

X-Trans2Cap: Cross-Modal Knowledge Transfer using Transformer for 3D Dense Captioning

Zhihao Yuan^{1,†}, Xu Yan^{1,†}, Yinghong Liao¹, Yao Guo², Guanbin Li³, Shuguang Cui¹, Zhen Li^{1,*}

¹The Chinese University of Hong Kong (Shenzhen), The Future Network of Intelligence Institute, Shenzhen Research Institute of Big Data, ²Shanghai Jiao Tong University, ³Sun Yat-sen University

{zhihaoyuan@link., xuyan1@link., lizhen@cuhk.edu.cn}

Abstract

3D dense captioning aims to describe individual objects in 3D scenes by natural language, where 3D scenes are usually represented as RGB-D scans or point clouds. However, only exploiting single modal information, e.g., point cloud, previous approaches fail to produce faithful descriptions. Though aggregating 2D features into point clouds may be beneficial, it introduces an extra computational burden, especially in the inference phase. In this study, we investigate a **cross-modal knowledge transfer** using **Transformer** for 3D dense captioning, namely **X-Trans2Cap**. Our proposed X-Trans2Cap effectively boost the performance of single-modal 3D captioning through the knowledge distillation enabled by a teacher-student framework. In practice, during the training phase, the teacher network exploits auxiliary 2D modality and guides the student network that only takes point clouds as input through the feature consistency constraints. Owing to the well-designed cross-modal feature fusion module and the feature alignment in the training phase, X-Trans2Cap acquires rich appearance information embedded in 2D images with ease. Thus, a more faithful caption can be generated only using point clouds during the inference. Qualitative and quantitative results confirm that X-Trans2Cap outperforms previous state-of-the-art by a large margin, i.e., about **+21** and **+16** CIDEr points on ScanRefer and Nr3D datasets, respectively.

1. Introduction

Hitherto, the computer vision community has witnessed significant progress in image captioning [3, 25, 33, 46, 50] and dense captioning [24–26, 30] under the success of deep learning techniques. Unlike image captioning describing a 2D image with a single sentence, dense captioning (DC) better interprets “A picture is worth a thousand words”.

* Corresponding author: Zhen Li. † Equal first authorship.

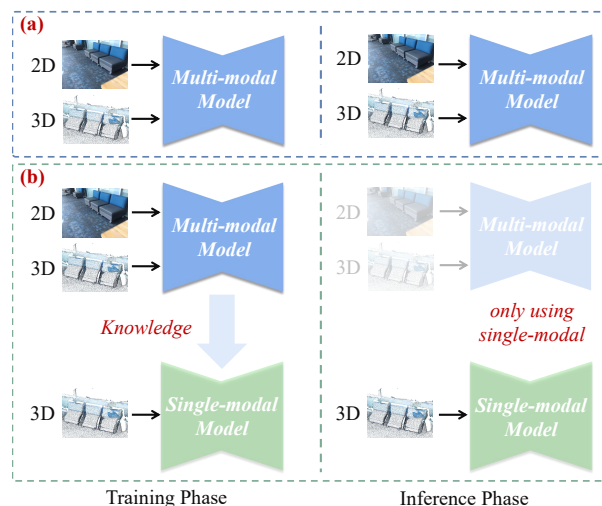


Figure 1. **The motivation of cross-modal knowledge transfer.** (a) Previous methods use the extra 2D modality as the input in both training and inference phases. (b) In contrast, we exploit a teacher-student framework with multi-modal data during training. For inference, the student network only takes 3D modality input.

That is to say, for DC task, each object in an image is first perceived, then is provided more customized and detailed descriptions according to its nature and context.

Most recently, 3D cross-modal learning in vision and language has gained an increasing amount of interest as well. Several datasets [1, 6, 15] and downstream applications [20, 56] are proposed and investigated. Unlike 2D images with regular grids and dense pixels, 3D data represented by a set of points are unordered and scattered in the 3D space, impeding the direct extension of 2D-based methods to 3D scenarios. To perform dense captioning on 3D point clouds, [9] proposes the first method, namely Scan2Cap, by directly combining 3D object detection with natural language generation. Specifically, Scan2Cap first employs a detection backbone to obtain object proposals,

and then applies a relational graph and a context-aware attention captioning module to learn object relations and generate tokens. Besides, multi-view features extracted by the pre-trained E-Net [36] are further projected onto the input point cloud to enhance final captioning. However, Scan2Cap still has several issues: **1)** The object representations in Scan2Cap are defective since they are solely learned from sparse 3D point clouds, thus failing to provide strong texture and color information compared with the ones generated from 2D images. **2)** It requires the extra 2D input in both training and inference phases, as shown in Figure 1 (a). However, the extra 2D information is usually computation-intensive and unavailable during inference. For instance, a model training with both 2D and 3D inputs cannot apply to LiDAR scenarios that only contains 3D point cloud.

To address the above issues, we explore how to ease the barrier of cross-modal learning on 2D and 3D data, and investigate how to effectively combine the merits of both modalities for 3D dense captioning in this paper. To this end, we first time present a flexible and novel cross-modal framework, namely **X-Trans2Cap**¹, which transfers color and texture-aware information from 2D image into 3D object representation using Transformer [44]. Concretely, all the instances in a given scene can be firstly extracted by 3D object detection. Subsequently, the 3D features of each instance and its 2D counterpart are processed by a teacher-student framework. Within this framework, the teacher network takes the multi-modal inputs, while the student one only leverages the 3D inputs. Considering different modalities for teacher and student streams, we innovatively design a Transformer-based knowledge transfer framework with more flexible input control and better representation. Moreover, to further enhance the knowledge transfer, a modified knowledge distillation operation with cross-modal fusion (CMF) module and cross-modal feature alignment objective is proposed for knowledge generalization. Owing to the end-to-end training scheme, the priors in the 2D modality can inherently improve the teacher network and the student as well, *i.e.*, our model takes advantage of the color and texture aware 2D representation and reduces the extra computational cost. Therefore, in the inference phase, X-Trans2Cap can perform superior captioning performance with only 3D inputs, as shown in Figure 1 (b).

