A Closer Look at Invariances in Self-supervised Pre-training for 3D Vision

Lanxiao $Li^{[0000-0003-3267-2525]}$ and Michael Heizmann^{[0000-0001-9339-2055]}

Institute of Industrial Information Technology, Karlsruhe Institute of Technology, Karlsruhe, Germany {lanxiao.li, michael.heizmann}@kit.edu

Abstract. Self-supervised pre-training for 3D vision has drawn increasing research interest in recent years. In order to learn informative representations, a lot of previous works exploit invariances of 3D features, e.q., perspective-invariance between views of the same scene, modalityinvariance between depth and RGB images, format-invariance between point clouds and voxels. Although they have achieved promising results, previous researches lack a systematic and fair comparison of these invariances. To address this issue, our work, for the first time, introduces a unified framework, under which various pre-training methods can be investigated. We conduct extensive experiments and provide a closer look at the contributions of different invariances in 3D pre-training. Also, we propose a simple but effective method that jointly pre-trains a 3D encoder and a depth map encoder using contrastive learning. Models pre-trained with our method gain significant performance boost in downstream tasks. For instance, a pre-trained VoteNet outperforms previous methods on SUN RGB-D and ScanNet object detection benchmarks with a clear margin.

Keywords: 3D Vision, Self-supervised Learning, Contrastive Learning, Invariances, Point Clouds, Depth Maps

1 Introduction

In order to cope with challenging tasks *e.g.*, object detection, scene understanding, and large-scale semantic segmentation, neural networks for 3D vision are continuously becoming deeper, more complicated, and thus, more data-hungry. In recent years, self-supervised pre-training has shown promising progress in natural language processing and computer vision. By learning powerful representations on non-annotated data, the models gain better performance and convergence in downstream tasks. Self-supervised pre-training is especially appealing in 3D vision because 3D annotation is more costly than the 2D counterpart.

Self-supervised pre-training for 3D vision has already gained some research interests. A lot of previous works use contrastive learning as a pretext task to pre-train models, as it has shown superior performance in other domains [22,29, 56,62]. One classic hypothesis in contrastive learning is that a powerful representation should model view-invariant factors. A common approach to creating



Fig. 1: Invariances in contrastive learning for 3D vision. Without loss of generality, we only consider the local correspondence here. Each column includes two views of the same scene. The exemplary correspondences across two views are illustrated with arrows, which means the two points/pixels have the same coordinate in the 3D space. In self-supervised pre-training, the similarity between corresponding local features is maximized, which forces networks to learn invariance between views. (a) **Perspective-invariance** in two views of the same scene from different view angles. We visualize RGD images instead of point clouds for better clarity. (b) **Modality-invariance** in an aligned image-point cloud pair. The data formats are also different in this case. But we still refer to it as modality-invariance to distinguish it from the format-invariance within a single modality. (c) **Format-invariance** between a depth map and a point cloud converted from it.

different views is data augmentation. Moreover, a 3D scene can be captured from various view angles, with different sensors (*e.g.*, RGB and depth cameras) and represented with different formats (*e.g.*, voxels, point clouds, and depth maps)¹, whereas the major semantic information in the scene is not changed by these factors. Thus, previous works exploit the perspective- [22, 56], modality- [29] and format-invariance [62] of 3D features in self-supervised learning, as shown in Fig. 1. Although these works have shown impressive results, the contribution of the invariances is still under-explored, and a fair and systematic comparison of them hasn't been performed yet.

In this work, we first establish a unified framework for 3D self-supervised learning. Our framework takes into account the local point/pixel-level correspondence as well as the global instance-level correspondence. Also, our frame-

¹ To avoid ambiguity, we use the term *data format* instead of *data representation* in this work.

work unifies contrastive learning with different input data formats and network structures, including Depth-Depth, Point-Point, Depth-Point, Image-Point, and Point-Voxel. By comparing various training strategies exploiting different invariances, we gain non-trivial results. The first insight of this work is that jointly pre-training a 3D encoder and a 2D encoder (Image-Point, Depth-Point) brings better performance than pre-training them separately or jointly pre-training two encoders with the same dimension (*e.g.*, a voxel and a point cloud encoder, which are both three dimensional).

