

PointMCD: Boosting Deep Point Cloud Encoders via Multi-view Cross-modal Distillation for 3D Shape Recognition

Qijian Zhang, Junhui Hou, *Senior Member, IEEE*, and Yue Qian

Abstract—As two fundamental representation modalities of 3D objects, 3D point clouds and multi-view 2D images record shape information from different domains of geometric structures and visual appearances. In the current deep learning era, remarkable progress in processing such two data modalities has been achieved through respectively customizing compatible 3D and 2D network architectures. However, unlike multi-view image-based 2D visual modeling paradigms, which have shown leading performance in several common 3D shape recognition benchmarks, point cloud-based 3D geometric modeling paradigms are still highly limited by insufficient learning capacity due to the difficulty of extracting discriminative features from irregular geometric signals. In this paper, we explore the possibility of boosting deep 3D point cloud encoders by transferring visual knowledge extracted from deep 2D image encoders under a standard teacher-student distillation workflow. Generally, we propose PointMCD, a unified multi-view cross-modal distillation architecture, including a pretrained deep image encoder as the teacher and a deep point encoder as the student. To perform heterogeneous feature alignment between 2D visual and 3D geometric domains, we further investigate visibility-aware feature projection (VAFP), by which point-wise embeddings are reasonably aggregated into view-specific geometric descriptors. By pair-wisely aligning multi-view visual and geometric descriptors, we can obtain more powerful deep point encoders without exhausting and complicated network modification. Experiments on 3D shape classification, part segmentation, and unsupervised learning strongly validate the effectiveness of our method. *The code and data will be publicly available at <https://github.com/keeganhk/PointMCD>.*

Index Terms—3D point cloud, multi-view images, cross-modal, knowledge distillation, 3D shape recognition

I. INTRODUCTION

DRIVEN by the recent advancements and popularization of 3D acquisition and perception technologies, 3D shape recognition has been attracting increasingly growing attention in both industry and academia. Generally, different from its scene-level counterparts such as place recognition [1], indoor [2]–[4] or outdoor [5]–[7] semantic segmentation, and LiDAR (i.e., light detection and ranging) driving environment object detection [8]–[10], shape recognition tasks focus on object-level models and thus require a more fine-grained understanding of 3D geometry, which pose special challenges and differentiated emphases in this field.

In contrast to 2D visual signals that are typically represented by images captured from ordinary optical cameras, 3D shape

information can be recorded by multiple alternative modalities generated by rich varieties of 3D sensors/scanners, which are widely deployed in different application scenarios. In the deep learning era, depending on different representation modalities of 3D objects, mainstream learning pipelines can be classified into voxel-based [11]–[15], image-based [16]–[23], and point-based [24]–[33] processing paradigms. Generally, voxel-based models employ regular volumetric grids to describe the spatial occupancy status of irregular geometric structures, such that standard 3D convolutional neural network (CNN) architectures can be naturally and seamlessly introduced. Unfortunately, due to the cubic growth of computational complexity and memory footprint, these methods are not suitable for dealing with high-resolution volumes with fine details, despite the exploration of more complicated adaptive and hierarchical spatial indexing strategies [13]–[15]. Instead of faithfully recording 3D spatial structures in the original geometric space, image-based models focus on visual appearances of object surfaces by rendering a collection of multi-view 2D images from different viewpoints. Benefiting from the maturity of powerful 2D modeling architectures [34]–[37] and the availability of large-scale annotated image databases [38]–[40], such multi-view learning paradigm has demonstrated dominating performance in common shape recognition benchmarks [12], [41] for diverse tasks of classification, retrieval, and pose estimation. Meanwhile, there also exist several aspects of limitations such as ignorance of interior spatial structures, difficulty in topology analysis, and inevitable loss of detailed surface textures.

More recently, point-based models that can directly work on unstructured point clouds have gained popularity. As the most straightforward representation modality for geometric signals, point clouds serve as raw outputs of many 3D data acquisition systems and faithful records of the original spatial information. However, different from images/voxels defined on regular grid domains, point cloud data are characterized by irregularity and unorderedness, causing much difficulty in designing expressive feature extraction operators with sufficient modeling capacity. Despite the proliferation of various fancy designs of deep set architectures in previous works, there still exists a lot of room for further performance boost. In practice, there is no unified modeling paradigm for 3D shape understanding, as each data modality and model architecture uniquely have different merits and limitations.

In fact, image-driven 2D visual modeling and point-driven 3D geometric modeling have a degree of complementarity. On one hand, image-based models built upon well-developed 2D

This project was supported by the Hong Kong Research Grants Council under Grants 11202320 and 11218121. *Corresponding author: Junhui Hou*

All authors are with the Department of Computer Science, City University of Hong Kong, Hong Kong SAR. Email: qijizhang3-c@my.cityu.edu.hk; jh.hou@cityu.edu.hk; yueqian4-c@my.cityu.edu.hk;

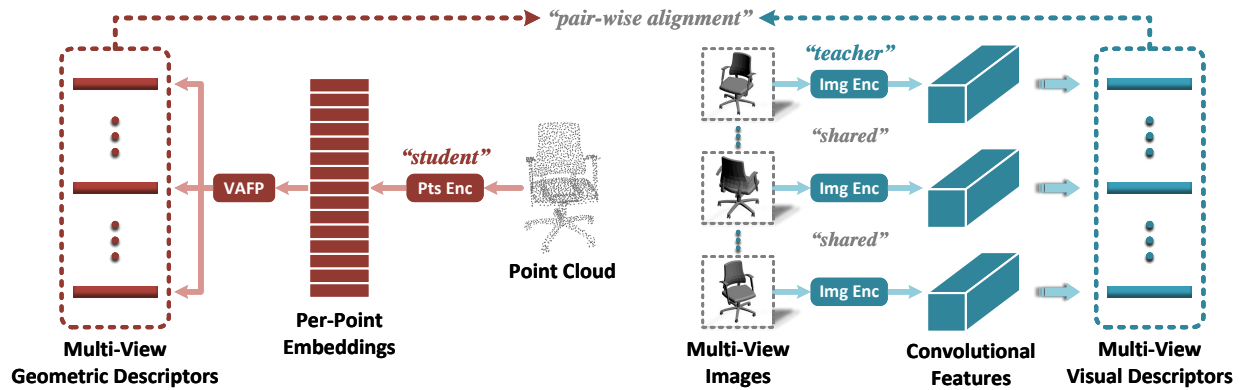


Fig. 1. Flowchart of the proposed PointMCD for multi-view cross-modal distillation. In the (right) teacher branch, a collection of multi-view images are fed into shared deep 2D CNNs to generate convolutional feature maps, which are further converted into vectorized multi-view visual descriptors. In the (left) student branch, the target deep 3D point encoder consumes the input point cloud and extracts high-dimensional point-wise embeddings, which are further aggregated via the proposed novel visibility-aware feature projection (VAFP) mechanism into vectorized multi-view geometric descriptors. Intuitively, VAFP selects points that are visible from a certain viewpoint, and then collects feature embeddings of the selected visible points for further aggregation. Thus, by pair-wisely aligning multi-view visual and geometric descriptors, the student model can be guided to learn more discriminative point features.

CNNs and sufficiently-large image datasets outperform point-based models in terms of network capacity and generalization ability. On the other hand, since 3D point clouds intrinsically contain more complete geometric information than multi-view 2D image renderings, we may expect that point-based models indicate a more promising direction with greater potential and convenience.

In this paper, instead of investigating new backbone network architectures for point feature extraction, we seek to construct a unified processing pipeline that can effectively boost existing mainstream deep point encoders to learn more discriminative features. Technically, we propose PointMCD that implements multi-view cross-modal distillation by transferring knowledge from a pretrained deep 2D image encoder as the teacher to a target deep 3D point encoder as the student, which is a highly non-trivial problem due to significant domain gaps in terms of both data modality and model structure. Fig. 1 presents the block diagram of the overall workflow. In the teacher branch, we feed a collection of multi-view images rendered from a pre-defined set of camera poses into a pretrained deep 2D image encoder, from which we obtain vectorized multi-view visual descriptors serving as teacher knowledge to be transferred. In the student branch, the target deep 3D point encoder consumes the corresponding point cloud model and produces point-wise embeddings. After that, to facilitate multi-view heterogeneous feature alignment between visual and geometric domains, we introduce a novel visibility-aware feature projection (VAFP) mechanism that aggregates vectorized multi-view geometric descriptors on the basis of view-specific per-point visibility computed under the same set of pre-defined camera poses. Finally, the distillation process can be achieved by pair-wisely imposing alignment constraints on the produced multi-view visual and geometric descriptors.

To verify the *effectiveness* and *universality* of our proposed PointMCD, we select three classic deep point encoders [24], [25], [30] as well as a more recent state-of-the-art framework named CurveNet [32] as the baseline student models, and conduct experiments on three common benchmark tasks, including

3D shape classification, part segmentation, and unsupervised learning, where we can observe stable and encouraging performance gains.

In summary, our main contributions are three-fold:

- we propose a unified multi-view cross-modal distillation framework, namely PointMCD, which aims at boosting deep 3D point cloud encoders without additional efforts on model structure modification;
- to facilitate multi-view heterogeneous (visual-geometric) feature alignment, we develop a novel VAFP mechanism, which reasonably aggregates point-wise embeddings into view-specific geometric descriptors on the basis of physical visibility property; and
- we achieve a highly encouraging performance boost on a rich variety of baseline models and downstream application scenarios.

Overall, the main technical novelty of this work lies in that we skillfully integrate two heterogeneous 3D data representation structures – 2D multi-view images and 3D point clouds – into a unified cross-modal knowledge distillation pipeline via customizing view-specific visual-geometric feature alignment. Our exploration implies a generic and promising direction for enhancing the learning capacity of point-based 3D geometric modeling architectures.