Sufficient experiments evaluated on the ScanRefer [6] and Nr3D [1] datasets have demonstrated the effectiveness of our proposed X-Trans2Cap. In specific, with the extra 2D priors and the novel framework design, X-Trans2Cap can effectively learn a better 3D object representation and boost the performance over the model without 2D priors, *i.e.*, improving the CIDEr points on ScanRefer from 75.75 to 87.09. This result also exceeds the previous state-of-the-art Scan2Cap by about **21 CIDEr**.

¹<https://github.com/CurryYuan/X-Trans2Cap>

In summary, our main contributions are threefold:

- We first time propose **X-Trans2Cap**, a simple but effective cross-modal knowledge transfer framework for 3D dense captioning, in which an enhanced 3D representation with 2D priors is achieved.
- X-Trans2Cap leverages a modified knowledge distillation method through a novel cross-modal fusion module and feature alignment techniques merged in Transformer, eliminating extra computation burdens during inference while achieving superior knowledge transfer.
- Our X-Trans2Cap gains significant performance boost on the ScanRefer [6] (**+21.0 CIDEr**) and Nr3D [1] (**+16.7 CIDEr**) datasets.

2. Related Work

2.1. Image Captioning and Dense Captioning

A broad collection of methods have been proposed in the field of image captioning in the last few years [13, 29, 33, 46, 50]. Recently, many methods focus on utilizing the attention mechanism to capture meaningful information in the image, *e.g.*, over grid regions [33, 50] and detected objects [3, 34]. Furthermore, some works attempt to combine attention with graph neural networks [14, 23, 52, 54] or Transformer [10] to boost performance.

For the dense captioning task, it needs to generate captions for all the detected objects. Johnson *et al.* [24] is the pioneer in this challenging field. Along this line, [51] considers the context outside the salient image regions and takes advantage of global image features. [26] further introduces the object relations among detected regions. However, due to the limited views of a single image, the performance of image-based dense captioning methods is significantly degraded when directly transferred to 3D scenarios.

2.2. 3D Vision and Language

Compared to image and language comprehension, 3D vision and language understanding is a relatively emerging research field. Existed works focus on using language to confine individual objects, *e.g.*, detecting referred 3D objects [8] or distinguishing objects according to language phrases [2]. Recently, ScanRefer [6] and ReferIt3D [1] introduce a task of localizing objects within a 3D scene given the linguistic descriptions, namely 3D visual grounding. TGNN [20] and InstanceRefer [56] follow the above settings and exploit panoptic segmentation to reduce the number of proposals. 3D dense captioning is proposed very lately in Scan2Cap [9]. It focuses on decomposing 3D scenes and describing the chromatic and spatial information of the objects. Very recently, [61] combines the above task of 3D grounding and caption to mutually enhance the performance of two tasks. Though promising, it

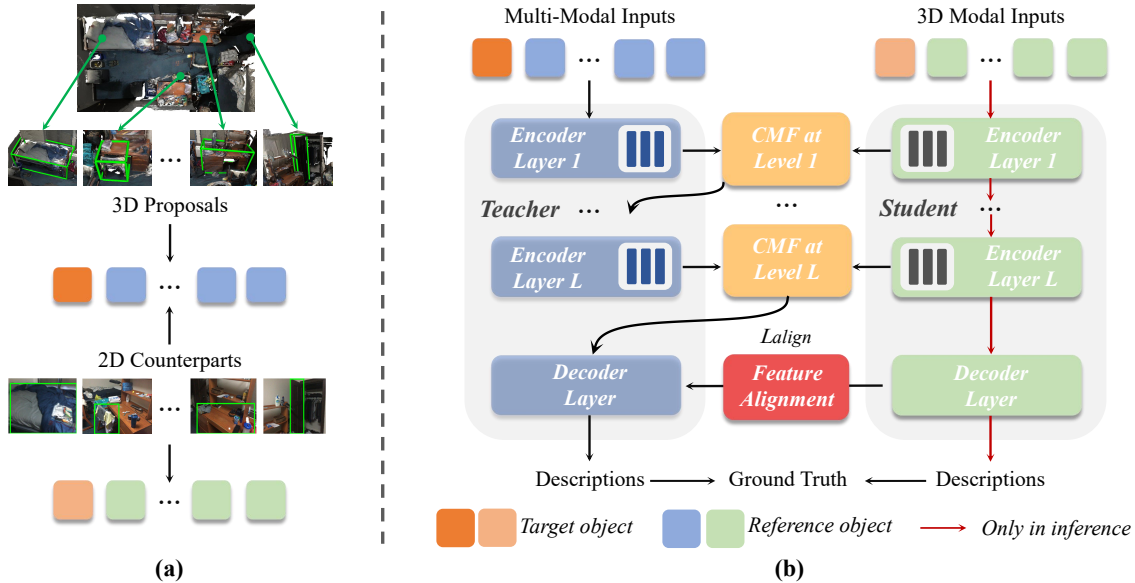


Figure 2. **X-Trans2Cap**. Part (a) describes the object representation, where the 3D proposals are utilized in the generation of pure 3D modal inputs. By integrating their 2D counterparts, multi-modal inputs are obtained. Part (b) illustrates the architecture of X-Trans2Cap, where the framework adheres to a teacher-student design. Both the teacher and student networks consist of several Transformer encoder layers and a decoder layer.

only takes point clouds as input to generate instance features. Compared with the well-organized 2D images containing stronger texture and color information, such representation inherently challenges the learning process.

2.3. Cross-modal Knowledge Transferring

Previous studies apply 2D images as the extra inputs to 3D tasks, *e.g.*, 3D object detection [28, 37, 39, 49], semantic segmentation [12, 19, 22] and object tracking [59, 60]. However, they require extra 2D information in both the training and inference phases. Thus, it inevitably augments computational burdens during evaluation and severely limits the efficiency in real-world applications. The concept of knowledge distillation was first shown by Hinton *et al.* [18]. Subsequent research [4, 7] enhanced distillation by matching intermediate representations in the networks along with outputs using different approaches. Zagoruyko *et al.* [57] proposed to align attentional activation maps between networks. Srinivas and Fleuret [43] improved it by applying Jacobian matching to networks. In recent years, cross-modal knowledge distillation [16, 47, 55, 58] extended knowledge distillation by applying it to transferring knowledge across different modalities. Very recently, there are works attempting to only utilize 2D images during training phase to address the above problems. Among them, the 2D-assisted pre-training [32], inflating 2D convolution kernels to 3D [48] and joint training with mask attention [53] are proposed. Unlike those, we adopt a well-designed teacher-

student framework with cross-modal fusion for more efficacious knowledge transfer, and the experiment results also demonstrate that our approach is much better than previous knowledge transferring.