Also, we propose the simple but effective idea to exploit the format-invariances between depth maps and point clouds/voxels. Our intuition is that depth maps are complementary to point clouds and voxels, although they contain almost the same information. The depth map format has the advantage that it's the natural view of the scene and clearly shows the perspective relationship between objects. Also, real-world depth maps usually contain bad pixels, which means the depth values are unmeasurable. In depth maps, the outlines of unmeasurable regions are sharp and clear, e.g., the chair leg in Fig. 1 (c). On the contrary, this information is lost if depth maps are lifted into 3D space. Moreover, thanks to its efficiency, the 2D encoder allows high-resolution depth maps as input, which preserves more fine-grained details in data. However, point or voxel-based networks usually take down-sampled or quantized input to avoid the excessive computational cost and memory usage, which results in inevitable information loss. On the other hand, point clouds and voxels are 3D formats and the corresponding networks can directly capture accurate 3D geometry, whereas depth map-based networks learn spatial relationships indirectly. Also, depth maps alone don't contain the information of camera calibrations. By contrasting the features extracted from two complementary data formats, the two networks learn appreciated properties from each other. This simple idea has less requirements on pre-training data and outperforms previous methods in our experiments.

The contribution of this work is many-fold:

- 1. We introduce a unified self-supervised pre-training framework for *all* major network types and data formats in 3D vision.
- 2. We provide a closer look at invariances in 3D pre-training, *e.g.*, format-, perspective- and modality-invariance.
- 3. We propose a novel approach for 3D self-supervised pre-training, which is based on the format-invariance between depth maps and points/voxels.
- 4. Our method reaches new SOTA results in multiple downstream tasks *e.g.*, object detection on SUN RGB-D dataset [48] and ScanNet [11] dataset.
- 5. The proposed method is also the first self-supervised pre-training approach for depth map-based networks.

2 Related Works

Feature Learning with 3D Data. PointNet [43] is the pioneer in deep learning methods for point clouds. To aggregate local information, PointNet++ [44] down-samples and groups point clouds hierarchically. Recent works [27, 51, 54]

define point convolution on point clouds. Voxel-based methods convert irregular point clouds to regular 3D grids and apply 3D convolution [34, 64] or deep sliding windows [49]. Also, some works [10, 16] introduce sparse CNNs to reduce the computational cost and memory footprint. Some other works use 2D CNNs to extract features from depth maps [26, 57, 58], LiDAR range images [4, 24, 28] or pseudo images [23]. Also, a lot of works use more than one format of 3D data [12, 23, 28, 33, 47]. They and our work share the same motivation to combine the advantages of different data formats. However, our method learns the appreciated property via contrastive learning in the pretext task. In fine-tuning for downstream tasks, only one format is used.

Self-supervised Pre-training in Computer Vision. A lot of pretext tasks for self-supervised learning has been proposed. Some generative approaches recover images under some corruption, e.g., auto-encoders for colorization [60,61] and denoising [50]. Some discriminative approaches generate pseudo-labels for e.g., rotation prediction [15], Jigsaw puzzle solving [36] and objects tracking [52]. Recently, contrastive learning achieved impressive performance in self-supervised learning [3,7,8,19,35,37]. Besides the instance-level discrimination, some works also exploit the local correspondence for better transfer in tasks which need dense features, e.g., object detection and semantic segmentation [5,39,53].

Self-supervised Pre-training for 3D Data. Some works [1,18,45,46] perform self-supervised learning on synthetic data *e.g.*, ShapeNet [55]. However, these approaches don't transfer well to real-world data [56]. PointContrast [56] first uses real-world point cloud data for self-supervised training. It learns perspectiveinvariance by predicting point-wise correspondence between two partially overlapping point clouds. Liu *et al.* [29] pre-train a 3D encoder by using a pre-trained 2D encoder as teacher. On the other hand, Liu *et al.* [32] propose a distillation pipeline to improve 2D encoders by using geometry guidance from 3D encoders. DepthContrast [62] extends the successful MoCo [8, 19] pipeline to 3D domain and exploits the cross-format contrast between point clouds and voxels. Hou *et al.* [22] propose spatial partition to improve the contrastive learning and investigate the data-efficiency and label-efficiency of pre-trained models.