The remainder of this paper is organized as follows. In Sec. II, we detailedly review three aspects of closely-related works. In Sec. III, we introduce the overall workflow of the proposed PointMCD learning framework. We report experimental results in Sec. IV, including evaluations on three benchmark tasks of 3D shape recognition, comparisons with baseline models and other state-of-the-art methods, and carefully designed ablation studies. In Sec. V, we re-emphasize the design philosophy of our specific technical implementations and further discuss new insights revealed by this work. In the end, we summarize our paper in Sec. VI.

II. RELATED WORK

In this section, we start by reviewing both multi-view image-based visual modeling paradigms in Sec. II-A and point-based geometric modeling paradigms in Sec. II-B, which serve as the two main ingredients in our proposed cross-modal processing pipeline. In II-C, we discuss the current research progress in knowledge distillation, especially for the most closely related works with focuses on big-gap cross-modal transfer between 2D and 3D domains.

A. Multi-View 2D CNNs for 3D Shape Recognition

As a natural extension of 2D deep learning frameworks for image processing [42], [43], multi-view 3D shape modeling is typically built upon various variants of multi-input 2D CNNs.

Pioneered by MVCNN [16] that separately consumes multi-view images rendered from multiple virtual camera poses and generates vectorized global shape signature through cross-view max-pooling, many follow-up studies are devoted to designing more advanced learning strategies of inter-view interaction and viewpoint selection. GVCNN [17] proposes a three-level hierarchical correlation modeling framework, which adaptively groups multi-view feature embeddings into separate clusters. MHBN [18] and RelationNet [20] further exploit patch-level interaction to enrich inter-view relationships. RotationNet [19] treats viewpoint indices as learnable latent variables and tends to jointly estimate object poses and semantic categories. EMV [21] presents a group convolution approach that operates on a discrete subgroup of rotation groups, which enable to extract rotation-equivalent shape descriptors. View-GCN [22] regards viewpoints as graph nodes to construct a directed view graph, on which graph convolution can be applied to learn inter-view relations. MVTN [23] introduces the differentiable rendering techniques to implement adaptive regression of optimal camera poses in an end-to-end trainable manner.

Generally, image-based methods have demonstrated leading performance in common shape recognition benchmarks [12], [41] for tasks like classification, retrieval, and pose estimation. Nevertheless, adapting such image-based learning paradigm to topology analysis and some fine-grained prediction scenarios (e.g., segmentation, normal estimation) is known to be highly non-trivial and cumbersome. Besides, interior spatial structures are totally ignored during surface rendering.

B. Deep Set Architectures for 3D Point Cloud Processing

In real-world applications, point cloud data play a key role in numerous low-level and high-level task scenarios [44]–[51]. Different from the maturity of 2D deep learning frameworks for image processing, the exploration of deep set architectures for 3D point cloud modeling is still at its fast-growing stage.

Pioneered by PointNet [24] that adopts point-wisely shared multi-layer perceptrons (MLPs) for permutation-invariant feature embedding, deep set architectures that directly operate on unstructured 3D point clouds have rapidly gained popularity in the geometry processing community. Inheriting the successful design experience of 2D CNNs, PointNet++ [25] incorporates local neighborhood aggregation and further adopts hierarchical

feature abstraction by applying progressive downsampling. A rich variety of highly-specialized point cloud feature aggregation operators have been developed in later studies. Typically, [27], [29], [52]–[54] mimic standard convolutions by learning adaptive weights to implement kernel matching. Following the design philosophy of specializing point convolutions, [28], [55]–[57] are devoted to designing various more complicated point feature aggregation strategies to enhance network capacity. Another classic work DGCNN [30] investigates a graph-based dynamic feature updating mechanism, which can capture global contextual information in a flexible manner. [58]–[60] explore learning-based, instead of heuristic, subset selection techniques to adaptively preserve the most informative points after downsampling. Driven by the latest trend in 2D vision and language community, transformer-style architectures [31], [61]–[64] are also adapted for point cloud modeling, and their corresponding self-supervised pre-training schemes [65], [66] show impressive performances in various downstream tasks.

Thanks to the recent academic attention, we have witnessed continuous progress in learning deep features directly from 3D point clouds. However, point-based methods still suffer from insufficient modeling capacity and relatively sub-optimal task performance. Moreover, since 3D labeling can be much more difficult and expensive than 2D image annotation, existing 3D shape repositories are of small scale and limited richness.

C. Knowledge Transfer between 2D and 3D Modalities

Ever since Hinton *et al.* [67] firstly proposed the processing paradigm of knowledge distillation, there have been numerous academic explorations of diverse technical variants and wide deployment in real applications. However, as analyzed in [68], transferring knowledge across multi-modalities still remains a highly challenging problem, especially when the domain gap is big (e.g., lacking of paired samples), and such practice turns to be even more non-trivial when conducted between 2D and 3D domains with an obvious dimensionality gap.

Typically, xMUDA [69] proposes to achieve unsupervised domain adaptation from the 2D source domain of single-view camera images to the 3D target domain of LiDAR point clouds by aligning outputs of 2D and 3D branches according to pixel-point correspondences. PPKT [70] constructs a 3D pretraining framework to leverage 2D pretrained knowledge by applying the contrastive learning strategy on positive and negative pixel-point pairs. In an opposite transferring direction, Pri3D [71] explores the potential of 3D-guided contrastive pretraining for boosting 2D perception capability. In addition to maintaining feature consistency between paired 2D pixels and 3D points, this work also aims at learning invariant pixel-wise descriptors across the image scans captured from different viewpoints. A more flexible 3D-to-2D distillation framework can be found in [72], which aligns statistical distributions of convolutional features from 2D and 3D CNNs through specialized dimension normalization techniques. To get rid of the dependence on fine-grained correspondences between 2D and 3D modalities that are usually expensive to acquire, this work further explores a semantics-aware adversarial training scheme, which enables to tackle unpaired 2D images and 3D volumetric grids.

Conclusively, existing studies on cross-modal (2D and 3D) knowledge transfer mainly focus on scene-level parsing, owing to the convenience of computing fine-grained correspondences (i.e., pixel-point/voxel pairs) from the corresponding benchmark datasets [4], [10], [73] with synchronized and calibrated camera and LiDAR device configuration, and actually we also need to shift more attention to shape recognition (object-level) scenarios. More importantly, heterogeneous feature alignment, which plays a key role in cross-modal distillation, still needs to be implemented in a more effective way that better boosts the distilled model.

III. PROPOSED METHOD

This section introduces the overall working mechanism and specific technical implementations of the proposed PointMCD. To facilitate understanding, we begin with an overview of our framework in Sec. III-A. Then we present general formulations of deep image encoders and deep point encoders respectively in Sec. III-B and Sec. III-C, based on which we construct a unified cross-modal distillation pipeline driven by multi-view visibility-aware feature alignment in Sec. III-D. In the end, we summarize the overall loss function and training strategy in Sec. III-E.

A. Problem Overview

We aim to transfer cross-modal knowledge from pretrained deep image encoders that extract 2D visual features to deep point encoders that extract 3D geometric embeddings.

Formally, we consider a given 3D point cloud $\mathcal{P} \in \mathbb{R}^{N \times 3}$ containing N spatial points and a collection of 2D multi-view images $\{\mathcal{I}_k \in \mathbb{R}^{H \times W \times 3}\}_{k=1}^K$ capturing view-specific visual appearances of the original 3D shape geometry when rendered from K pre-defined camera positions $\{\mathbf{c}_k = (\theta_k^{az}, \theta_k^{el}, \mu)\}_{k=1}^K$, where θ_k^{az} and θ_k^{el} denote azimuth and elevation angles relative to the object centroid at viewpoint \mathbf{c}_k and μ is specified as a constant observation distance.

Following the standard teacher-student architectural design, the proposed PointMCD is composed of a 2D teacher branch $\mathcal{B}_t = \{\mathcal{M}_t, \mathcal{H}_t\}$ and a 3D student branch $\mathcal{B}_s = \{\mathcal{M}_s, \mathcal{H}_s\}$, in which \mathcal{M}_t and \mathcal{M}_s respectively represent the pretrained deep image encoder to be transferred and the deep point encoder to be distilled. We perform distillation by aligning geometric embeddings learned by \mathcal{M}_s with visual features exported from \mathcal{M}_t . The subsequent task-specific head networks are denoted as \mathcal{H}_t and \mathcal{H}_s , which are excluded from the actual distillation process. Functionally, we emphasize that PointMCD serves as a universal processing pipeline that is compatible with generic deep set architectures.

Remark. In particular, it is worth reminding that our method should be differentiated from multi-modal fusion frameworks that consume data of multi-modalities as inputs during both training and inference stages. In other words, the core concern of this work lies in information transfer from source domain to target domain, instead of multi-source feature fusion and interaction (e.g., [74]).

B. Teacher Network for 2D Image Modeling

Driven by the maturity of deep convolutional architectures, we resort to powerful 2D CNNs in the visual modeling branch for image feature extraction. On one hand, the proliferation of numerous off-the-shelf deep architectures [34], [35], [75], [76] allows us to conveniently deploy 2D image backbone encoders without additional efforts on model design. On the other hand, benefiting from the common practice of large-scale pretraining (e.g., on ImageNet [38]), the resulting deep feature extractors demonstrate satisfactory generalization ability when fine-tuned on downstream visual recognition application scenarios. These excellent properties make 2D CNNs an optimal choice of the teacher model of the deep image encoder \mathcal{M}_t .

Formally, the teacher encoder consumes multi-view images in parallel as inputs and correspondingly produce a collection of high-dimensional 2D convolutional feature maps. A general abstraction of the teacher model behaviors is formulated as

$$\{\mathcal{V}_k\}_{k=1}^K = \mathcal{M}_t(\{\mathcal{I}_k\}_{k=1}^K), \quad (1)$$

where $\mathcal{V}_k \in \mathbb{R}^{H' \times W' \times C_t}$ with subscript of k represents a view-specific 2D convolutional feature map extracted from image \mathcal{I}_k . This corresponds to partial encoding of visual appearances of the original 3D geometric shape when observed from the viewpoint \mathbf{c}_k .

