. To address this issue, we propose a semantic-geometry disentangled query (SGDQ) design. Different from the origin query, SGDQ disentangles each query to a semantic part and a geometry part. The semantic part is supervised by semantic sub-tasks such as object classification and semantic instance segmentation. The geometric part is supervised by the geometric sub-tasks such as pose estimation and mesh reconstruction. With the distinct focuses on semantic and geometric sub-tasks, our SGDQ empowers disentangled learning for the task-specific features.

Following the above, we propose our method as Disentangled Object-Centric TRansformer (DOCTR). Our DOCTR consists of a backbone, a disentangled Transformer decoder with the SGDQ design, a prediction head for the sub-tasks of point scene understanding, and a shape decoder. It facilitates learning with multiple objects for the multiple sub-tasks in a unified manner. Compared to DIMR, our DOCTR has a much more compact pipeline as shown in Figure 2. Meanwhile, different from Mask3D (Schult et al. 2023), our SGDQ design allows learning for multiple different sub-tasks that relates to either semantic or geometric attributes of the scene objects. For training our DOCTR model, we employ a hybrid bipartite matching strategy during the matching process between ground truths and SGDQs. We also propose a mask-enhanced box refinement module that leverages segmentation to improve pose estimation. Extensive experiments are performed with the real-world large-scale ScanNet dataset in comparison to the SOTA methods.

The contributions of this paper are as follows:

- As far as we know, our DOCTR is the first to introduce an object-centric Transformer-based network for the point scene understanding task that allows learning with multiple objects and multiple sub-tasks in a unified manner.



S: Semantic; G: Geometry; M: Mask; C: Class; B: Box; SH: SHape.

Figure 2: Comparison of the pipelines between DIMR (Tang et al. 2022) and our DOCTR. DIMR first segments each object instance by clustering and then processes them independently with multiple stages for the different sub-tasks. Our DOCTR pipeline is much more compact that represents each object instance as a query and optimize for multiple objects and multiple sub-tasks in a unified manner.

- We propose semantic-geometry disentangled query (SGDQ) that enables the DOCTR network to extract semantic and geometric features for the different sub-tasks to effectively use disentangled representations.
- Qualitative and quantitative experimental results on the challenging ScanNet dataset show that our proposed method achieves superior performance than previous SOTA methods, especially for the challenging cases such as cluttered scenes with nearby objects.

Related Work

Point scene understanding. Many real-world applications require understanding both semantic and geometric attributes of 3D scenes. Among different types of inputs, point scene understanding is related to 3D deep learning using only scene’s point cloud as input. While point cloud is commonly used for object-level tasks such as object completion (Huang, Chen, and Li 2019; Wang, Ang Jr, and Lee 2020; Xie et al. 2020; Yuan et al. 2018), less work (Zhong and Zeng 2020; Rist et al. 2020, 2021) use point cloud with neural networks for scene understanding. Early works (Garbade et al. 2019; Roldao, de Charette, and Verroust-Blondet 2020; Wu et al. 2020; Yan et al. 2021) often use voxelized input and then predict the semantic label of each voxel in both visible and occluded regions. They aim to jointly estimate the complete geometry and semantic labels from partial input. Working with voxels, computational cost and resolution need to be balanced which limits to reconstruct high fidelity object meshes.

To solve such issues, recent point scene understanding methods explore the relationship of semantic and geometric information, containing object localization and reconstruction. RfD-Net (Nie et al. 2021) is among the first to propose a reconstruction-from-detection framework to detect