Multi-modal Feature Fusion. The idea of learning from two complementary sources is similar to data fusion. In 3D computer vision, a common practice is to fuse the color and geometry information. A lot of fusing approaches have been proposed, *e.g.*, for object detection [21, 40, 42, 59] and salient object detection [14, 25, 57, 63]. Some other works use self-supervised pre-training to improve the feature fusion [30, 31]. The difference between fusing and contrasting multi-modal features is that fusion enriches features by combining complementary information from different modalities, while contrastive learning maximizes the shared information between modalities.

3 Method

In this work, we intend to research the invariances in 3D self-supervised learning, including perspective-, modality- and format-invariance. For a fair comparison,

it's meaningful to investigate them under a unified framework. In this section, we first briefly revisit some representative works. Then, we introduce a unified framework to which all previous methods fit. Moreover, we introduce several contrastive learning methods under the unified framework. At last, we provide technical details of the framework.

3.1 Unified Framework for 3D Constrastive Learning

In this work, we pay attention to three previous works. (1) PointContrast [56]: it generates two views of the same scene from different perspectives and learns the local correspondence between 3D points using a contrastive loss. (2) Depth-Contrast [62]: following the successful MoCo pipeline [8, 19], it augments two views of the same point cloud to build positive pairs and learn global correspondences by distinguishing the positive samples from a large number of negative samples. Also, it proposes to exploit cross-format contrast between point clouds and voxels. (3) Pixel-to-point [29]: its overall pipeline is similar to PointContrast. However, it learns local correspondences between point clouds and RGB images, in order to benefit from strong pre-trained RGB encoders.

Therefore, a unified framework must support both the local and global correspondence of 3D data and at least two different input types, either from different modalities (*e.g.*, RGB images and point clouds) or with different data formats (*e.g.*, depth maps and point clouds).



Fig. 2: A unified framework for 3D contrastive learning. Here, α and β refer to data types, *e.g.*, point clouds, images, and depth maps.

We show our framework in Fig. 2, which uses a single-view depth map or RGB-D image for simplicity. However, experiments show that the pre-trained weights generalize well on reconstructed multi-view 3D scans. Without loss of

generality, we assume the input of the framework is a depth map in this section. We randomly crop the input to get crop C_1 , which is further randomly augmented and converted into views α_1 and β_1 . Here α and β refer to data formats (*e.g.*, depth maps and point clouds, as visualized in Fig. 2). Then, α_1 and β_1 go through respective encoders and are encoded into pixelwise or pointwise features $\{\mathbf{x}_{\alpha}\}$ and $\{\mathbf{x}_{\beta}\}$. Note that the α and β encoders are usually different networks matching the input formats. But in the case of $\alpha = \beta$, they share the weights, following [56]. As α_1 and β_1 are generated from the same crop C_1 , the dense correspondence between $\{\mathbf{x}_{\alpha}\}$ and $\{\mathbf{x}_{\beta}\}$ can be easily calculated without camera extrinsic. In this work, we follow [56] and use InfoNCE loss to train dense local correspondence, which is further explained in Sec. 3.3.

In order to learn informative representations, our framework also considers the global correspondence between views. Following [62], we perform instance discrimination based on global features \mathbf{q}_{α} and \mathbf{q}_{β} , which are globally pooled and projected from $\{\mathbf{x}_{\alpha}\}$ and $\{\mathbf{x}_{\beta}\}$. To preserve a large number of negative samples for effective contrastive learning, we use memory banks and momentum encoders, following the successful MoCo pipeline [8,19]. However, in supplementary material, we further show that our methods can be generalized to other pipelines, *e.g.* BYOL [17] and SimSiam [9].

Analog to crop C_1 , we randomly crop C_2 from the same depth map, generate α_2 and β_2 and feed them to the momentum encoders. We refer to the globally pooled and projected features from momentum encoders as \mathbf{k}_{α} and \mathbf{k}_{β} , respectively. They are dynamically saved and updated in memory banks during training. Note that unlike [62] our work only contrasts features from different input formats, as we empirically found the gains from additional contrast within the formats are marginal.