3. Method

Our X-Trans2Cap is developed upon a teacher-student framework [18], which is widely exploited in the knowledge distillation research field. The detailed architecture of X-Trans2Caps is presented in Figure 2. X-Trans2Cap takes two types of features as input, *i.e.*, pure 3D modal input for student and multi-modal input for teacher respectively. We first introduce the details of the above feature representation in Section 3.1. Then we propose a baseline model for 3D dense captioning through Transformers [44] in Section 3.2, named TransCap. In Section 3.3, we illustrate how X-Trans2Cap transfers the 2D priors to the 3D representations, in which a cross-modal fusion (CMF) module is proposed. The details of training objectives are presented in Section 3.4. Finally, by incorporating the above components in a whole architecture, we illustrate the data flow of X-Trans2Cap in training and inference phases in Section 3.5.

3.1. Object Representation

As shown in Figure 2 (a), our framework takes object-level representation as input, and each object feature is de-

noted as a token. Given that there are M objects in the 3D scene, in the remaining section, the objects set is represented as $\mathbf{O} = \{O_m\}_{m=1}^M$, in which O_m and O_m^{att} are depicted as the m -th object and the attribute of the m -th object, respectively. In each iteration, we randomly choose an object as the target object (O_*) to be described as in [9]. The other $M - 1$ objects, *i.e.*, $\{O_m \in \mathbf{O}\} \cap \{O_m \neq O_*\}$, are treated as the reference objects, and only provide the cues of locations or relations to the target object.

For the 3D modal input, each object is considered from the perspective of its 3D feature, semantic, size as well as relative position to the target object. Specifically, the object representation is computed as follows:

$$F_m^{3d} = \mathcal{T}_1([O_m^{f3d}; O_m^{cls}; W_1 O_m^{b3d}; W_2 O_m^{pe}]), \quad (1)$$

where $[\cdot; \cdot]$ indicates the concatenation operation. O_m^{f3d} is the output feature extracted by a 3D network, *e.g.*, PointNet++ [40], and O_m^{cls} is a one-hot vector for their predicted semantic class. O_m^{b3d} is the 3D bounding box of the object, consisting of the bounding box center (x, y, z) and size (w, h, l) . To a better object representation, we further design a positional encoding O_m^{pe} for the m -th object as:

$$O_m^{pe} = [O_m^x - O_*^x; O_m^y - O_*^y; O_m^z - O_*^z; O_m^w/O_*^w; O_m^h/O_*^h; O_m^l/O_*^l]. \quad (2)$$

The first three elements in the positional encoding calculate the center offset between the target object and m -th object, and the others denote their relative size. Two learnable projection matrices W_1 and W_2 in Eqn. (1) then transform the dimensions of O_m^{b3d} and O_m^{pe} to d . Finally, a transformation function \mathcal{T}_1 generates the final object feature F_m^{3d} for the m -th object.

Apart from 3D information in the multi-modal input, the corresponding 2D feature O_m^{f2d} and 2D bounding box O_m^{b2d} are introduced for the m -th object as follows:

$$F_m^{\text{multi}} = \mathcal{T}_2([O_m^{f3d}; O_m^{cls}; W_1 O_m^{b3d}; W_2 O_m^{pe}; O_m^{f2d}; W_3 O_m^{b2d}]). \quad (3)$$

Concretely, for each object, its ground truth of the 3D bounding box is projected onto the original ScanNet videos [11] to obtain the corresponding 2D boxes. In each training step, an image is randomly selected from the video sequences to generate an extra input. Features in the 2D box area are extracted by the Faster-RCNN detector [41] pre-trained on the Visual Genome [27] dataset, which are regarded as 2D features for the m -th instance, *i.e.*, O_m^{f2d} . Finally, as shown in Eqn. (3), by applying linear and nonlinear transformations W_3 and T_2 , a d -dimensional multi-modal representation F_m^{multi} is generated for the m -th object.

As shown in Figure 2 (a), the multi-modal and 3D modal inputs have the same format, and each of them is a set of

object features with the shape of $\mathbb{R}^{M \times d}$, in which the features of the target object and $M - 1$ reference objects are entailed. For convenience, we denote the multi-modal input as $\mathbf{F}^{\text{multi}} = \{F_m^{\text{multi}}\}_{m=1}^M$, and the 3D one as $\mathbf{F}^{3d} = \{F_m^{3d}\}_{m=1}^M$. Then, these two inputs are fed to the teacher and student networks for the cross-modal knowledge transfer.

3.2. Baseline Model: TransCap

To apply the 3D and multi-modal object representation to our framework, we first introduce a baseline model, named TransCap, which employs Transformer [44] structure to generate the descriptions of the target object. The student network in Figure 2 (b) displays the architecture of TransCap. It contains L encoder layers and one decoder layer. In each encoder layer, a self-attention mechanism is exploited to obtain a permutation invariant encoding for the input feature. Inspired by [10], we design the self-attention operator $\text{SA}(\cdot)$ as follows:

$$\begin{aligned} \text{SA}(\mathbf{X}) &= \text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}), \\ \mathbf{Q} &= W_q \mathbf{X}, \mathbf{K} = [W_k \mathbf{X}; M_k], \mathbf{V} = [W_v \mathbf{X}; M_v], \end{aligned} \quad (4)$$

where $\mathbf{X} \in \mathbb{R}^{M \times D}$ is a D -dimensional sequence, and W_q , W_k and W_v are matrices of learnable weights. Different from the traditional attention mechanism [44], the two persistent memory vectors M_k and M_v are appended to learn the prior knowledge.