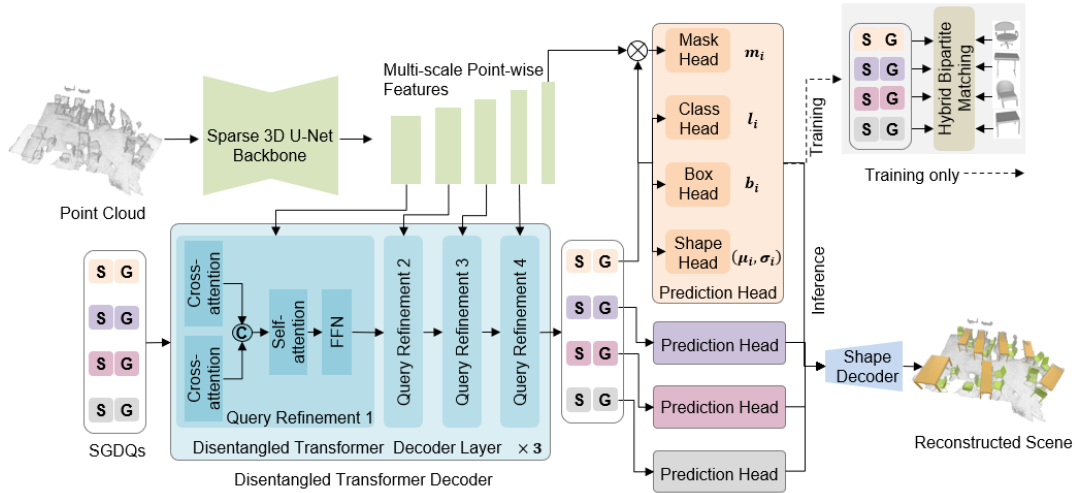


Figure 3: Our proposed DOCTR pipeline consists of a backbone, a disentangled Transformer decoder (DTD), a prediction head, and a shape decoder. We propose semantic-geometry disentangled queries (SGDQs) to represent scene objects. The DTD is trained to attend the SGDQs to multi-scale point-wise features extracted from the sparse 3D U-Net backbone. In inference, each object SGDQ is passed to the prediction head to predict the object’s mask, class, box (pose), and shape code. The shape code is input to the shape decoder to reconstruct the object’s complete mesh, which is then aligned by the estimated pose to the scene.

objects and generate complete object meshes. It shows that the tasks of object detection and object completion are complementary. Current state-of-the-art is achieved by DIMR (Tang et al. 2022) that proposes to use segmentation network instead of detection network to improve the accuracy of object recognition. DIMR also disentangles the tasks of object completion and mesh generation, mitigating the ambiguity of learning complete shapes from incomplete point cloud observations. In contrast to existing methods, our proposed approach introduces an object-centric Transformer-based network for the point scene understanding task that enables simultaneous learning with multiple objects and multiple sub-tasks in a unified manner.

Object-centric learning. Learning with object-centric representation is a promising approach to understanding real-world complex scenes with multiple objects. With the help of set prediction formulation, DETR (Carion et al. 2020) proposes to learn object-centric representations with the Transformer architecture for 2D object detection. In (Locatello et al. 2020), a fully-unsupervised approach based on slot attention for object-centric learning is proposed. Object-centric learning has achieved great success in 2D tasks (Zhu et al. 2020; Liu et al. 2022; Zhang et al. 2022). Recently, there have been notable advancements in applying object-centric learning to 3D point cloud segmentation or detection tasks (Schult et al. 2023; Sun et al. 2023; Zhu et al. 2023). These methods have demonstrated performance that rivals or even surpasses the current SOTA approaches in these fields. Different from the recent object-centric models for 3D tasks such as Mask3D (Schult et al. 2023), our work proposes a semantic-geometry disentangled query design that allows object-centric Transformer to deal with multiple different se-

manic and geometric sub-tasks simultaneously.

Methods

An overview of our posed DOCTR pipeline is illustrated in Figure 3. Our DOCTR consists of a sparse 3D U-Net backbone, a disentangled Transformer decoder (DTD), a prediction head, and a shape decoder. All objects are represented by our designed SGDQs and each SGDQ corresponds to an object instance. During training, the Transformer decoder is trained to attend the SGDQs to the multi-scale point-wise features extracted by the backbone. To supervise the training, our proposed hybrid bipartite matching is used to assign either a ‘no object’ class or an object’s ground-truth attributes to the corresponding object SGDQ. In inference, for each object’s SGDQ, the prediction head predicts its mask, class, pose (box), and a shape code. The predicted shape code is decoded by a shape decoder to reconstruct the object’s complete mesh. Then the mesh is transformed by the estimated pose to align to the scene’s coordinate frame as the object’s reconstruction. All the reconstructed objects comprise the final reconstructed scene. Next we’ll describe the details for each component of the pipeline.

Sparse 3D Backbone

From the input point cloud $\mathbf{P} \in \mathbb{R}^{N \times 3}$, we employ a sparse 3D U-Net (Choy, Gwak, and Savarese 2019) to extract multi-scale point-wise features $\{\mathbf{F}_l\} \in \mathbb{R}^{N_l \times D_l^f}$, where $l = 0, \dots, L$ is scale level, N_l and D_l^f are spatial dimension and feature dimension of \mathbf{F}_l respectively. Especially, the spatial dimension of \mathbf{F}_0 and \mathbf{P} are the same. Different from DIMR, our method does not add any additional semantic branch or offset branch for instance segmentation.




Figure 4: Disentangled Transformer Decoder Layer

Disentangled Transformer Decoder

Point scene understanding comprises sub-tasks involving different types of information. Specifically, the sub-tasks of instance segmentation and classification are more related to semantic information. Meanwhile, the sub-tasks of pose estimation and shape reconstruction are more related to geometry information. To ensure optimal learning for all the sub-tasks, we propose disentangled Transformer decoder (DTD), which aims to decouple the learning of semantic features and geometric features. This allows our model to learn the most relevant information for each sub-task.