After backbone feature extraction, we feed $\{\mathcal{V}_k\}_{k=1}^K$ into a global average pooling (GAP) layer to deduce a collection of vectorized multi-view visual descriptors $\{\mathbf{v}_k \in \mathbb{R}^{C_t}\}_{k=1}^K$ that serve as the teacher knowledge to be transferred into the target student model in the actual distillation process.

C. Student Network for 3D Point Cloud Modeling

In contrast to the maturity of 2D image modeling networks, 3D deep learning on raw point clouds is still a newly-emerging yet fast-growing research area. Comparatively, existing point cloud learning networks are faced with the following aspects of restrictions:

- Due to the difficulty of collecting and labeling 3D shape models, the current geometry community still lacks large-scale richly-annotated 3D shape repositories comparable to their 2D counterparts (e.g., [38], [39]).
- Limited by the insufficiency of training data, mainstream point cloud networks are actually designed to be far from “deep”, in order to relieve parameter overfitting.
- Different from the 2D vision community where “pretrain-and-finetune” has become a default practice, mainstream point cloud networks are still trained from scratch.

Consequently, point-based learning paradigms demonstrate insufficient feature learning capacity, and unstable and sub-optimal downstream task performance.

Though there is no unified point feature extraction schemes, we can formulate a general abstraction of the actual behaviors of various existing deep set architectures. Formally, the student deep 3D point encoder \mathcal{M}_s consumes an input set of 3D points and outputs high-dimensional point embedding vectors (from the last abstraction level), which can be described as

$$\mathcal{G} = \mathcal{M}_s(\mathcal{P}), \quad (2)$$



Fig. 2. Illustration of multi-view visibility checking, where we apply hidden point removal [77] to select visible points observed from the specified viewpoint. In the first and second rows, we respectively display the rendered multi-view images and the selected visible points (subsets of the complete point cloud) generated from the same pre-defined views. For better visual effect, we further provide visualizations of the selected visible points when observed from the opposite azimuth angles in the third row.

where $\mathcal{G} \in \mathbb{R}^{N_a \times C_s}$ represents a set of C_s -dimensional feature vectors extracted from the original N input points. The actual learning process operating on \mathcal{P} can be further detailed as

$$\mathcal{M}_s(\mathcal{P}) = \varphi_{\text{bm}}(\mathcal{A}), \quad (3)$$

where $\mathcal{A} \in \mathbb{R}^{N_a \times 3}$ denotes N_a (what we call) feature anchors, and $\varphi_{\text{bm}}(\cdot)$ represents a learnable feature mapping that pointwisely embeds each of the 3D feature anchor point (in \mathcal{A}) to its corresponding feature vector (in \mathcal{G}). Here, we propose the key concept of “feature anchors” for convenience of forming a concise and general description for the student behaviors. In practice, depending on specific network architectures and task properties, there exist different definitions and implementations of the feature anchors \mathcal{A} . For example, in hierarchical learning frameworks where $N_a < N$, feature anchors \mathcal{A} can be a downsampled subset of input points (as conducted in [25], [28], [57]) or a sparse set of adaptively learned and generated anchor positions (as conducted in [55], [59], [60]). By contrast, many other approaches [24], [30], [52] also choose to deduce point-wise embeddings for all input points, such that $N_a = N$, which means that the feature anchors \mathcal{A} are exactly the input point cloud \mathcal{P} itself.

D. Multi-View Visibility-aware Feature Alignment

Conventional development protocols of knowledge distillation studies mainly focus on homogeneous model architectures (e.g., deep and shallow CNNs) and the same data structure but different contents (e.g., RGB and depth images). Differently, here we are faced with significant heterogeneity of both model and data across 2D and 3D domains. Essentially, our goal is to establish connections in a reasonable manner between the K view-specific visual descriptors $\{\mathbf{v}_k\}_{k=1}^K$ (extracted from \mathcal{M}_t) corresponding to viewpoints $\{\mathbf{c}_k\}_{k=1}^K$ and the N_a anchor embedding vectors in \mathcal{G} (extracted from \mathcal{M}_s).

Intuitively, each individual \mathbf{v}_k encodes partial shape appearance that is visible from the viewpoint \mathbf{c}_k , while point features of \mathcal{G} capture complete 3D geometry information of the whole

shape scope. Such asymmetry strongly drives us to investigate a symmetric visual-geometric feature alignment procedure via view-specific feature projection. Formally, we attempt to distill the teacher knowledge into the student model by imposing the following pair-wise alignment constraint

$$\mathcal{L}_{\text{dist}} = \sum_{k=1}^K \|\mathbf{v}_k - \phi_k(\mathcal{G})\|_1, \quad (4)$$

where $\phi_k(\cdot)$ denotes a learnable feature projector aggregating point embeddings in \mathcal{G} into a vectorized geometric descriptor $\mathbf{g}_k \in \mathbb{R}^{C_t}$, which is supposed to particularly encode the partial shape geometry within the field of view when observed from the corresponding viewpoint \mathbf{c}_k .

Thus, by view-wisely aligning each pair of visual-geometric descriptors in $\{\mathbf{v}_k\}_{k=1}^K$ and $\{\mathbf{g}_k = \phi_k(\mathcal{G})\}_{k=1}^K$, the multi-view teacher knowledge can be adequately exploited for distillation in an interpretable manner.

VAFP Mechanism. Obviously, under our proposed knowledge transfer framework, the core problem is to design an effective view-specific geometric feature projector ϕ_k . To this end, here we design a novel VAFP mechanism to generate $\{\mathbf{g}_k\}_{k=1}^K$ from point embeddings in \mathcal{G} on the basis of point-wise visibility of feature anchors \mathcal{A} .

To implement visibility checking for each of the N_a feature anchor points in \mathcal{A} , we employ a simple and fast hidden point removal algorithm [77], i.e., a classic computational approach proposed earlier in the computer graphics community. Fig. 2 visually illustrates the selected visible points corresponding to varying manually specified viewpoints. This geometrically-meaningful selection mechanism enables us to aggregate the N_a feature anchor points in \mathcal{A} view-wisely into K independent visible subsets $\{\tilde{\mathcal{A}}_k \in \mathbb{R}^{N_k \times 3}\}_{k=1}^K$ with respect to viewpoints $\{\mathbf{c}_k\}_{k=1}^K$, where we have $\tilde{\mathcal{A}}_k \subseteq \mathcal{A}$.

Furthermore, based on the one-to-one mapping relationships (i.e., row-wise correspondence) between feature anchor points in \mathcal{A} and their high-dimensional embeddings in \mathcal{G} , we further obtain K view-specific clusters of feature embeddings denoted




Fig. 3. Illustration of the teacher network pretrained for 2D multi-view image-based 3D shape classification. Given a collection of multi-view images at the input end, we extract the corresponding vectorized (512-dimensional) visual descriptors, which are aggregated via view pooling [16] and then mapped to predict category scores. In addition, we make independent predictions from each single-view descriptor via side-output supervision.

as $\{\tilde{\mathcal{G}}_k \in \mathbb{R}^{N_k \times C_s}\}$, where we have $\tilde{\mathcal{G}}_k \subseteq \mathcal{G}$. More concretely, for viewpoint \mathbf{c}_k , suppose that the i^{th} point in $\tilde{\mathcal{A}}_k$ corresponds to the j^{th} point in \mathcal{A} , then the j^{th} embedding vector in \mathcal{G} will be collected into $\tilde{\mathcal{G}}_k$.

Intuitively, the above view-specific feature grouping process simply finds the embedding vectors of the visible anchor points and then collect them together. After that, we can deduce the desired multi-view geometric descriptors $\{\mathbf{g}_k\}_{k=1}^K$ by

$$\mathbf{g}_k = \phi_k(\mathcal{G}) = \tilde{h}_{\text{align}}(\rho_{\max}(\tilde{\mathcal{G}}_k)), \quad (5)$$

where ρ_{\max} means applying channel-wise max-pooling on $\tilde{\mathcal{G}}_k$ to generate a vectorized C_s -dimensional codeword, and \tilde{h}_{align} is implemented as a single fully-connected (FC) layer, which aligns the number of feature channels (from C_s to C_t) so that we can impose L_1 loss between \mathbf{v}_k and \mathbf{g}_k (as formulated in Eq. (4)). From a different perspective, the existence of \tilde{h}_{align} also indicates that we tend to perform visual-geometric feature alignment in another hidden space, instead of directly forcing point cloud features to be close to image features.

E. Training Strategy

Just like standard knowledge distillation pipelines, given a specific 3D shape recognition task, we start by pretraining the teacher branch \mathcal{B}_t on 2D multi-view images before distillation. After that, we train the target student branch \mathcal{B}_s on 3D point clouds while in the meantime transferring knowledge exported from the fixed deep image encoder \mathcal{M}_t into the deep point cloud encoder \mathcal{M}_s .

The overall training objective $\mathcal{L}_{\text{overall}}$ can be formulated as a weighted summation of the task loss $\mathcal{L}_{\text{task}}$ and the distillation constraint $\mathcal{L}_{\text{dist}}$ (Eq. (4)) as given in the following

$$\mathcal{L}_{\text{overall}} = \omega_t \cdot \mathcal{L}_{\text{task}} + \omega_d \cdot \mathcal{L}_{\text{dist}}, \quad (6)$$

where we empirically set $\omega_t = 0.1$ and $\omega_d = 1/K$ in all our experiments. Functionally, $\mathcal{L}_{\text{task}}$ drives the optimization of the whole student learning branch, including both the encoder \mathcal{M}_s and the head \mathcal{H}_s , while $\mathcal{L}_{\text{dist}}$ only has impact on the encoder \mathcal{M}_s by constraining its feature outputs. In particular, for some deep point cloud encoders (e.g., [24]) that require optimizing auxiliary regularizers, we uniformly conform to their original weighting factors applied on regularization terms.

TABLE I
PERFORMANCE COMPARISON OF SHAPE CLASSIFICATION ON MODELNET40 DATASET IN TERMS OF OVERALL ACCURACY WITHOUT ANY INFERENCE VOTING. NOTE THAT “*” MEANS NORMAL VECTORS ARE EXPLOITED AS ADDITIONAL INPUTS.