The captioning decoder is conditioned on previously generated words and features from the encoder layers to generate the next token. Specifically, it integrates the features from different encoder layers and performs the cross-attention on the generated tokens.

$$\text{Decoder}(\hat{\mathbf{X}}_l, \mathbf{Y}) = \sum_{l=1}^L \alpha_l \odot \text{CA}(\hat{\mathbf{X}}_l, \mathbf{Y}), \quad (5)$$

where $\text{CA}(\cdot, \cdot)$ stands for the encoder-decoder cross-attention [44], computed using queries from the decoder output \mathbf{Y} , and keys and values from the l -th layer encoder output $\hat{\mathbf{X}}_l$. α are learnable weights having the same size as the cross-attention results. In this manner, TransCap takes a sequence of object features as the input, and generates the description for the target object.

3.3. Cross-Modal Fusion Module

The Cross-modal fusion (CMF) module enables cross-modal feature interaction between pure 3D and multi-modal feature representations. As shown in Figure 3, it is designed to construct interaction from the student network to the teacher at the same encoder level, thus building a bridge to fuse features between single and multiple modalities. Moreover, to further enhance the ability of the student network to learn the multi-modal representation, we exploit a random

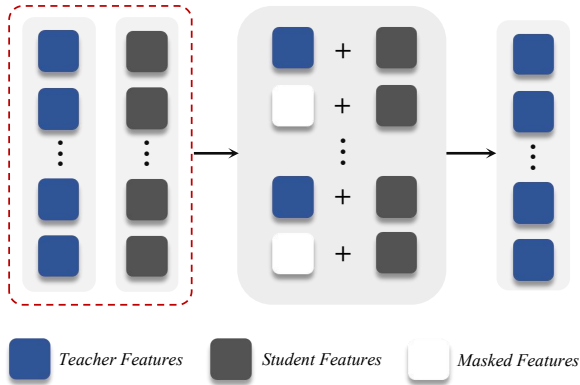


Figure 3. **Cross-Modal Fusion (CMF) Module.** It is designed to construct the interaction from the student to the teacher network at the same encoder level, interacting the features of different modalities. During the training phrase, the features from the teacher network are randomly masked with a specific probability.

mask on the features from the teacher network. Owing to this framework, the strengths of multi-modal representation can be assimilated to reinforce the student network via an end-to-end training protocol. Specifically, we element-wise add the student features with the masked teacher features.

$$\hat{\mathbf{X}}_l^* = \hat{\mathbf{X}}_l^{\text{student}} \oplus \mathcal{I}(p)\hat{\mathbf{X}}_l^{\text{teacher}}, \quad (6)$$

where the $\hat{\mathbf{X}}_l^{\text{student}}$ and $\hat{\mathbf{X}}_l^{\text{teacher}}$ denote features from l -th encoder layer of student and teacher networks, respectively. The notation \oplus means element-wise addition. The mask indicator $\mathcal{I}(p)$ is initialized with 1 and has the probability of p change to 0. After that, we feed the fused features into the next encoder layer of the teacher network. It should be highlighted that, since our CMF module employs the single-directional connection from the student to the teacher, the teacher network can be discarded during inference, *i.e.*, it introduces no extra computation for the student network. Moreover, various designs for the CMF module, including ablations, are shown in Section. 4.4.

3.4. Objective Function

Feature alignment loss. Following a standard practice in knowledge transfer, we use Huber loss $\mathcal{L}_{\text{align}}$ (*i.e.*, Smooth-L1 regression loss) to align decoder features between teacher and student networks.

Captioning loss. As in the previous work [9], we apply a conventional cross entropy loss function \mathcal{L}_{ce} on the generated token probabilities in both teacher and student networks. Furthermore, to further boost the performance, we propose an enhanced version model X-Trans2Cap (C) by applying the CIDEr-D score [3] $\mathcal{L}_{\text{CIDEr}}$ as reward. Following the previous work [10], we baseline the reward using the

mean of the rewards rather than greedy decoding following previous methods [3, 42].

Total objective loss. We combine all three loss terms linearly as our final objective loss function:

$$\mathcal{L} = \alpha\mathcal{L}_{\text{align}} + \beta\mathcal{L}_{\text{ce}} + \gamma\mathcal{L}_{\text{CIDEr}}, \quad (7)$$

where α , β and γ are the weights for each individual loss. To guarantee the loss terms are roughly of the same magnitude, we fine-tune weights on the validation split, and set those to $\alpha = 1$, $\beta = 1$, and $\gamma = 0.1$ empirically in the experiments.

3.5. Training and Inference Schemes

The black and red arrows in the Figure 2 (b) illustrate the information flow of the X-Trans2Cap for training and inference. It is noteworthy that teacher and student networks are trained from scratch. In the training phase, both networks are exploited (see the black and red arrows in Figure 2 (b)), and CMF modules between corresponding encoder layers and feature alignment are conducted to enhance mutual representation. During the inference, if only the 3D modality exists, we only apply the student network (see the red arrow in Figure 2 (b)). However, if the auxiliary 2D information is available as well, the stronger teacher framework will be exploited. In our experiment, we demonstrate that our architecture can both enhance the performance of teacher and student networks with and without additional modality.

4. Experiment

We compare our method with Scan2Cap and 2D baselines proposed in their paper. Extending from [9], we further compare all methods on Nr3D dataset [1]. More experiment results including subjective evaluation and ablations are in the supplementary material.

4.1. Datasets

ScanRefer. The ScanRefer dataset [6] annotates 800 3D indoor scenes in the ScanNet [11] dataset with 51,583 language queries. It follows the official ScanNet splits and contains 36,665, 9,508, and 5,410 samples in train/val/test sets, respectively. Since the dataset is initially used in visual grounding and the labels of the test set are inaccessible, we follow the same setting as in [9] to form the train and val sets for training and testing.

Nr3D. The Natural Reference in 3D (Nr3D) [1] has the same train/val split as ScanRefer. It contains 41,503 queries annotated by Amazon Mechanical Turk (AMT) workers. Compared with ScanRefer dataset, Nr3D is more challenging since it does not contain the fixed or redundant sentence patterns, *i.e.*, declarative sentences starting with “*this is*” or “*that is*”. We do not compare our method on its counterpart, the Spatial Reference in 3D (Sr3D) dataset, since it is totally generated by the machine templates.