Semantic-geometry disentangled query. As a key design of the DTD, we pose semantic-geometry disentangled query (SGDQ). All scene objects are represented by a set of SGDQs and each SGDQ corresponds to one object instance. Unlike the query used in the vanilla object-centric framework, the feature of each SGDQ is divided into a semantic part and a geometric part. In training, the disentanglement of semantic and geometric parts is achieved through separate supervision with the different sub-task prediction heads. The semantic part is supervised by the sub-tasks of object classification and semantic instance segmentation. The geometric part is supervised by the sub-tasks of pose estimation and mesh reconstruction. We choose the same fixed number of SGDQs for all scenes, whose value is often much larger than the number of objects in the scene. SGDQs are randomly initialized and iteratively learned to be representations of objects in the scene.

Disentangled transformer decoder layer. Our DTD consists of multiple DTD layers, and each DTD layer contains multiple query refinement modules. Each query refinement module corresponds to a different level of the multi-scale features $\{\mathbf{F}_l\}$ for $l = 1, \dots, L$ as shown in Figure 4. The semantic parts and geometric parts of SGDQs are separately refined by cross-attending to the scene’s point-wise features and fused at the object level through self-attention. Starting from M randomly initialized SGDQs, through the series

of query refinement modules, the SGDQs are optimized in a coarse-to-fine manner progressively. In our model, we use three DTD layers and four levels ($L = 4$) in each DTD layer.

Different from the original Transformer decoder layer, our DTD layer is designed to learn disentangled semantic and geometric representations of SGDQs. At each level l in a decoder layer, each SGDQ competes for attending to the features \mathbf{F}_l from one object instance through attention mechanism. We denote \mathbf{Q}^l as an input SGDQ of level l , and \mathbf{Q}^{l+1} as its output SGDQ. Each SGDQ is composed of a semantic part and a geometric part, denoted as $\mathbf{Q}^l = (\mathbf{Q}_s^l, \mathbf{Q}_g^l)$. In each query refinement module of DTD layer, we firstly utilize two cross-attention modules to learn the semantic and geometric information separately from the multi-scale point-wise features at level l . In the semantic aware cross-attention module, we first apply learnable linear transformations, denoted as $\mathbf{q}_s, \mathbf{k}_s, \mathbf{v}_s$ respectively, to features \mathbf{F}_p and semantic part queries \mathbf{Q}_s^l to obtaining query $\hat{\mathbf{Q}}_s^l \in \mathbb{R}^{M \times D}$, key $\mathbf{K}_s^l \in \mathbb{R}^{N \times D}$ and value $\mathbf{V}_s^l \in \mathbb{R}^{N \times D}$. Then matrix multiplication between $\hat{\mathbf{Q}}_s^l$ and \mathbf{K}_s^l gives the correlation matrix between M queries and N_l input features. Here we use the masked cross-attention where each query only attends to features within binary instance mask \mathbf{m} predicted by the mask prediction head using \mathbf{Q}_s^{l-1} , which can be formulated as:

$$\hat{\mathbf{Q}}_s^{l+1} = \text{softmax}(\underbrace{\hat{\mathbf{Q}}_s^l}_{\mathbf{q}_s(\mathbf{Q}_s^l)} \cdot \underbrace{\mathbf{K}_s^l}_{\tilde{\mathbf{m}}(\mathbf{F}, \mathbf{m})})^T \underbrace{\mathbf{V}_s^l}_{\mathbf{v}_s(\mathbf{F})}), \quad (1)$$

where

$$\tilde{\mathbf{m}}(\mathbf{F}, \mathbf{m})_{j,k} = \begin{cases} \mathbf{F}_{j,k}, & \mathbf{m}_{j,k} = 0, \\ 0, & \mathbf{m}_{j,k} = 1. \end{cases} \quad (2)$$

The same attention method is used in the geometric part \mathbf{Q}_g^l :

$$\hat{\mathbf{Q}}_g^{l+1} = \text{softmax}(\mathbf{q}_g(\mathbf{Q}_g^l) \cdot \mathbf{k}_g(\tilde{\mathbf{m}}(\mathbf{F}, \mathbf{m}))^T) \mathbf{v}_g(\mathbf{F}). \quad (3)$$

These two cross-attentions empower the queries to extract semantic and geometric information from the point-wise features. Then, the semantic and geometric features will be concatenated and fed into a self-attention module and a Feed Forward Network (FFN) to gather context information:

$$\begin{aligned} \hat{\mathbf{Q}}^{l+1} &= \text{self-attention}((\hat{\mathbf{Q}}_s^{l+1}, \hat{\mathbf{Q}}_g^{l+1})), \\ \mathbf{Q}^{l+1} &= \text{FFN}(\hat{\mathbf{Q}}^{l+1}). \end{aligned} \quad (4)$$