Method	Data Modality	OAcc (%)
VoxNet [24]	Voxels; 32^3	85.9
O-CNN [14]	Voxels; 64^3	89.9
VRN-Single [78]	Voxels; 32^3	91.3
MVCNN [16]	Images; $80 \times$	90.1
MVCNN-New [79]	Images; $12 \times$	95.0
RotationNet [19]	Images; $20 \times$	97.4
View-GCN [22]	Images; $20 \times$	97.6
PointNet [24]	Points; 1K	89.2
PointNet++ [25]	Points; 1K	90.7
PointNet++ [25] *	Points; 5K	91.9
SpiderCNN [26] *	Points; 1K	92.4
SO-Net [55] *	Points; 5K	93.4
KPCConv [29]	Points; 7K	92.9
DGCNN [30]	Points; 1K	92.9
RS-CNN [28]	Points; 1K	92.9
PACConv [54]	Points; 1K	93.6
[ST] Point-BERT [65]	Points; 8K	93.8
CurveNet [32]	Points; 1K	93.8
[ST] Point-MAE [66]	Points; 8K	94.0
PointMLP [33]	Points; 1K	94.1
MCD-PointNet	Points; 1K	91.1 (+1.9)
MCD-PointNet++	Points; 1K	93.3 (+2.6)
MCD-DGCNN	Points; 1K	93.7 (+0.8)
MCD-CurveNet	Points; 1K	94.3 (+0.5)

IV. EXPERIMENTS

We conducted experiments to validate the effectiveness and universality of the proposed PointMCD architecture on three 3D shape recognition scenarios, including shape classification (Sec. IV-A), part segmentation (Sec. IV-B), and unsupervised learning (Sec. IV-C). Within each sub-section, we introduced the corresponding benchmark dataset, evaluation protocol, and detailed technical implementation. In Sec. IV-D, we provided more comprehensive comparisons and ablation experiments.




Fig. 4. Illustration of the teacher network pretrained for semantic segmentation of 2D rendered images. Instead of jointly learning from multiple views, here we independently perform segmentation on each single view image. After extracting a vectorized (512-dimensional) visual descriptor from the input single-view image, we append an auxiliary categorical vector to indicate object-level semantics and further deploy a stack of deconvolutional layers, following the classic U-Net-style [37] encoder-decoder network design, to predict pixel-wise labels on a full-resolution segmentation map.

A. Evaluations on Shape Classification

As one of the most fundamental benchmark tasks, 3D shape classification has been widely adopted to evaluate the learning capacity of deep set architectures. Generally, this task aims at extracting from an input object point cloud a vectorized global shape signature, which can be further mapped through a stack of FC layers to predict category scores.

We conducted experiments on the ModelNet40 [12] dataset, which totally consists of 12311 synthetic mesh models of man-made 3D objects covering 40 semantic categories. Following the official split, we have 9843 shapes for training and the rest 2468 for testing. We applied Poisson disk sampling (PDS) to uniformly discretize 1024 spatial points from the original mesh faces. Note that we did not preserve ground-truth point-wise normals as additional input attributes.

For the teacher learning branch, we adopted a mature multi-view learning framework (i.e., MVCNN [16]) built upon multi-input CNNs while making a few modifications to both encoder \mathcal{M}_t and head \mathcal{H}_t for performance boost. Fig. 3 illustrates the overall workflow and core components of the teacher learning branch \mathcal{B}_t for multi-view image-based 3D shape classification. Specifically, we introduced MobileNetV2 [76], a light-weight 2D CNN backbone network, as the teacher image encoder \mathcal{M}_t . After backbone feature extraction, view-wise visual descriptors are on one hand aggregated to predict the final class logits and on the other hand separately mapped to form side-outputs. For the rendering pipeline, we used the same processing procedure as recommended in [79], in which $K = 12$, $\theta_k^{az} \in \{k\pi/6\}_{k=1}^{12}$, $\theta_k^{el} = \pi/6$, and $\mu = 1$. After pretraining, the resulting teacher branch achieves 96.7% overall classification accuracy on the ModelNet40 testing set, and its backbone encoder parameters are fixed when applied to distill different students.

In practice, we selected four representative deep point cloud encoders as the target 3D student models (i.e., \mathcal{M}_s), including three classic learning frameworks (PointNet [24], PointNet++ [25], DGCNN [30]) and a more recent state-of-the-art method CurveNet [32]. It is worth mentioning that, during performance comparison, we did not involve the results [28], [32], [33], [54] obtained by extensively voting a large amount of trials during inference, which turns to be unstable and highly cumbersome.

Table I quantitatively compares classification performances of different modeling paradigms and specific learning models, from which we can draw several aspects of useful conclusions. First, it is observed that, despite the continuous progress made by the point cloud community, image-based visual modeling

TABLE II
PERFORMANCE COMPARISON OF OBJECT PART SEGMENTATION ON SHAPE-NETPART DATASET UNDER THE MEASUREMENT OF INSTANCE-AVERAGED MEAN INTERSECTION-OVER-UNION. NOTE THAT “*” MEANS NORMAL VECTORS ARE EXPLOITED AS ADDITIONAL INPUTS.

Method	mIoU (%)
PointNet [24]	83.7
PointNet++ [25] *	85.1
SpiderCNN [26] *	85.3
SO-Net [55] *	84.9
KPConv [29]	86.4
DGCNN [30]	85.1
RS-CNN [28]	85.8
PACConv [54]	86.0
CurveNet [32]	86.6
PointMLP [33]	86.1
MCD-PointNet	85.9 (+2.2)
MCD-PointNet++	86.4 (+1.3)
MCD-DGCNN	86.8 (+1.7)
MCD-CurveNet	87.1 (+0.5)

still outperforms voxel-based and point-based paradigms with significant margins. Second, it turns out that PointMCD is able to produce stable and obvious performance boost for various types of deep point encoders. Notably, the distilled PointNet++ and DGCNN are boosted to achieve 93.3% and 93.7% overall accuracy, which are comparable to more complicated network structures investigated in follow-up works. Even for CurveNet, a much more powerful baseline model that originally achieves 93.8% state-of-the-art performance, our distillation framework still brings satisfactory gains. Third, an interesting observation is that the actual performance gap between some earlier classic methods (e.g., [25], [30]) and more recently proposed methods (e.g., [32], [33]) has been largely reduced via the distillation of PointMCD. This phenomenon implies that the actual potential of some previous network design paradigms for point cloud learning may not have been sufficiently explored and realized.

B. Evaluations on Part Segmentation

To further demonstrate the effectiveness of our approach on more fine-grained per-point semantic understanding scenarios, we experimented with part segmentation of 3D objects on the ShapeNetPart [80] dataset, consisting of 16881 point-wisely annotated 3D object models covering totally 50 different parts




Fig. 5. Illustration of the teacher network pretrained for single-view image reconstruction. The overall learning framework is designed as a standard deep convolutional image auto-encoder. The encoding stage represents an input image into a compact 512-dimensional feature vector. After feature channel lifting and reshaping, the decoding phase further reconstructs a full-resolution image.

defined on 16 object categories. Following common evaluation protocols, there exist 14007 training models and 2874 testing models, each of which contains 2048 labeled spatial points. Similar to our preceding shape classification experiments, here we did not use ground-truth point-wise normals as additional input attributes, and our inference performance was obtained without any voting strategy.

For the teacher learning branch, it is worth mentioning that, in contrast to global geometry understanding applications such as classification/retrieval in which multi-view 2D image-based processing architectures have already been richly investigated, there is relatively little such effort on 3D shape segmentation scenarios. Hence, instead of resorting to existing frameworks, here we simply designed a standard 2D encoder-decoder [37] convolutional architecture that consumes an input single-view image and accordingly predict a full-resolution segmentation map, as illustrated in Figure 4. More specifically, we selected VGG11 [34] as backbone image encoder \mathcal{M}_t while removing the last spatial max-pooling layer to enlarge the final feature map resolution. Following common practice in previous works on part segmentation, we inserted a categorical vector between the encoded vectorized visual descriptors and the subsequent deconvolutional layers to supplement object-level semantics. The rendering pipeline basically inherits our preceding shape classification experiments, while we additionally expanded the interval of azimuth angles from $\pi/6$ to $\pi/4$ and further added another elevation angle of $-\pi/6$, such that $K = 16$, $\theta_k^{az} \in \{k\pi/4\}_{k=1}^8$, $\theta_k^{el} = \{\pm\pi/6\}$, and $\mu = 1$. For the generation of pixel-wisely annotated ground-truth maps, we projected point-wise labels onto 2D image planes, where another new category of “background” that corresponds to empty pixels during the rendering process need to be added into the original annotation set. During our explorations, we did not quantitatively evaluate the actual segmentation performance of such a teacher learning branch, because there is no viable way to determine semantic labels of interior points located inside shape surface.

We experimented with the same four representative deep set architectures [24], [25], [30], [32] as selected in our preceding shape classification experiments. As reported in Table II, our PointMCD consistently brings obvious performance improvement. In particular, PointNet can be significantly boosted from the original 83.7% to the distilled 85.9%, which is on par with more advanced frameworks [28], [33], [54]. After distillation, the classic DGCNN framework also reaches 86.8% with large gains of 1.7%, which even outperforms the original CurveNet. For the most powerful CurveNet baseline, our approach still brings 0.5% performance boost. These results demonstrate the

TABLE III

PERFORMANCE COMPARISON OF UNSUPERVISED LEARNING (TRANSFER CLASSIFICATION) ON MODELNET40 DATASET. NOTE THAT “†” MEANS DISTILLING VISUAL KNOWLEDGE FROM AN IMAGENET PRETRAINED BACKBONE IMAGE ENCODER WITHOUT TASK-SPECIFIC PRETRAINING OF THE WHOLE TEACHER LEARNING BRANCH ON THE SOURCE DATASET. PARTICULARLY, MULTIPLE DIFFERENT POINT CLOUD BACKBONES ARE INVOLVED IN [81]–[84], AND BELOW WE REPORT THE RESULTS OBTAINED FROM THE POINTNET [24] BACKBONE, WHOSE REPRESENTATION POWER IS BASICALLY CLOSE TO THE POINT CLOUD ENCODER USED IN FOLDINGNET [85].