Table 1. Comparison of 3D dense captioning results obtained by X-Trans2Cap and previous methods using ground truth instances on ScanRefer and Nr3D datasets. We introduce the conventional captioning metrics, *i.e.*, CIDEr (C), BLEU-4 (B-4), METEOR (M) and ROUGE (R) for evaluation. The column ‘Extra 2D’ means whether using extra 2D modality in the *inference* phase. X-Trans2Cap (C) denotes the original model exploiting extra $\mathcal{L}_{\text{CIDEr}}$ loss in the final objective function.

Method	Extra 2D	ScanRefer				Nr3D			
		C	B-4	M	R	C	B-4	M	R
Scan2Cap [9]	✗	65.79	38.54	28.81	61.93	63.36	32.07	28.92	64.56
Scan2Cap (Inst)	✗	64.44	36.89	28.42	60.42	61.89	32.02	28.88	64.17
TransCap	✗	75.75	42.06	28.82	62.62	70.60	35.99	29.04	66.00
X-Trans2Cap	✗	87.09	44.12	30.67	64.37	80.02	37.90	30.48	67.64
X-Trans2Cap (C)	✗	89.46	44.46	30.71	64.55	81.44	39.08	30.79	68.15
Scan2Cap [9]	✓	67.95	41.49	29.23	63.66	64.13	32.98	29.75	65.24
Scan2Cap (Inst)	✓	70.04	41.57	29.67	64.10	64.00	33.19	29.53	65.29
TransCap	✓	88.72	44.24	30.95	64.70	77.55	37.25	30.63	67.43
X-Trans2Cap	✓	89.73	44.25	31.00	64.50	85.38	39.52	31.23	68.18
X-Trans2Cap (C)	✓	106.11	49.07	32.25	65.54	85.40	40.51	31.36	68.84

4.2. Tasks and Metrics

Tasks. In our experiment, we follow [9] and design two protocols to evaluate the generated caption:

- *Dense captioning with ground truth instances* (Oracle DC): In this setting, the point cloud of each instance is given. Then one needs to generate faithful captions based on their attribute information and spatial relationships.
- *Dense captioning with 3D scans* (Scan DC): This setting is more challenging. One needs to detect objects from the 3D scans first and then generate captions for each object according to the detection results.

Metrics. In Oracle DC, we directly apply CIDEr [45], BLEU-4 [35], METEOR [5] and ROUGE [31] averagely on all instances as metrics. For brevity, we simplify them as C, B-4, M and R, correspondingly.

In Scan DC, to jointly measure the quality of generated captions and detected bounding boxes, we evaluate them by combining above metrics with Intersection-over-Union (IoU) scores between predicted bounding boxes and GT bounding boxes. Specifically, we follow [9] and define the combined metrics as $m@k\text{IoU} = \frac{1}{N} \sum_{i=1}^N m_i u_i$, where $u_i \in \{0, 1\}$ is set to 1 if the IoU score for the i -th box exceeds k , otherwise 0. We use m to represent the above captioning metrics, *e.g.*, CIDEr. N is the number of detected object bounding boxes. We also use mean average precision (mAP) thresholded by IoU as the object detection metric.

4.3. 3D Dense Captioning Results

Oracle dense captioning. The results of Oracle DC task are displayed in Table 1. In the upper part, we compare results without the extra 2D input for inference. Scan2Cap and Scan2Cap (Inst) denote the methods that exploit ground-truth (GT) boxes and GT instances as input, respectively.

Merely using the baseline model TransCap, we improve the captioning result by a large margin compared to Scan2Cap (+9.96 and +7.24 CIDEr points on ScanRefer and Nr3D). Utilizing our cross-modal knowledge transfer training strategy further boost the performance on all the captioning metrics. Specifically, after using our teacher student framework, X-Trans2Cap achieves 11.04 and 9.42 CIDEr improvement over TransCap on the ScanRefer and Nr3D datasets, where the performance on both datasets are about **20** CIDEr scores higher than those of Scan2Cap. The bottom part of Table 1 illustrates the result using extra 2D input in both the training and inference phases. Though using the extra 2D input for inference, the performance of Scan2Cap is still inferior to that of our proposed X-Trans2Cap only exploiting 3D modal input, let alone using multi-modal. Besides, X-Trans2Cap is better than TransCap when both using the extra 2D input, especially on Nr3D (85.38 vs 77.55 CIDEr), which illustrates that training with student network can even improve the result of teacher network. Furthermore, with CIDEr-D score optimization, *i.e.*, X-Trans2Cap (C) model, the performance of captioning can be further improved.

Scan dense captioning. In Table 2, we compare the result of Scan DC, which shows results without and with extra 2D input in the inference phase. The method for proposal generation is listed in the third column. Among these methods, *2D-3D Proj.* and *3D-2D Proj.* are two baseline methods proposed in [9]. *2D-3D Proj.* applies Mask R-CNN [17] to generate 2D proposals in images, where the corresponding 2D bounding boxes and features are fed into the description generation module [9]. On the contrary, *3D-2D Proj.* exploits VoteNet [38] to extract 3D proposals, which are projected back to 2D images. Then the projected 2D proposals are finally adopted [50] to generate captions. As shown in Table 2, 2D-based methods obtain lowest caption-

Table 2. Comparison of 3D dense captioning obtained by X-Trans2Cap and previous methods, taking 3D Scans as the input on ScanRefer dataset. We average the scores of the above captioning metrics, which are with the IoU percentage between the predicted bounding box and the ground truth over 0.25 and 0.5, respectively. The ‘Extra 2D’ means that whether using the extra 2D modality as above. *2D-3D Proj.* and *3D-2D Proj.* represent the method in [9], i.e., 2D proposal mapped to 3D and 3D proposal projected to 2D images, respectively. ‘Proposals’ shows the methods exploited to obtain 2D or 3D proposals.