Prediction head. After DTD, each SGDQ $[\mathbf{Q}]_i$ for $i = 0, \dots, M - 1$ attends to the features of a specific object instance. These SGDQs are then fed into a prediction head whose weights are shared across queries. The prediction head consists of four MLP-based task heads, including mask head, class head, box head, and shape head. The mask head takes semantic part $[\mathbf{Q}_s]_i$ and the point-wise features from the sparse 3D U-Net as input. It maps $[\mathbf{Q}_s]_i$ through an MLP $f_{\text{mask}}(\cdot)$, and outputs object binary mask \mathbf{m}_i by computing dot products between the mapped query features and point-wise features \mathbf{F}_0 with a threshold at 0.5:

$$[\mathbf{m}]_i = \{\sigma(\mathbf{F}_0 \cdot f_{\text{mask}}([\mathbf{Q}_s]_i)^T) > 0.5 \in \{0, 1\}^N\}. \quad (5)$$

The class head also takes the semantic part $[\mathbf{Q}_s]_i$ as input, and outputs classification logits $\mathbf{l}_i \in \mathbb{R}^C$ for each instance, where C is the total class number.

The box head takes the geometric part $[\mathbf{Q}_g]_i$ as input, and outputs box vector \mathbf{b}_i , including rotation angle $r_i \in [-\pi, \pi]$ along the z-axis following (Nie et al. 2021), center of bounding box $\mathbf{c}_i = (x_i, y_i, z_i)$ and size of bounding box $\mathbf{s}_i = (s_i^x, s_i^y, s_i^z)$ and IoU score $\text{score}_i^{\text{IoU}}$, denoted as $\mathbf{b}_i = (r_i, \mathbf{c}_i, \mathbf{s}_i, \text{score}_i^{\text{IoU}})$.

For shape completion, we assume that the complete shape is sampled from a latent Gaussian distribution and learn it through the reparameterization trick (Kingma and Welling 2013). The shape head takes the geometric part $[\mathbf{Q}_g]_i$ as input and regresses mean and standard deviation $\boldsymbol{\mu}_i, \boldsymbol{\sigma}_i \in \mathbb{R}^{D_{\text{shape}}}$ for the shape latent Gaussian distribution, where D_{shape} is the latent shape code dimension. We decode the latent code $\mathbf{z}_i \sim \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\sigma}_i)$ to mesh only at inference time through a pretrained shape decoder of BSP-Net (Chen, Tagliasacchi, and Zhang 2021, 2020). Details of the shape decoder can be found in (Tang et al. 2022).

Training Design

Hybrid bipartite matching. During training, we utilize ground truth annotations for each object instance, including the mask, class, box, and mesh. For determining the correspondence between the queries and the ground truths, we use bipartite matching following DETR (Carion et al. 2020). For each SGDQ, a pair of mask and box predicted by two task heads may be inconsistent with each other. To obtain accurate and consistent matching, we employ a mixed cost of predicted mask, box, and class when performing bipartite matching. We construct a cost matrix \mathcal{C} defined as follows:

$$\mathcal{C} = \lambda_1 \mathcal{L}_{\text{dice}}^{\text{mask}} + \lambda_2 \mathcal{L}_{\text{IoU}}^{\text{box}} + \lambda_3 \mathcal{L}^{\text{class}} = \{\mathcal{C}_{jl}\}, \quad (6)$$

in which \mathcal{C}_{jl} is the similarity of the j -th query and the l -th ground truth. We set the weights to $\lambda_1 = 5$ and $\lambda_2 = \lambda_3 = 2$. The meaning of each component will be explained in the upcoming discussion of training losses. After matching, those queries that are not assigned to any ground truth are given to the ‘no object’ category and only participate in the classification task. Benefiting from the one-to-one matching mechanism, each object instance in the scene will be represented by only one query, which allows our network to directly perform sparse prediction.

Training losses. The loss \mathcal{L} consists of semantic loss \mathcal{L}^{sem} and geometric loss \mathcal{L}^{geo} .

Semantic tasks are supervised by :

$$\mathcal{L}^{\text{sem}} = \mathcal{L}_{\text{BCE}}^{\text{mask}} + \mathcal{L}_{\text{dice}}^{\text{mask}} + \mathcal{L}^{\text{class}}, \quad (7)$$

in which $\mathcal{L}_{\text{BCE}}^{\text{mask}}$ is the binary cross-entropy loss for mask, $\mathcal{L}_{\text{dice}}^{\text{mask}}$ is the Dice loss for mask, and $\mathcal{L}^{\text{class}}$ is the cross-entropy loss for object-wise semantic label.