Method	Data Modality	OAcc (%)
TL-Net [86]	Voxels; 20 ³	74.4
VConv-DAE [87]	Voxels; 30 ³	75.5
3D-GAN [88]	Voxels; 64 ³	83.3
VIP-GAN [89]	Images; 12×	90.2
Latent-GAN [90]	Points; 2K	85.7
STRL [81]	Points; 2K	88.3
FoldingNet [85]	Points; 2K	88.4
CrossPoint [82]	Points; 2K	89.1
3D-PointCapsNet [91]	Points; 2K	89.3
SelfContrast [92]	Points; 2K	89.6
GTIF [93]	Points; 2K	89.7
SelfCorrection [83]	Points; 2K	89.9
ParAE [84]	Points; 1K	90.3
GSIR [94]	Points; 1K	90.4
MCD-FoldingNet	Points; 2K	90.0 (+1.6)
MCD-FoldingNet †	Points; 2K	89.8 (+1.4)

effectiveness of our method on fine-grained per-point labeling tasks, despite that the actual distillation process only operates on global feature vectors without pixel-to-point alignment.

C. Evaluations on Unsupervised Learning

Previous experiments on classification and segmentation are targeted at supervised learning scenarios with requirements of domain-specific data annotations. Taking a step forward, here we further explored the possibility of transferring generic 2D visual knowledge extracted from unsupervised image feature learning frameworks to boost 3D geometric modeling. Practically, we adopted one of the most popular evaluation protocols (known as transfer classification) employed by many previous works [85], [91], [93]. The overall evaluation protocol relies on a source dataset (i.e., ShapeNetCore [41]) for representation learning in an unsupervised manner and another target dataset (i.e., ModelNet40 [12]) for linear probing. More specifically, the first step is to pretrain the unsupervised learning model on the whole source dataset. Then, we directly employ the trained

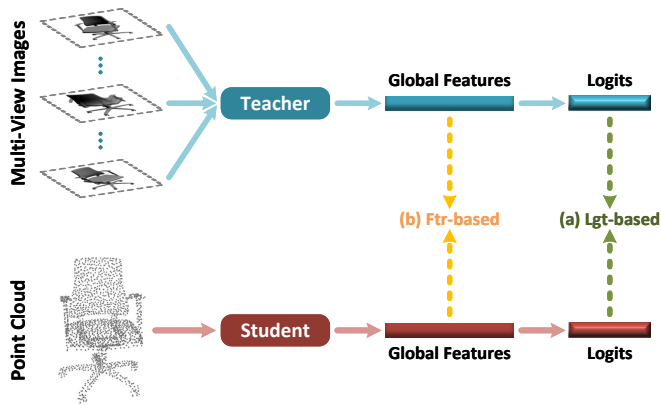


Fig. 6. Illustration of two conventional (a) logit-based and (b) feature-based distillation baselines, which impose alignment constraints over the final class logits and the intermediate global feature vectors, respectively.

model to convert each of the shape models in the whole target dataset into a global feature vector, during which the network parameters are fixed. After that, we train and test a linear support vector machine (SVM) classifier using the extracted feature vectors of all training and testing models. In our experiments, each point cloud in both the source and the target datasets contains 2048 spatial points uniformly discretized from the original mesh faces.

We selected FoldingNet [85], a well-known reconstruction-driven unsupervised point cloud representation learning model, as the student branch, which consists of a graph-based encoder (i.e., \mathcal{M}_s) and a folding-based decoder (i.e., \mathcal{H}_s). Accordingly, we designed a standard convolutional auto-encoder for single image reconstruction, as illustrated in Figure 5. In the encoding stage, we deployed the same backbone encoder as adopted in our preceding part segmentation experiments. The decoding stage comprises a stack of deconvolutional layers to generate a full-resolution image. The rendering pipeline is the same as described in Sec. IV-B to generate multi-view images on the source dataset of ShapeNetCore.

Table III quantitatively compares transfer classification performances of different learning frameworks. It is observed that multi-view image-based modeling paradigm can still achieve dominating overall accuracy. As a representative point-based unsupervised learning model, FoldingNet also performs favorably against other more advanced state-of-the-art approaches. After distillation, its accuracy is boosted from 88.4% to 90.0%, which is comparable to [89] and outperforms [91], [93].

Additionally, we are interested in the potential of distilling generic visual knowledge learned from natural image statistics without task-specific pretraining for the teacher model. To this end, we did not pretrain the whole teacher learning branch on the source dataset for single image reconstruction. Instead, we employed the original VGG11 backbone image encoder with ImageNet pretrained model parameters to produce the teacher knowledge during the distillation process. As shown in the last row of Table III, this variant suffers from slight performance degradation from 90.0% to 89.8%, but still turns to be highly competitive and encouraging, considering the big domain gap between natural images and shape renderings.

TABLE IV
PERFORMANCE VERIFICATION OF CONVENTIONAL LOGIT-BASED AND FEATURE-BASED DISTILLATION BASELINES.

Method	PointMCD	Lgt-based	Ftr-based
<i>MCD-PointNet</i>	91.1	90.2	90.4
<i>MCD-PointNet++</i>	93.3	92.5	92.8
<i>MCD-DGCNN</i>	93.7	93.0	93.2
<i>MCD-CurveNet</i>	94.3	93.8	93.9

D. Ablation Study

In order to facilitate more comprehensive understanding of the whole processing pipeline of the proposed PointMCD, we conducted ablation studies on ModelNet40 classification for distilling the preceding four student architectures [24], [25], [30], [32] under different experimental conditions.

First, we constructed and evaluated two conventional distillation paradigms to validate the necessity and superiority of our PointMCD that is particularly customized for multi-view image-to-point distillation. Second, we investigated different hyperparameter configurations and their influences to the final performance of the whole processing pipeline. Third, one more important experiment was added to evaluate the performance of purely point-deduced baseline pipelines, which can strongly reveal the great potential of transferring knowledge from 2D image domains to 3D point cloud domains.

1) *Conventional distillation paradigm:* In the research area of knowledge distillation, there are two representative learning paradigms, namely response-based [67] and feature-based [95] methods. The former type of works aims at performing vector alignment between the final class logits outputted from the last layers of teacher and student learning branches. For the latter type of works, attentions are shifted from output distribution mimicking to feature-level alignment within intermediate network layers.

Here, as illustrated in Figure 6, we accordingly adapted the above two classic distillation paradigms into our development protocol, leading to the logit-based and feature-based baseline frameworks. Table IV compares performance of the proposed PointMCD against these two baselines. First, our experimental results strongly validate that multi-view image-to-point knowledge transfer serves as a robust and stable way of empowering various point cloud learning frameworks. Note that even the most straightforward distillation approach can bring different degrees of task performance enhancement in almost all of our experimental setups. Second, it turns out that the feature-level teacher guidance is more informative than the soft targets (i.e., logits), based on the observation that the feature-based variants uniformly outperform their logit-based counterparts.

2) *Hyperparameter influence:* There are two critical influencing factors within the overall processing pipeline, i.e., 1) the number of views during the rendering process (i.e., K); 2) the weighting scheme between task loss and distillation loss used to formulate the overall training objective (i.e., Eq. (6)).

First, as reported in Table V, we experimented with different number of views. During the generation of multi-view images, we reduced the number of virtual cameras from the complete

TABLE V
ABLATION STUDIES ON THE INFLUENCE OF DIFFERENT NUMBER OF VIEWS DURING MULTI-VIEW IMAGE RENDERING.

Method	Comp-12	Redu-6	Redu-4	Rand-1
<i>MCD-PointNet</i>	91.1	91.0	90.8	90.7
<i>MCD-PointNet++</i>	93.3	93.0	92.9	92.7
<i>MCD-DGCNN</i>	93.7	93.4	93.2	93.3
<i>MCD-CurveNet</i>	94.3	94.1	94.0	94.0

TABLE VI
ABLATION STUDIES ON THE INFLUENCE OF DIFFERENT TASK LOSS WEIGHTS WITHIN THE OVERALL TRAINING OBJECTIVE.

Method	$\omega_t = 0.1$	$\omega_t = 0.01$	$\omega_t = 1.0$
<i>MCD-PointNet</i>	91.1	81.9	90.8
<i>MCD-PointNet++</i>	93.3	83.5	92.9
<i>MCD-DGCNN</i>	93.7	87.1	93.2
<i>MCD-CurveNet</i>	94.3	89.6	93.9

setting of $K = 12$ to $K = \{6, 4, 1\}$, leading to three variants, namely: Redu-6, Redu-4, and Rand-1. More specifically, under the same elevation angle $\theta_k^e = \pi/6$ and observation distance $\mu = 1$, we changed azimuth angles to $\theta_k^{az} \in \{k\pi/3\}_{k=1}^6$ and $\theta_k^{az} \in \{k\pi/2\}_{k=1}^4$ in Redu-6 and Redu-4, respectively. As for Rand-1, throughout the whole training process, we randomly selected one of the original 12 views at each iteration. We can observe from the quantitative comparisons that using fewer rendering views inevitably causes performance degradation, owing to the loss of visual information captured from those missing viewpoints. Particularly, the Rand-1 variant with the random view selection strategy also shows satisfactory performance, which is viewed as a good trade-off between training cost and inference accuracy.

Second, to investigate the influences of different weighting schemes in the overall training objective $\mathcal{L}_{\text{overall}}$, we maintained $\omega_d = 1/K$ while experimenting with different values of ω_t . As reported in Table VI, $\omega_t = 0.1$ turns to be a robust choice that uniformly leads to the optimal performance on all the distilled models. Given a much smaller value of 0.01, the task loss cannot adequately converge, resulting in significant performance degradation. Inversely, enlarging its value to 1.0 weakens the constraining effect of the distillation loss, which also causes relatively sub-optimal results.