Method	Extra 2D	Proposals	C@0.25	B-4@0.25	M@0.25	R@0.25	C@0.5	B-4@0.5	M@0.5	R@0.5	mAP@0.5
Scan2cap [9]	✗	VoteNet	50.71	33.01	25.47	53.60	33.53	21.58	21.04	43.03	32.46
TransCap	✗	VoteNet	55.36	32.46	25.64	53.19	40.08	22.86	21.72	44.04	33.34
X-Trans2Cap	✗	VoteNet	58.81	34.17	25.81	54.10	41.52	23.83	21.90	44.97	34.68
2D-3D Proj. [50]	✓	Mask R-CNN	18.29	10.27	16.67	33.63	8.31	2.31	12.54	25.93	10.50
3D-2D Proj. [50]	✓	VoteNet	19.73	17.86	19.83	40.68	11.47	8.56	15.73	31.65	31.83
Scan2cap [9]	✓	VoteNet	56.82	34.18	26.29	55.27	39.08	23.32	21.97	44.78	32.21
TransCap	✓	VoteNet	60.04	35.04	26.27	54.46	43.12	24.25	22.15	44.72	34.34
X-Trans2Cap	✓	VoteNet	61.83	35.65	26.61	54.70	43.87	25.05	22.46	45.28	35.31

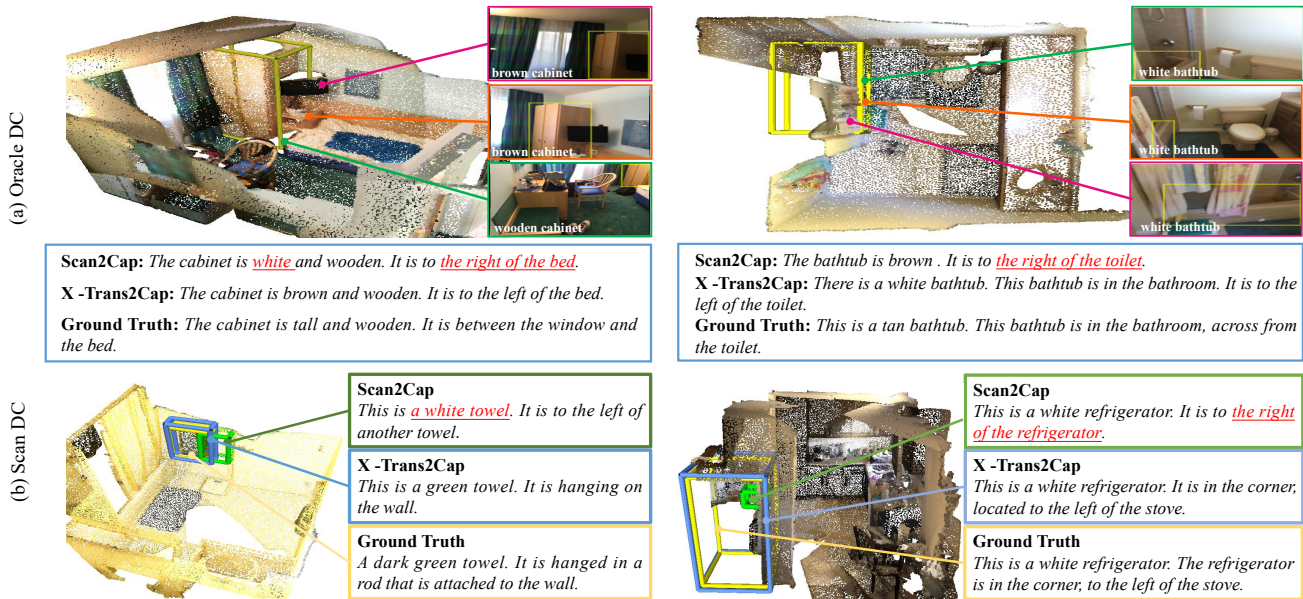


Figure 4. **Qualitative comparisons on ScanRefer dataset.** The upper illustrates the results applying ground truth instances (i.e., Oracle DC task), and the lower part presents the results using object detection (i.e., Scan DC task). Best viewed in color.

ing scores, which reveals that they cannot directly handle the 3D dense captioning task. Though Scan2Cap achieves better results than these 2D based methods, it is much inferior to X-Trans2Cap without the assistance of appealing 2D priors and dedicated network structure. Surprisingly, we observe that the detection performance of X-Trans2Cap is improved as well, though there is no extra 2D input fed into the detector during training and testing. It confirms that our X-Trans2Cap is not only capable of faithful caption generation, but also acquires the knowledge mining capacity within multi-modalities for more complex applications, i.e., digging out the information embedded into language description for 3D visual detection. The results of Scan DC on Nr3D dataset are illustrated in the supplementary.

Visualization. Figure 4 presents the visualization results

of X-Trans2Cap, which demonstrates great improvements upon Scan2Cap for more faithful captions. Furthermore, we present the corresponding 2D counterparts within each 3D scene. Regarding 2D images, they can provide stronger texture and color information obviously when compared with sparse point cloud.

Comparison for knowledge transfer. To further verify the effectiveness of our proposed method upon common teach-student architecture and other cross-modal manners, we compare X-Trans2Cap with typical approaches of knowledge transfer in Table 3. Among all the methods, Hinton *et al.* [18] and Huang *et al.* [21] are pure knowledge distillation designs, where the former is the pioneer for the research filed and the latter is newly proposed. As shown in the table, pure knowledge distillation manners cannot be

Table 3. Comparison for knowledge transfer. The results are obtained in Oracle DC on the ScanRefer dataset, where we compare both single-modal and cross-modal knowledge transfer methods.

Method (Year)	C	B-4	M	R
Hinton <i>et al.</i> [18] (2015)	81.43	42.85	30.40	64.07
Huang <i>et al.</i> [21] (2021)	78.61	41.93	30.14	63.78
Pixel-to-point [32] (2021)	77.82	41.98	29.42	62.79
2D SAT [53] (2021)	80.13	41.13	30.00	63.16
TransCap (baseline)	75.75	42.06	28.82	62.62
X-Trans2Cap (pre-trained)	79.41	42.78	29.88	63.41
X-Trans2Cap	87.09	44.12	30.67	64.37

directly adopted on the 3D DC scenario, and their improvement upon the baseline model is limited. Very recently, approaches [32] and [53] adopt cross-modal knowledge transfer technique in the 3D tasks. The core idea of [32] is using extra 2D input to conduct 3D pre-training. We modify it by first training a TransCap with multi-modal input, then use its pre-trained parameters as initialization weights for pure 3D input training. For 2D semantic-assisted training (SAT) [53], it treats 2D features as additional tokens in (*i.e.*, concatenated in sequence dimension) the same model, and then exploits an attention mask in Transformer layers. This mask only ignores the attention from 3D to the 2D. However, both methods cannot boost the performance as no mutual enhancement is introduced.