Geometry tasks are supervised by :

$$\mathcal{L}^{\text{geo}} = \mathcal{L}_{\text{center}}^{\text{box}} + \mathcal{L}_{\text{size}}^{\text{box}} + \mathcal{L}_{\text{angle}}^{\text{box}} + \mathcal{L}_{\text{IoU}}^{\text{box}} + \mathcal{L}^{\text{shape}}, \quad (8)$$

where $\mathcal{L}_{\text{center}}^{\text{box}}$ is the Huber loss for bounding box center, $\mathcal{L}_{\text{size}}^{\text{box}}$ is the Huber loss for bounding box size, $\mathcal{L}_{\text{angle}}^{\text{box}}$ contains the

cross-entropy loss for angle label and Huber loss for residual angle, $\mathcal{L}_{\text{IoU}}^{\text{box}}$ is the 3D GIoU loss for box regression, $\mathcal{L}^{\text{shape}}$ is the latent shape distribution loss for predicted shape latent distribution. The ground-truth latent shape distributions are obtained from a shape encoder of BSP-Net pretrained with ground-truth meshes.

Mask Enhanced Box Refinement

For an object instance i , it is observed that when the network predicts its mask m_i accurately, the box $\mathbf{B}_i^m = (\mathbf{c}_i^p, \mathbf{s}_i^p)$ calculated from the object’s point cloud tends to be more accurate than the network’s predicted box $\mathbf{B}_i^p = (\mathbf{c}_i^m, \mathbf{s}_i^m)$. We consider the mask m_i is ‘‘sufficiently accurate’’, when the distance $d(\mathbf{B}_i^p, \mathbf{B}_i^m) = |\mathbf{s}_i^p - \mathbf{s}_i^m|_{\infty}$ between \mathbf{B}_i^m and \mathbf{B}_i^p is less than $d_0 = 0.1\text{m}$. For each object point cloud \mathbf{o}_i , we apply the estimated rotation angle r_i to transform it to a canonical coordinate system for mask enhanced box refinement (MEBR). For MEBR, denote the box calculated from the object point cloud \mathbf{o}_c in canonical coordinate system as \mathbf{B}_i^m and the box predicted by network as \mathbf{B}_i^p , then the refined box \mathbf{B}_i^f is as follows:

$$\mathbf{B}_i^f = \begin{cases} \mathbf{B}_i^p, & \text{if } d(\mathbf{B}_i^p, \mathbf{B}_i^m) \leq d_0, \\ \mathbf{B}_i^m, & \text{otherwise.} \end{cases} \quad (9)$$

Experiment

Experimental Settings

Datasets. As previous works for point scene understanding, our experiments are conducted with the ScanNet V2 dataset (Dai et al. 2017) (abbreviated as ScanNet for short). The ScanNet dataset is a richly-annotated dataset that contains point clouds of 1513 real world indoor scenes labeled at instance level. The Scan2CAD dataset (Avetisyan et al. 2019) provides additional geometric instance annotations by aligning CAD models from ShapeNet (Chang et al. 2015) to instances in ScanNet. To deal with data inconsistency issue, DIMR (Tang et al. 2022) relabels the datasets based on a compatible label system including 8 object categories. We follow exactly the same data split and pre-processing method as DIMR. A minor exception is when searching for the best matches between CAD models in Scan2CAD and instances in ScanNet based on box IoU, we further require that the matched pairs belong to the same category. This excludes some cases of mismatching caused by overlapping object boxes. The annotation variations will be made publicly available. For fair comparison, we re-evaluate previous RfD-Net and DIMR models on the refined annotations.

Metrics. Following previous works (Nie et al. 2021; Tang et al. 2022), we evaluate performance from two aspects including completion quality and mapping quality.

(i) *Completion quality.* It is measured by the distance between the reconstructed meshes and the aligned models in the ground truth. We employ three metrics to evaluate completion quality following (Tang et al. 2022): the voxel-based 3D Intersection over Union (IoU) metric, the point-based Chamfer Distance (CD) metric, and the mesh-based Light

	IoU@0.25	IoU@0.5	CD@0.1	CD@0.047	LFD@5000	LFD@2500	PCR@0.5
RfD-Net(Nie et al. 2021)	42.52	14.35	46.37	19.09	28.59	7.80	43.49
RfD-Net(refined data)	45.84	15.60	46.97	18.41	29.90	6.83	43.49
DIMR(Tang et al. 2022)	46.34	12.54	52.39	25.71	29.47	8.55	56.76
DIMR(refined data)	53.18	13.61	54.20	26.99	33.14	8.47	56.76
DOCTR w/o MEBR (ours)	55.18	17.88	56.82	29.17	33.17	12.00	55.70
DOCTR w/ MEBR (ours)	58.25	19.60	59.61	31.46	33.61	10.9	59.89

Table 1: Comparisons on mesh completion quality and mapping quality. We report mean average precision for different metric@threshold. For IoU and PCR, higher thresholds are more difficult. For CD and LFD, smaller thresholds are more difficult.