3) *Purely point-deduced performance*: In all our preceding experiments, there are two aspects of auxiliary data used in the training process. First, off-the-shelf deep image encoders are pretrained on ImageNet, which means implicit exploitation of large-scale labeled image data. Second, multi-view rendered images are computed from manually created polygon meshes, which contain more informative surface properties than point clouds. Although we emphasize that, while finishing training, the actual inference phase is a purely point-based architecture without requirement of any other auxiliary inputs, we are still interested in the possibility of building our processing pipeline completely upon point clouds. To this end, we correspondingly made the following two aspects of modifications to our original

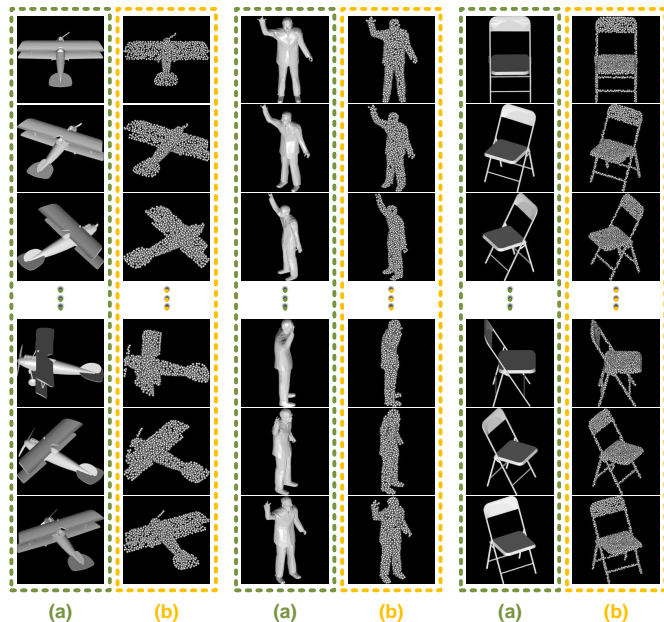


Fig. 7. Visual comparison of (a) mesh-based and (b) point-based multi-view rendering pipelines, where the resulting images are generated from polygon meshes and raw point clouds, respectively.

settings for constructing point-deduced distillation baselines:

- we rendered multi-view images from raw point clouds, rather than polygon meshes. Figure 7 visually compares mesh-based and point-based rendering results; and
- we did not initialize the deep image encoder (i.e., \mathcal{M}_t) with ImageNet-pretrained network parameters during the pretraining process of the whole teacher learning branch.

Table VII quantitatively compares our original framework (i.e., M-Rend) and two point-deduced baselines, in which the first modification produces the P-Rend variant, combining the above two modifications produces the P-Rend-Scratch variant. From these results, we can draw some important conclusions that are worth further emphasizing:

- Comparing M-Rend and P-Rend, we can observe that the overall performance degradation is actually insignificant. In particular, for the distillation of PointNet, P-Rend even slightly outperforms the original M-Rend. This indicates that PointMCD can be implemented in a more practical and flexible way without being excessively demanding on the preceding rendering process.
- Comparing P-Rend and P-Rend-Scratch, we surprisingly notice that pretraining teachers from scratch is comparable to fine-tuning from ImageNet-pretrained parameters. For the distillation of PointNet++ and CurveNet, P-Rend-Scratch is inferior to P-Rend with performance decrease of 0.1% and 0.2%, respectively. In the meanwhile, for the distillation of PointNet and DGCNN, P-Rend-Scratch instead outperforms P-Rend with accuracy boost of 0.3% and 0.1%, respectively. We reason that this phenomenon reveals the great potential of learning from heterogeneous feature representations of cross-modal data.

Overall, we remind that the above ablative experiments and the corresponding purely point-based performances produced

TABLE VII
PERFORMANCE COMPARISON OF OUR ORIGINAL MESH-DEDUCED POINTMCD AND ITS TWO POINT-DEDUCED BASELINES.

Method	M-Rend	P-Rend	P-Rend-Scratch
<i>MCD-PointNet</i>	91.1	91.2	91.5
<i>MCD-PointNet++</i>	93.3	93.0	92.9
<i>MCD-DGCNN</i>	93.7	93.5	93.6
<i>MCD-CurveNet</i>	94.3	94.2	94.0

TABLE VIII
PERFORMANCE OF REAL-SCANNED OBJECT CLASSIFICATION ON SCANOBJECTNN (OBJ-BG) DATASET USING THE DGCNN MODEL, WHERE THE ORIGINAL AND THE DISTILLED MODELS ARE ABBREVIATED AS "ORGN." AND "DIST.", RESPECTIVELY.

DGCNN	ScanObjectNN
<i>Orgn.</i>	82.8
<i>Dist.</i>	89.7

by P-Rend-Scratch are critical and informative, demonstrating that our PointMCD still achieves significant performance boost without any other auxiliary information being introduced in the overall processing pipeline.

4) *Evaluations on real-scanned data:* The preceding experiments use datasets [12], [41], [80] of synthetic shape models. Here, we further evaluated the effectiveness of our approach on the real-scanned object classification dataset of ScanObjectNN [96]. Since mesh models are not available in this dataset, we used point-based rendering for the training of the 2D teacher network. As reported in Table VIII, our PointMCD still brings a significant performance boost from the original 82.8% to the distilled 89.7%.

5) *Explorations of universal teacher knowledge:* Performing cross-modal distillation for different downstream tasks using a single universal 2D teacher network is a highly interesting and promising direction. To verify this, we directly used an ImageNet-pretrained 2D CNN backbone encoder [34], which was also used in our preceding unsupervised learning experiments, to generate the teacher knowledge during the distillation process. As reported in Table IX, it turns out that using such a universal teacher network can still bring varying degrees of performance improvements, although the relative performance gains are relatively smaller compared with distilling with task-specific teacher networks.

V. DISCUSSION

1) *Design Philosophy:* Throughout the whole methodology, our ultimate principle is to make the technical implementation of multi-view cross-modal distillation as concise as possible, while avoiding introducing complicated learning mechanisms in each core component and processing procedure. We believe that a simple but effective overall workflow will better validate the potential of multi-view cross-modal distillation. Therefore, we can reasonably expect that more advanced visual-geometric feature alignment/interaction techniques will further boost the current PointMCD learning framework.

TABLE IX
PERFORMANCE COMPARISON ON MULTIPLE BENCHMARK TASKS, WHERE THE TARGET STUDENT DGCNN MODEL IS DISTILLED BY THE TEACHER KNOWLEDGE EXPORTED FROM AN IMAGENET-PRETRAINED 2D CNN IMAGE ENCODER [34].

ModelNet40		ScanObjectNN		ShapeNetPart	
<i>Orgn.</i>	<i>Dist.</i>	<i>Orgn.</i>	<i>Dist.</i>	<i>Orgn.</i>	<i>Dist.</i>
92.9	93.2	82.8	84.4	85.1	86.0

2) *Experimental Evaluation:* When evaluating the value of our method, we must particularly remind that major attention should be focused on whether PointMCD is able to contribute robust and stable performance gains when applied to different types of deep point encoders on diverse types of task scenarios. In other words, as PointMCD is claimed to serve as a universal plug-in component, one should observe whether the distilled students achieve better task-specific performance, instead of directly making comparisons with some other different learning frameworks. In the preceding experimental sections (e.g., Tables I, II, and III), we listed and compared different paradigms of methods only for demonstrating the significance of the actual performance gains brought by PointMCD.

In addition, during our development of experimental verifications, we noticed that previous works of multi-view learning only benchmarked over global geometry understanding tasks of either classification or retrieval, due to the inconvenience of preparing domain-specific datasets. This work makes efforts to extend additional application scenarios of fine-grained part segmentation and unsupervised learning, forming a much more comprehensive and persuasive evaluation protocol.

3) *In-depth Insights:* Our encouraging experimental results indicate that, in addition to network design, we may as well pay attention to cross-modal data exploitation. In contrast to the availability of large-scale richly-annotated 2D visual data, collecting and labeling 3D geometric shapes can be much more inconvenient and expensive. Under such a context, the proposed image-to-point distillation approach is a promising and highly economical way of boosting 3D shape recognition performance.

4) *Promising Extensions:* In terms of general categorization of knowledge distillation paradigms (as summarized in [68]), our proposed framework can be attributed to the most classic *offline distillation*, which is composed of two separate stages, i.e., 1) pretrain the teacher model before distillation; 2) deploy the pretrained teacher model to guide the subsequent training of the student model. As we look towards the future, we are naturally motivated to investigate more diverse distillation paradigms, including *online distillation* where the teacher model and the student model are simultaneously updated, as well as other data types [97] and richer application scenarios [98].

Moreover, for another different application scenario of multiple modality fusion, we can also investigate multi-modal joint learning frameworks, where both visual and geometric signals are consumed as inputs in the inference phase, by introducing appropriate interaction mechanisms between image encoders and point cloud encoders.

Considering recent progress in visual-language joint modeling [99], [100], it is also highly valuable to explore a unified representation learning framework simultaneously working on language, vision, and geometry signals.

5) *Limitations*: Our current technical implementations may exhibit several typical aspects of potential limitations. First, in our viewpoint configuration scheme, all virtual cameras are positioned outside the underlying surface of 3D objects. For 3D shapes with severe self-occlusion or complex inner structures, the overall distillation process may not effectively impact the feature learning of all input points. Therefore, more adaptive viewpoint specification techniques are required to fully leverage the distillation power. Second, while our approach demonstrates satisfactory effectiveness in object-level task scenarios, adapting the same workflow to scene-level data domains, such as LiDAR point clouds with extremely sparse and non-uniform distributions, poses challenges due to difficulties in performing high-quality image rendering and accurate point-wise visibility checking. Consequently, more flexible feature alignment strategies for 2D image and 3D point cloud data modalities are expected to be further studied.