We also offer an offline distillation design, preparing a pre-trained teacher network before training the student network, called X-Trans2Cap (pre-trained) in the bottom part of the Table 3. It can be noticed that using pre-trained teacher network results in a performance drop of 7.38 CIDEr, which may result from the distribution gap between multi-modality data. To the end, in the Table 3, X-Trans2Cap significantly performs better, which illustrates the effectiveness of the teacher-student framework and cross-modal fusion (CMF) module.

4.4. Analysis and Ablation Studies

Does knowledge transfer help? As results shown in Table 1 and 2, when we adopt 2D prior during the training phase (X-Trans2Cap), it can greatly improve the performance upon the baseline model (TransCap).

Does our proposed components help? To further verify the effectiveness of different components, we conduct the ablation studies in the Table 4. As shown in Table 4, model A is our baseline model (TransCap), and model B is our entire architecture of X-Trans2Cap. The model C is the ablated architecture that discards the feature alignment loss \mathcal{L}_{align} . It can be found out that there is a great performance drop from 87.09 to 79.58 in terms of CIDEr metric. Fortunately, due to the advantage of CMF module, it still has 3.83 improvement upon the baseline model. Similarly, the per-

Table 4. Ablation study for applying different knowledge transfer designs. The results are obtained in Oracle DC on the ScanRefer dataset. The upper part shows ablated results without specific components, and the lower illustrates results of CMF designs.

	Design	C	B-4	M	R
A	TransCap (baseline)	75.75	42.06	28.82	62.62
B	X-Trans2Cap	87.09	44.12	30.67	64.37
C	w/o \mathcal{L}_{align}	79.58	41.47	30.07	63.59
D	w/o CMF	80.54	43.15	30.19	63.59
E	Concatenation	79.87	43.28	30.22	64.88
F	w/o Randon Mask	85.36	42.57	30.52	64.18
G	Attention [44]	80.85	44.75	30.45	64.91

formance drop is appearing (-5.25 CIDEr) when removing the CMF module (model D). This result demonstrates that both framework architecture and CMF module play important roles in X-Trans2Cap.

How to design cross-modal fusion? We illustrate the results from different designs of CMF module in the bottom part of the Table 4. On the one hand, discarding the random mask hampers the caption results, as shown in model F. On the other hand, exploiting more complicated operations such as concatenation and attention mechanism cannot effectively improve the performance. There is only a slight improvement on the metrics of BLEU-4 and Rough for adopting attention. However, it will greatly increase the model complexity while make the CIDEr decline.

5. Conclusion

In this work, we propose an enhanced 3D dense captioning method via cross-modal knowledge transfer, named **X-Trans2Cap**. By designing the network architecture and knowledge distillation method carefully, our X-Trans2Cap outperforms previous methods by a large margin on multiple datasets with more faithful captions. We believe that our work can be applied to a wider range of 3D vision and language scenarios, and provide a solution to the comprehension of 3D scenes with severe texture details missing, *i.e.*, leveraging the 2D priors and the cross-modal knowledge transfer to improve the performance.

Acknowledgment

This work was supported in part by NSFC-Youth 61902335, by Key Area R&D Program of Guangdong Province with grant No.2018B030338001, by the National Key R&D Program of China with grant No.2018YFB1800800, by Shenzhen Outstanding Talents Training Fund, by Guangdong Research Project No.2017ZT07X152, by Guangdong Regional Joint Fund-Key Projects 2019B1515120039, by the NSFC 61931024&81922046, by helixon biotechnology company Fund and CCF-Tencent Open Fund.

References

- [1] Panos Achlioptas, Ahmed Abdelreheem, Fei Xia, Mohamed Elhoseiny, and Leonidas Guibas. Referit3d: Neural listeners for fine-grained 3d object identification in real-world scenes. In *European Conference on Computer Vision*, pages 422–440. Springer, 2020. 1, 2, 5
- [2] Panos Achlioptas, Judy Fan, Robert Hawkins, Noah Goodman, and Leonidas J Guibas. Shapeglot: Learning language for shape differentiation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8938–8947, 2019. 2
- [3] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. 1, 2, 5
- [4] Lei Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? *NeurIPS*, page 2654–2662, 2014. 3
- [5] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. 6
- [6] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. *16th European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 5
- [7] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 742–751, 2017. 3
- [8] Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. Text2shape: Generating shapes from natural language by learning joint embeddings. In *Asian Conference on Computer Vision*, pages 100–116. Springer, 2018. 2
- [9] Zhenyu Chen, Ali Gholami, Matthias Nießner, and Angel X Chang. Scan2cap: Context-aware dense captioning in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3193–3203, 2021. 1, 2, 4, 5, 6, 7
- [10] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10578–10587, 2020. 2, 4, 5
- [11] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. pages 5828–5839, 2017. 4, 5
- [12] Angela Dai and Matthias Nießner. 3dmv: Joint 3d-multi-view prediction for 3d semantic scene segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 452–468, 2018. 3
- [13] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015. 2
- [14] Lizhao Gao, Bo Wang, and Wenmin Wang. Image captioning with scene-graph based semantic concepts. In *Proceedings of the 2018 10th International Conference on Machine Learning and Computing*, pages 225–229, 2018. 2
- [15] Ankit Goyal, Kaiyu Yang, Dawei Yang, and Jia Deng. Rel3d: A minimally contrastive benchmark for grounding spatial relations in 3d. *Advances in Neural Information Processing Systems*, 33, 2020. 1
- [16] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2827–2836, 2016. 3
- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 6
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *NeurIPS Workshops*, 2014. 3, 7, 8
- [19] Wenbo Hu, Hengshuang Zhao, Li Jiang, Jiaya Jia, and Tien-Tsin Wong. Bidirectional projection network for cross dimension scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14373–14382, 2021. 3
- [20] Pin-Hao Huang, Han-Hung Lee, Hwann-Tzong Chen, and Tyng-Luh Liu. Text-guided graph neural networks for referring 3d instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1610–1618, 2021. 1, 2
- [21] Zhen Huang, Xu Shen, Jun Xing, Tongliang Liu, Xinmei Tian, Houqiang Li, Bing Deng, Jianqiang Huang,