	Prec.@0.25	Prec.@0.5
RfD-Net(refined data)	22.99	7.92
DIMR(refined data)	36.23	11.91
DOCTR w/o MEBR (ours)	47.75	21.08
DOCTR w/ MEBR (ours)	49.14	21.63

Table 2: Object Recognition Precision. We report the precision of object recognition at different IoU thresholds.

Field Distance (LFD) metric. We adopt these metrics with the same thresholds as these in DIMR to determine whether a predicted mesh can match a ground-truth mesh. We also report the mean precision of IoU over all classes.

(ii) *Mapping quality*. We adopt the Point Coverage Ratio (PCR) proposed in (Tang et al. 2022) to measure the distance between the reconstructed mesh and the original point cloud, which computes the nearest distance from each observed instance point to the corresponding mesh surface.

Implementation details. The Minkowski Res16UNet34C (Choy, Gwak, and Savarese 2019) is used as the 3D U-Net backbone. The SGDQ in our DTD is represented by 256D features, of which the first 128 dimensions represent semantic features, and the latter 128 dimensions represent geometric features. Following (Sun et al. 2023), we also reduce memory consumption by computing the dot product between instance queries and aggregated point features within segments which are obtained from a graph-based segmentation (Felzenszwalb and Huttenlocher 2004). During training, we use the AdamW (Loshchilov and Hutter 2017) optimizer for 600 epochs with a batch size of 5 on a single Nvidia RTX A6000 GPU for all the experiments. One-cycle learning rate schedule (Smith and Topin 2019) is utilized with a maximum learning rate of 10^{-4} and a minimum learning rate of 10^{-6} . Standard data augmentation are performed on point cloud including horizontal flipping, random rotations around the z-axis, elastic distortion, and random scaling.

Comparisons to the State-of-the-Arts

Table 1 shows the quantitative comparisons between our method with RfD-Net (Nie et al. 2021) and DIMR (Tang et al. 2022) in metrics of completion quality and mapping quality. As described in previous section, some wrong annotations are removed, so we use the officially released models of RfD-Net and DIMR to re-evaluate the metrics for fair comparison. For reader’s convenience, Table 1 in-

cludes both the metrics from the original paper report and our re-evaluation results. We provide results for our DOCTR method in two versions, one is with mask enhanced box refinement (MEBR) and the other is without MEBR. As shown in Table 1, our DOCTR (w/o MEBR) achieves improved performance on most metrics. Especially in the IoU@0.5, our DOCTR (w/o MEBR) surpasses RfD-Net by a large margin, while DIMR shows slightly worse performance than RfD-Net. Also, our DOCTR (w/o MEBR) obtains consistent improvements on all the CD and LFD metrics. A reason for the lower PCR@0.5 is that the PCR metric is highly related to point cloud and a slight box deviation can lead to a significant decrease, even if our reconstructed shape is of good quality. With adding the MEBR to enhance box quality, our DOCTR (with MEBR) achieves further improvement and surpasses DIMR in all the metrics including the PCR.

Table 2 reports the precision of object recognition at different IoU thresholds. The Prec@0.25 and Prec@0.5 are both improved more than 10 points by our DOCTR, indicating that our model greatly reduces false positive proposals. This shows that our method enables direct sparse predictions, and even without using NMS, it yields fewer false positives compared to existing methods with NMS.

Ablation Study

We conduct ablation studies to analyze the effectiveness of our method on ScanNet. As shown in Table 3, the first row shows the performance of a baseline model that is modified from Mask3D (Schult et al. 2023) by directly adding box and shape prediction heads. Both the IoU@0.5 and CD@0.1 metrics of this baseline are inferior to those of DIMR, which shows that straight forward applying an object-centric Transformer-based model is not sufficient. Adding hybrid bipartite matching module to the baseline model brings an improvement in the CD@0.1 metric, but the IoU@0.5 metric shows a slight decrease. As shown in the third row, our proposed SGDQ for DTD layers achieves apparent improvement in both metrics compared to baseline model, and both metrics surpass those of DIMR. By combining the SGDQ design with hybrid bipartite matching, both metrics are further improved and reach the best results.