VI. CONCLUSION


We investigated the possibility of boosting deep 3D point cloud encoders by distilling discriminative cross-modal visual knowledge extracted from multi-view rendered images for 3D shape recognition task scenarios. Technically, we proposed a novel image-to-point distillation framework called PointMCD, serving as a universal plug-in component for generic deep set architectures. To facilitate view-specific heterogeneous feature alignment of paired visual and geometric descriptors, we also customized a simple but effective VAFP mechanism based on visibility checking. Extensive experiments and ablation studies strongly demonstrated the superiority and universality of our method. We believe that our work brings new insights in point cloud community and will motivate more explorations along this promising direction.

REFERENCES

- [1] M. A. Uy and G. H. Lee, "Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition," in *Proc. CVPR*, 2018, pp. 4470–4479.
- [2] B.-S. Hua, Q.-H. Pham, D. T. Nguyen, M.-K. Tran, L.-F. Yu, and S.-K. Yeung, "Scenenn: A scene meshes dataset with annotations," in *Proc. 3DV*, 2016, pp. 92–101.
- [3] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese, "3d semantic parsing of large-scale indoor spaces," in *Proc. CVPR*, 2016, pp. 1534–1543.
- [4] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *Proc. CVPR*, 2017, pp. 5828–5839.
- [5] T. Hackel, N. Savinov, L. Ladicky, J. D. Wegner, K. Schindler, and M. Pollefeys, "Semantic3d.net: A new large-scale point cloud classification benchmark," *ISPRS*, 2017.
- [6] X. Roynard, J.-E. Deschaud, and F. Goulette, "Paris-lille-3d: A large and high-quality ground-truth urban point cloud dataset for automatic segmentation and classification," *IJRR*, vol. 37, no. 6, pp. 545–557, 2018.
- [7] W. Tan, N. Qin, L. Ma, Y. Li, J. Du, G. Cai, K. Yang, and J. Li, "Toronto-3d: A large-scale mobile lidar dataset for semantic segmentation of urban roadways," in *Proc. CVPRW*, 2020, pp. 202–203.
- [8] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Proc. CVPR*, 2012, pp. 3354–3361.
- [9] P. Sun, H. Kretschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proc. CVPR*, 2020, pp. 2446–2454.
- [10] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proc. CVPR*, 2020, pp. 11 621–11 631.
- [11] D. Maturana and S. Scherer, "Voxnet: A 3d convolutional neural network for real-time object recognition," in *Proc. IROS*, 2015, pp. 922–928.
- [12] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3d shapenets: A deep representation for volumetric shapes," in *Proc. CVPR*, 2015, pp. 1912–1920.
- [13] G. Riegler, A. Osman Ulusoy, and A. Geiger, "Octnet: Learning deep 3d representations at high resolutions," in *Proc. CVPR*, 2017, pp. 3577–3586.
- [14] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, and X. Tong, "O-cnn: Octree-based convolutional neural networks for 3d shape analysis," *ACM Trans. Graph.*, vol. 36, no. 4, pp. 1–11, 2017.
- [15] R. Klokov and V. Lempitsky, "Escape from cells: Deep kd-networks for the recognition of 3d point cloud models," in *Proc. ICCV*, 2017, pp. 863–872.
- [16] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller, "Multi-view convolutional neural networks for 3d shape recognition," in *Proc. CVPR*, 2015, pp. 945–953.
- [17] Y. Feng, Z. Zhang, X. Zhao, R. Ji, and Y. Gao, "Gvcnn: Group-view convolutional neural networks for 3d shape recognition," in *Proc. CVPR*, 2018, pp. 264–272.
- [18] T. Yu, J. Meng, and J. Yuan, "Multi-view harmonized bilinear network for 3d object recognition," in *Proc. CVPR*, 2018, pp. 186–194.
- [19] A. Kanezaki, Y. Matsushita, and Y. Nishida, "Rotationnet: Joint object categorization and pose estimation using multiviews from unsupervised viewpoints," in *Proc. CVPR*, 2018, pp. 5010–5019.
- [20] Z. Yang and L. Wang, "Learning relationships for multi-view 3d object recognition," in *Proc. ICCV*, 2019, pp. 7505–7514.
- [21] C. Esteves, Y. Xu, C. Allen-Blanchette, and K. Daniilidis, "Equivariant multi-view networks," in *Proc. ICCV*, 2019, pp. 1568–1577.
- [22] X. Wei, R. Yu, and J. Sun, "View-gcn: View-based graph convolutional network for 3d shape analysis," in *Proc. CVPR*, 2020, pp. 1850–1859.
- [23] A. Hamdi, S. Giancola, and B. Ghanem, "Mvtn: Multi-view transformation network for 3d shape recognition," in *Proc. ICCV*, 2021, pp. 1–11.
- [24] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proc. CVPR*, 2017, pp. 652–660.
- [25] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *Proc. NeurIPS*, 2017, pp. 5105–5114.
- [26] Y. Xu, T. Fan, M. Xu, L. Zeng, and Y. Qiao, "Spidernn: Deep learning on point sets with parameterized convolutional filters," in *Proc. ECCV*, 2018, pp. 87–102.
- [27] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen, "Pointcnn: Convolution on χ -transformed points," in *Proc. NeurIPS*, 2018, pp. 828–838.
- [28] Y. Liu, B. Fan, S. Xiang, and C. Pan, "Relation-shape convolutional neural network for point cloud analysis," in *Proc. CVPR*, 2019, pp. 8895–8904.
- [29] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "Kpconv: Flexible and deformable convolution for point clouds," in *Proc. ICCV*, 2019, pp. 6411–6420.
- [30] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *ACM Trans. Graph.*, vol. 38, no. 5, pp. 1–12, 2019.
- [31] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun, "Point transformer," in *Proc. ICCV*, 2021, pp. 16 259–16 268.
- [32] T. Xiang, C. Zhang, Y. Song, J. Yu, and W. Cai, "Walk in the cloud: Learning curves for point clouds shape analysis," in *Proc. ICCV*, 2021, pp. 915–924.
- [33] X. Ma, C. Qin, H. You, H. Ran, and Y. Fu, "Rethinking network design and local geometry in point cloud: A simple residual MLP framework," in *Proc. ICLR*, 2022.
- [34] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015.
- [35] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.