- and Xian-Sheng Hua. Revisiting knowledge distillation: An inheritance and exploration framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3579–3588, 2021. 7, 8
- [22] Maximilian Jaritz, Jiayuan Gu, and Hao Su. Multi-view pointnet for 3d scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3
- [23] Hu Jingpeng, Li Zhuo, Chen Zhihong, Li Zhen, Wan Xiang, and Chang Tsung-Hui. Graph enhanced contrastive learning for radiology findings summarization. In *Proceedings of Association for Computational Linguistics (ACL)*, 2022. 2
- [24] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4565–4574, 2016. 1, 2
- [25] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015. 1
- [26] Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. Dense relational captioning: Triple-stream networks for relationship-based captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6271–6280, 2019. 1, 2
- [27] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yanis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 4
- [28] Jean Lahoud and Bernard Ghanem. 2d-driven 3d object detection in rgb-d images. In *Proceedings of the IEEE international conference on computer vision*, pages 4622–4630, 2017. 3
- [29] Jianan Li, Yunchao Wei, Xiaodan Liang, Fang Zhao, Jianshu Li, Tingfa Xu, and Jiashi Feng. Deep attribute-preserving metric learning for natural language object retrieval. pages 181–189, 2017. 2
- [30] Xiangyang Li, Shuqiang Jiang, and Jungong Han. Learning object context for dense captioning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8650–8657, 2019. 1
- [31] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 6
- [32] Yueh-Cheng Liu, Yu-Kai Huang, Hung-Yueh Chiang, Hung-Ting Su, Zhe-Yu Liu, Chin-Tang Chen, Ching-Yu Tseng, and Winston H Hsu. Learning from 2d: Pixel-to-point knowledge transfer for 3d pretraining. *arXiv preprint arXiv:2104.04687*, 2021. 3, 8
- [33] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 375–383, 2017. 1, 2
- [34] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural baby talk. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7219–7228, 2018. 2
- [35] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002. 6
- [36] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. Enet: A deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147*, 2016. 2
- [37] Charles R Qi, Xinlei Chen, Or Litany, and Leonidas J Guibas. Imvotenet: Boosting 3d object detection in point clouds with image votes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4404–4413, 2020. 3
- [38] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. pages 9277–9286, 2019. 6
- [39] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018. 3
- [40] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. pages 5099–5108, 2017. 4
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. 4
- [42] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7008–7024, 2017. 5
- [43] Suraj Srinivas and François Fleuret. Knowledge transfer with jacobian matching. In *International Confer-*

- ence on Machine Learning*, pages 4723–4731. PMLR, 2018. [3](#)
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. pages 5998–6008, 2017. [2](#), [3](#), [4](#), [8](#)
- [45] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. [6](#)
- [46] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015. [1](#), [2](#)
- [47] Lichen Wang, Jiayang Wu, Shao-Lun Huang, Lizhong Zheng, Xiangxiang Xu, Lin Zhang, and Junzhou Huang. An efficient approach to informative feature extraction from multimodal data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5281–5288, 2019. [3](#)
- [48] Chenfeng Xu, Shijia Yang, Bohan Zhai, Bichen Wu, Xiangyu Yue, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with pretrained 2d convnets. *arXiv preprint arXiv:2106.04180*, 2021. [3](#)
- [49] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 244–253, 2018. [3](#)
- [50] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015. [1](#), [2](#), [6](#), [7](#)
- [51] Linjie Yang, Kevin Tang, Jianchao Yang, and Li-Jia Li. Dense captioning with joint inference and visual context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2193–2202, 2017. [2](#)
- [52] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10685–10694, 2019. [2](#)
- [53] Zhengyuan Yang, Songyang Zhang, Liwei Wang, and Jiebo Luo. Sat: 2d semantics assisted training for 3d visual grounding. 2021. [3](#), [8](#)
- [54] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 684–699, 2018. [2](#)
- [55] Shanxin Yuan, Bjorn Stenger, and Tae-Kyun Kim. Rgb-based 3d hand pose estimation via privileged learning with depth images. *arXiv preprint arXiv:1811.07376*, 2018. [3](#)
- [56] Zhihao Yuan, Xu Yan, Yinghong Liao, Ruimao Zhang, Sheng Wang, Zhen Li, and Shuguang Cui. Instancerefer: Cooperative holistic understanding for visual grounding on point clouds through instance multi-level contextual referring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1791–1800, 2021. [1](#), [2](#)
- [57] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *ICLR*, 2017. [3](#)
- [58] Long Zhao, Xi Peng, Yuxiao Chen, Mubbasir Kapadia, and Dimitris N Metaxas. Knowledge as priors: Cross-modal knowledge generalization for datasets without superior knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6528–6537, 2020. [3](#)
- [59] Chaoda Zheng, Xu Yan, Jiantao Gao, Weibing Zhao, Wei Zhang, Zhen Li, and Shuguang Cui. Box-aware feature enhancement for single object tracking on point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13199–13208, 2021. [3](#)
- [60] Chaoda Zheng, Xu Yan, Haiming Zhang, Baoyuan Wang, Shenghui Cheng, Shuguang Cui, and Zhen Li. Beyond 3d siamese tracking: A motion-centric paradigm for 3d single object tracking in point clouds. *arXiv preprint arXiv:2203.01730*, 2022. [3](#)
- [61] Dave Zhenyu Chen, Qirui Wu, Matthias Nießner, and Angel X Chang. D3net: A speaker-listener architecture for semi-supervised dense captioning and visual grounding in rgb-d scans. *arXiv e-prints*, pages arXiv–2112, 2021. [2](#)