Qualitative Comparisons

We present a qualitative comparison in Figure 5. The version of our model to yield the results is the DOCTR (with MEBR). The first two rows in Figure 5 demonstrate that

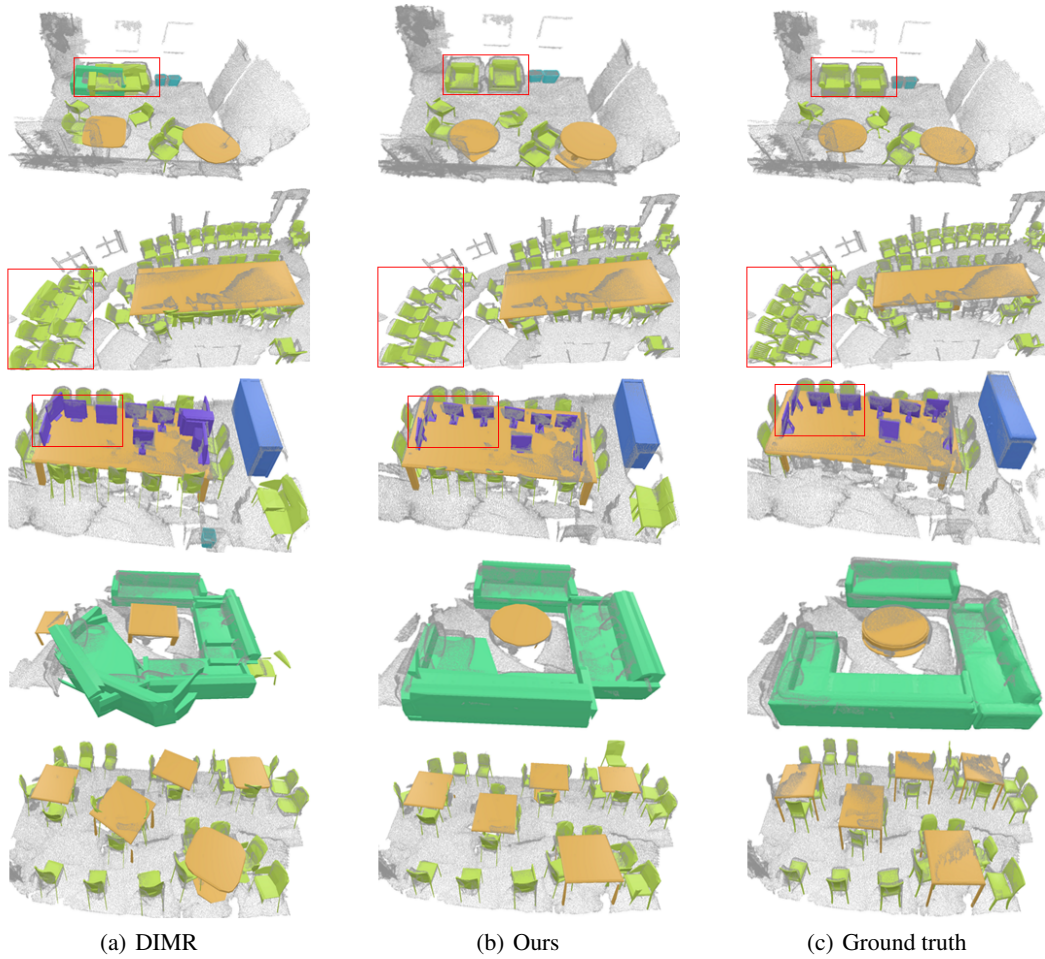


Figure 5: Qualitative comparison on the ScanNet dataset

SGDQ	Hybrid Matching	MEBR	IoU@0.25	CD@0.1
✗	✗	✗	52.59	52.51
✗	✓	✗	52.51	53.31
✓	✗	✗	53.93	54.79
✓	✓	✗	55.18	56.82
✓	✓	✓	58.25	59.61

Table 3: Ablation Study.

DIMR tends to yield erroneous segmentation when there are multiple objects in close proximity, resulting in multiple repetitive false positives. In comparison, our method achieves accurate segmentation and correct object reconstruction even for the closely placed objects. The third row shows that DIMR predicts incorrect orientations for the monitors and the fifth row shows that DIMR predicts some noticeable angular errors of the tables. In contrast, our method predicts correct results. Our method also outperforms DIMR in object shapes, such as for distinguishing between square and round tables in the fifth row. Moreover,

DIMR is less accurate in predicting objects of limited presence in training set such as L-shaped sofas in the fourth row.

As the visualization shows, our DOCTR demonstrates significant and consistent improvements than DIMR in the reconstructed scene in terms of object instance segmentation, pose accuracy, and shape quality.

Conclusion

This paper introduces a novel Disentangled Object-Centric Transformer (DOCTR) that allows learning with multiple objects for multiple sub-tasks in a unified manner. In particular, our design of semantic-geometry disentangled query is proved to be effective to improve performance on the different semantic and geometric sub-tasks. Extensive experiments demonstrate our superior performance than previous SOTA methods. At present, our method relies on a pre-trained shape decoder to generate object meshes, which can be integrated within the main network for collaborative optimization in future work. We hope our work inspires more future works to explore unified learning models to understand rich attributes of multiple objects in complex scenes.

Acknowledgments

We would like to acknowledge contributions through helpful discussions from Hui Zhang and Yi Zhou.