- [36] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. CVPR*, 2015, pp. 3431–3440.
- [37] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015, pp. 234–241.
- [38] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. NeurIPS*, vol. 25, 2012, pp. 1097–1105.
- [39] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.
- [40] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. CVPR*, 2016, pp. 3213–3223.
- [41] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su *et al.*, "Shapenet: An information-rich 3d model repository," *arXiv preprint arXiv:1512.03012*, 2015.
- [42] J. Wang, L. Zhang, and Z. Yi, "Mixture 2d convolutions for 3d medical image segmentation," *Int. J. Neural Syst.*, pp. 2 250 059–2 250 059, 2022.
- [43] T. Si, F. He, Z. Zhang, and Y. Duan, "Hybrid contrastive learning for unsupervised person re-identification," *IEEE Trans. Multimedia*, 2022.
- [44] P. de Oliveira Rente, C. Brites, J. Ascenso, and F. Pereira, "Graph-based static 3d point clouds geometry coding," *IEEE Trans. Multimedia*, vol. 21, no. 2, pp. 284–299, 2018.
- [45] L. Li, Z. Li, S. Liu, and H. Li, "Efficient projected frame padding for video-based point cloud compression," *IEEE Trans. Multimedia*, vol. 23, pp. 2806–2819, 2020.
- [46] M. Zhang, H. You, P. Kadam, S. Liu, and C.-C. J. Kuo, "Pointhop: An explainable machine learning method for point cloud classification," *IEEE Trans. Multimedia*, vol. 22, no. 7, pp. 1744–1755, 2020.
- [47] H. Liu, Y. Guo, Y. Ma, Y. Lei, and G. Wen, "Semantic context encoding for accurate 3d point cloud segmentation," *IEEE Trans. Multimedia*, vol. 23, pp. 2045–2055, 2020.
- [48] D. Valsesia, G. Fracastoro, and E. Magli, "Learning localized representations of point clouds with graph-convolutional generative adversarial networks," *IEEE Trans. Multimedia*, vol. 23, pp. 402–414, 2020.
- [49] C. Chen, S. Qian, Q. Fang, and C. Xu, "Hapgn: Hierarchical attentive pooling graph network for point cloud segmentation," *IEEE Trans. Multimedia*, vol. 23, pp. 2335–2346, 2020.
- [50] A. F. Guarda, N. M. Rodrigues, and F. Pereira, "Constant size point cloud clustering: A compact, non-overlapping solution," *IEEE Trans. Multimedia*, vol. 23, pp. 77–91, 2020.
- [51] S. Qiu, S. Anwar, and N. Barnes, "Geometric back-projection network for point cloud classification," *IEEE Trans. Multimedia*, vol. 24, pp. 1943–1955, 2021.
- [52] B.-S. Hua, M.-K. Tran, and S.-K. Yeung, "Pointwise convolutional neural networks," in *Proc. CVPR*, 2018, pp. 984–993.
- [53] N. Verma, E. Boyer, and J. Verbeek, "Feastnet: Feature-steered graph convolutions for 3d shape analysis," in *Proc. CVPR*, 2018, pp. 2598–2606.
- [54] M. Xu, R. Ding, H. Zhao, and X. Qi, "Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds," in *Proc. CVPR*, 2021, pp. 3173–3182.
- [55] J. Li, B. M. Chen, and G. H. Lee, "So-net: Self-organizing network for point cloud analysis," in *Proc. CVPR*, 2018, pp. 9397–9406.
- [56] W. Wu, Z. Qi, and L. Fuxin, "Pointconv: Deep convolutional networks on 3d point clouds," in *Proc. CVPR*, 2019, pp. 9621–9630.
- [57] Z. Zhang, B.-S. Hua, and S.-K. Yeung, "Shellnet: Efficient point cloud convolutional neural networks using concentric shells statistics," in *Proc. ICCV*, 2019, pp. 1607–1616.
- [58] J. Yang, Q. Zhang, B. Ni, L. Li, J. Liu, M. Zhou, and Q. Tian, "Modeling point clouds with self-attention and gumbel subset sampling," in *Proc. CVPR*, 2019, pp. 3323–3332.
- [59] E. Nezhadarya, E. Taghavi, R. Razani, B. Liu, and J. Luo, "Adaptive hierarchical down-sampling for point cloud classification," in *Proc. CVPR*, 2020, pp. 12956–12964.
- [60] X. Yan, C. Zheng, Z. Li, S. Wang, and S. Cui, "Pointasnl: Robust point clouds processing using nonlocal neural networks with adaptive sampling," in *Proc. CVPR*, 2020, pp. 5589–5598.
- [61] L. Hui, H. Yang, M. Cheng, J. Xie, and J. Yang, "Pyramid point cloud transformer for large-scale place recognition," in *Proc. ICCV*, 2021, pp. 6098–6107.
- [62] K. Mazur and V. Lempitsky, "Cloud transformers: A universal approach to point cloud processing tasks," in *Proc. ICCV*, 2021, pp. 10715–10724.
- [63] M. Krivokuca, E. Miandji, C. Guillemot, and P. Chou, "Compression of plenoptic point cloud attributes using 6-d point clouds and 6-d transforms," *IEEE Trans. Multimedia*, 2021.
- [64] S. Jiayao, S. Zhou, Y. Cui, and Z. Fang, "Real-time 3d single object tracking with transformer," *IEEE Trans. Multimedia*, 2022.
- [65] X. Yu, L. Tang, Y. Rao, T. Huang, J. Zhou, and J. Lu, "Point-bert: Pre-training 3d point cloud transformers with masked point modeling," in *Proc. CVPR*, 2022, pp. 19313–19322.
- [66] Y. Pang, W. Wang, F. E. Tay, W. Liu, Y. Tian, and L. Yuan, "Masked autoencoders for point cloud self-supervised learning," in *Proc. ECCV*, 2022, pp. 604–621.
- [67] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [68] J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *International Journal of Computer Vision*, vol. 129, no. 6, pp. 1789–1819, 2021.
- [69] M. Jaritz, T.-H. Vu, R. d. Charette, E. Wirbel, and P. Pérez, "xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation," in *Proc. CVPR*, 2020, pp. 12 605–12 614.
- [70] Y.-C. Liu, Y.-K. Huang, H.-Y. Chiang, H.-T. Su, Z.-Y. Liu, C.-T. Chen, C.-Y. Tseng, and W. H. Hsu, "Learning from 2d: Pixel-to-point knowledge transfer for 3d pretraining," *arXiv preprint arXiv:2104.04687*, 2021.
- [71] J. Hou, S. Xie, B. Graham, A. Dai, and M. Nießner, "Pri3d: Can 3d priors help 2d representation learning?" in *Proc. ICCV*, 2021, pp. 5693–5702.
- [72] Z. Liu, X. Qi, and C.-W. Fu, "3d-to-2d distillation for indoor scene parsing," in *Proc. CVPR*, 2021, pp. 4464–4474.
- [73] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "Semantickitti: A dataset for semantic scene understanding of lidar sequences," in *Proc. ICCV*, 2019, pp. 9297–9307.
- [74] J. Xu, R. Zhang, J. Dou, Y. Zhu, J. Sun, and S. Pu, "Rpnnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation," in *Proc. ICCV*, 2021, pp. 16 024–16 033.
- [75] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. CVPR*, 2016, pp. 2818–2826.
- [76] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proc. CVPR*, 2018, pp. 4510–4520.
- [77] S. Katz, A. Tal, and R. Basri, "Direct visibility of point sets," in *ACM SIGGRAPH*, 2007, pp. 24–es.
- [78] A. Brock, T. Lim, J. M. Ritchie, and N. Weston, "Generative and discriminative voxel modeling with convolutional neural networks," *arXiv preprint arXiv:1608.04236*, 2016.
- [79] J.-C. Su, M. Gadelha, R. Wang, and S. Maji, "A deeper look at 3d shape classifiers," in *Proc. ICCV*, 2018, pp. 0–0.
- [80] L. Yi, V. G. Kim, D. Ceylan, I.-C. Shen, M. Yan, H. Su, C. Lu, Q. Huang, A. Sheffer, and L. Guibas, "A scalable active framework for region annotation in 3d shape collections," *ACM Trans. Graph.*, vol. 35, no. 6, pp. 1–12, 2016.
- [81] S. Huang, Y. Xie, S.-C. Zhu, and Y. Zhu, "Spatio-temporal self-supervised representation learning for 3d point clouds," in *Proc. ICCV*, 2021, pp. 6535–6545.
- [82] M. Afham, I. Dissanayake, D. Dissanayake, A. Dharmasiri, K. Thilakarathna, and R. Rodrigo, "Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding," in *Proc. CVPR*, 2022, pp. 9902–9912.
- [83] Y. Chen, J. Liu, B. Ni, H. Wang, J. Yang, N. Liu, T. Li, and Q. Tian, "Shape self-correction for unsupervised point cloud understanding," in *Proc. ICCV*, 2021, pp. 8382–8391.
- [84] B. Eckart, W. Yuan, C. Liu, and J. Kautz, "Self-supervised learning on 3d point clouds by learning discrete generative models," in *Proc. CVPR*, 2021, pp. 8248–8257.
- [85] Y. Yang, C. Feng, Y. Shen, and D. Tian, "Foldingnet: Point cloud auto-encoder via deep grid deformation," in *Proc. CVPR*, 2018, pp. 206–215.
- [86] R. Girdhar, D. F. Fouhey, M. Rodriguez, and A. Gupta, "Learning a predictable and generative vector representation for objects," in *Proc. ECCV*, 2016, pp. 484–499.
- [87] A. Sharma, O. Grau, and M. Fritz, "Vconv-dae: Deep volumetric shape learning without object labels," in *Proc. ECCV*, 2016, pp. 236–250.

- [88] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, "Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling," *Proc. NeurIPS*, vol. 29, 2016.
- [89] Z. Han, M. Shang, Y.-S. Liu, and M. Zwicker, "View inter-prediction gan: Unsupervised representation learning for 3d shapes by learning global shape memories to support local view predictions," in *Proc. AAAI*, vol. 33, no. 01, 2019, pp. 8376–8384.
- [90] P. Achlioptas, O. Diamanti, I. Mitliagkas, and L. Guibas, "Learning representations and generative models for 3d point clouds," in *Proc. ICML*, 2018, pp. 40–49.
- [91] Y. Zhao, T. Birdal, H. Deng, and F. Tombari, "3d point capsule networks," in *Proc. CVPR*, 2019, pp. 1009–1018.
- [92] B. Du, X. Gao, W. Hu, and X. Li, "Self-contrastive learning with hard negative sampling for self-supervised point cloud learning," in *Proc. ACM MM*, 2021, pp. 3133–3142.
- [93] S. Chen, C. Duan, Y. Yang, D. Li, C. Feng, and D. Tian, "Deep unsupervised learning of 3d point clouds via graph topology inference and filtering," *IEEE Trans. Image Process*, vol. 29, pp. 3183–3198, 2019.
- [94] H. Chen, S. Luo, X. Gao, and W. Hu, "Unsupervised learning of geometric sampling invariant representations for 3d point clouds," in *Proc. ICCV*, 2021, pp. 893–903.
- [95] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "Fitnets: Hints for thin deep nets," in *Proc. ICLR*, 2015.
- [96] M. A. Uy, Q.-H. Pham, B.-S. Hua, T. Nguyen, and S.-K. Yeung, "Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data," in *Proc. ICCV*, 2019, pp. 1588–1597.
- [97] M. Carranza-García, F. J. Galán-Sales, and J. M. Luna-Romera, "Object detection using depth completion and camera-lidar fusion for autonomous driving," *Integr Comput Aided Eng*, no. Preprint, pp. 1–18, 2022.
- [98] W. Tang, F. He, and Y. Liu, "Ydtr: infrared and visible image fusion via y-shape dynamic transformer," *IEEE Trans. Multimedia*, 2022.
- [99] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *Proc. ICML*, 2021, pp. 8748–8763.
- [100] Y. Li, F. Liang, L. Zhao, Y. Cui, W. Ouyang, J. Shao, F. Yu, and J. Yan, "Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm," in *Proc. ICLR*, 2022.



Qijian Zhang received the B.S. degree in Electronic Information Science and Technology from Beijing Normal University, Beijing, China, in 2019. Currently, he is a Ph.D. student with the Department of Computer Science, City University of Hong Kong, Hong Kong SAR. His research interests include 3D point cloud processing and geometric deep learning.



Junhui Hou holds a B.Eng. degree in information engineering (Talented Students Program) from the South China University of Technology (SCUT), Guangzhou, China (2009), an M.Eng. degree in signal and information processing from Northwestern Polytechnical University (NPU), Xi'an, China (2012), and a Ph.D. degree from the School of Electrical and Electronic Engineering, Nanyang Technological University (NTU), Singapore (2016). He joined the Department of Computer Science at CityU as an Assistant Professor in 2017 and was promoted

to Associate Professor in 2023. His research interests are multi-dimensional visual computing.

Dr. Hou was the recipient of several prestigious awards, including the Chinese Government Award for Outstanding Students Study Abroad from China Scholarship Council in 2015 and the Early Career Award (3/381) from the Hong Kong Research Grants Council in 2018. He is an elected member of IEEE MSA-TC, VSPC-TC, and MMSP-TC. He is currently serving as an Associate Editor for *IEEE Transactions on Circuits and Systems for Video Technology*, *IEEE Transactions on Image Processing*, *Signal Processing: Image Communication*, and *The Visual Computer*. He also served as the Guest Editor for the *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* and *Journal of Visual Communication and Image Representation* and an Area Chair of multiple conferences.



Yue Qian received the Ph.D. degree in Computer Science at the City University of Hong Kong, Hong Kong in 2022, and the B.S. and M.Phil. degrees in Mathematics from The Chinese University of Hong Kong, in 2014 and 2016, respectively. Her research interests include 3D point cloud and geometric processing.