References

- Avetisyan, A.; Dahnert, M.; Dai, A.; Savva, M.; Chang, A. X.; and Nießner, M. 2019. Scan2CAD: Learning CAD model alignment in RGB-D scans. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 2614–2623.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; et al. 2015. ShapeNet: An information-rich 3D model repository. *arXiv preprint arXiv:1512.03012*.
- Chen, Z.; Tagliasacchi, A.; and Zhang, H. 2020. BSP-Net: Generating compact meshes via binary space partitioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 45–54.
- Chen, Z.; Tagliasacchi, A.; and Zhang, H. 2021. Learning mesh representations via binary space partitioning tree networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Choy, C.; Gwak, J.; and Savarese, S. 2019. 4D Spatio-Temporal ConvNets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3075–3084.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5828–5839.
- Felzenszwalb, P. F.; and Huttenlocher, D. P. 2004. Efficient graph-based image segmentation. *International journal of computer vision*, 59: 167–181.
- Garbade, M.; Chen, Y.-T.; Sawatzky, J.; and Gall, J. 2019. Two stream 3D semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 0–0.
- Huang, H.; Chen, H.; and Li, J. 2019. Deep neural network for 3D point cloud completion with multistage loss function. In *2019 Chinese Control And Decision Conference (CCDC)*, 4604–4609. IEEE.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Liu, S.; Li, F.; Zhang, H.; Yang, X.; Qi, X.; Su, H.; Zhu, J.; and Zhang, L. 2022. DAB-DETR: Dynamic anchor boxes are better queries for DETR. *arXiv preprint arXiv:2201.12329*.
- Locatello, F.; Weissenborn, D.; Unterthiner, T.; Mahendran, A.; Heigold, G.; Uszkoreit, J.; Dosovitskiy, A.; and Kipf, T. 2020. Object-centric learning with slot attention. *Advances in Neural Information Processing Systems*, 33: 11525–11538.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Nie, Y.; Hou, J.; Han, X.; and Nießner, M. 2021. RfD-Net: Point scene understanding by semantic instance reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4608–4618.
- Rist, C. B.; Emmerichs, D.; Enzweiler, M.; and Gavrilu, D. M. 2021. Semantic scene completion using local deep implicit functions on LiDAR data. *IEEE transactions on pattern analysis and machine intelligence*, 44(10): 7205–7218.
- Rist, C. B.; Schmidt, D.; Enzweiler, M.; and Gavrilu, D. M. 2020. SCSSnet: Learning spatially-conditioned scene segmentation on LiDAR point clouds. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, 1086–1093. IEEE.
- Roldao, L.; de Charette, R.; and Verroust-Blondet, A. 2020. LMSCNet: Lightweight multiscale 3D semantic completion. In *2020 International Conference on 3D Vision (3DV)*, 111–119. IEEE.
- Schult, J.; Engelmann, F.; Hermans, A.; Litany, O.; Tang, S.; and Leibe, B. 2023. Mask3D: Mask Transformer for 3D Semantic Instance Segmentation.
- Smith, L. N.; and Topin, N. 2019. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, 369–386. SPIE.
- Sun, J.; Qing, C.; Tan, J.; and Xu, X. 2023. Superpoint transformer for 3D scene instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2393–2401.
- Tang, J.; Chen, X.; Wang, J.; and Zeng, G. 2022. Point scene understanding via disentangled instance mesh reconstruction. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, 684–701. Springer.
- Wang, X.; Ang Jr, M. H.; and Lee, G. H. 2020. Cascaded refinement network for point cloud completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 790–799.
- Wu, S.-C.; Tateno, K.; Navab, N.; and Tombari, F. 2020. SC-Fusion: Real-time incremental scene reconstruction with semantic completion. In *2020 International Conference on 3D Vision (3DV)*, 801–810. IEEE.
- Xie, H.; Yao, H.; Zhou, S.; Mao, J.; Zhang, S.; and Sun, W. 2020. GRNet: Gridding residual network for dense point cloud completion. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX*, 365–381. Springer.
- Yan, X.; Gao, J.; Li, J.; Zhang, R.; Li, Z.; Huang, R.; and Cui, S. 2021. Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 3101–3109.

- Yuan, W.; Khot, T.; Held, D.; Mertz, C.; and Hebert, M. 2018. PCN: Point completion network. In *2018 international conference on 3D vision (3DV)*, 728–737. IEEE.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L. M.; and Shum, H.-Y. 2022. DINO: DETR with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*.
- Zhong, M.; and Zeng, G. 2020. Semantic point completion network for 3D semantic scene completion. In *ECAI 2020*, 2824–2831. IOS Press.
- Zhu, B.; Wang, Z.; Shi, S.; Xu, H.; Hong, L.; and Li, H. 2023. Conquer: Query contrast voxel-DETR for 3D object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9296–9305.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable DETR: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.