3.2 Variants of Strategies

As the overall framework is shown, we now introduce various contrastive learning strategies under this framework. As shown in Fig. 3, we investigate the following variants in this work:

- 1. DPCo (Depth-Point Contrast), our proposed method, which learns *formatinvariance* between depth maps and point clouds.
- 2. DVCo (Depth-Voxel Contrast), our proposed method, which learns *formatinvariance* between depth maps and voxels.
- 3. PVCo (Point-Voxel Contrast), which learns *format-invariance* between point clouds and voxels. It's extended from PointContrast [62].
- 4. PPCo (Point-Point Contrast), which only uses point clouds as input. It serves as a baseline method as it only learns invariance against data augmentation.
- 5. IPCo (Image-Point Contrast), which learns *modality-invarince* between RGB images and point clouds. It's inspired by Pixel-to-point [29].
- 6. PointContrast [56], which learns *perspective-invariance* between view angles. It can be interpreted as a special case of our unified framework, since it generates the crops C_1 and C_2 from two overlapping depth maps from different view angles and only considers the local correspondence.





Fig. 3: Contrastive learning strategies under a unified framework.

In this work, we propose to contrast a 3D format and a 2D format of the same geometric data (*i.e.*, DPCo and DVCo). Although they represent the same 3D scene, the two formats are complementary to some extent. As discussed in Sec. 1, point clouds and voxels directly represent 3D geometry while having inevitable information loss due to sampling and bad pixels. On the contrary, depth maps reserve more information but only represent the 3D scene indirectly. Experiment results show that our methods bring significantly better performance than PPCo and PVCo, which contrast only 3D formats.

3.3 Details

Point Cloud Encoder. We use a U-shaped PointNet++ [44] and follow the network configuration in [41], which consists of 4 down-sampling and 2 up-sampling modules. We use 20K points as input in the pre-training. The number of output points is fixed to 1024.

Voxel Encoder. We use a sparse residual U-Net [10] with 34 convolution layers to encode voxel inputs, following previous works [22, 56]. We use the implementation of sparse convolution in [10]. For geometry-only input, we set all input features to 1. For input with colors, we use normalized RGB values as input features. In pre-training, we quantize inputs with the voxel size of 2.5 cm. The output of the voxel encoder has the same resolution as the input.

Depth Map Encoder. We use the U-shaped 2D CNN in [26] as depth map encoder. The network is a modified ResNet-34 [20] with relative depth convolution [26] and extra up-sampling layers. The input is resized and zero-padded to 352×352 . The output is a feature map down-sampled with factor 8.

Pipeline

Color Image Encoder. Analog to the depth map encoder, we use a ResNet-34 with extra up-sampling layers to encoder the RGB images. We initialize this encoder with the pre-trained weights on ImageNet [13], following the setup in [29]. **Momentum Encoders and Projection Heads.** The momentum encoders have the same structure as the encoders. Their weights are updated via exponential moving average (EMA) from the corresponding encoders instead of back-propagation. We use global max pooling to aggregate the global features. The pooling layer is followed by an MLP consisting of 3 fully connected layers. The intermediate and the output layer have 512 and 128 channels, respectively. The projection heads of momentum encoders are updated via EMA as well.

Loss Functions. Our loss function consists of a local sub-loss $L_{\rm l}$ and a global sub-loss $L_{\rm g}$. The local sub-loss is an InfoNCE loss which optimizes the local dense correspondence:

$$L_{l,\alpha\beta} = -\sum_{i} \log \frac{\exp(\mathbf{x}_{\alpha,i} \cdot \mathbf{x}_{\beta,i}/\tau)}{\exp(\mathbf{x}_{\alpha,i} \cdot \mathbf{x}_{\beta,i}/\tau) + \sum_{j \neq i} \exp(\mathbf{x}_{\alpha,i} \cdot \mathbf{x}_{\beta,j}/\tau)}$$
(1)

With $\mathbf{x}_{\alpha,i} \in {\{\mathbf{x}_{\alpha}\}}$ and $\mathbf{x}_{\beta,j} \in {\{\mathbf{x}_{\beta}\}}$. If the corresponding 3D coordinates of feature vector $\mathbf{x}_{\alpha,i}$ and $\mathbf{x}_{\beta,j}$ are close, they're considered as a positive pair and have i = j. The temperature τ is a hyperparameter and is set to 0.07 in this work. All features are L2-normalized before being fed into the loss function.

The global sub-loss is applied to optimize an instance discrimination task:

$$L_{\mathrm{g},\,\alpha\beta} = -\log \frac{\exp(\mathbf{q}_{\alpha} \cdot \mathbf{k}_{\beta}/\tau)}{\exp(\mathbf{q}_{\alpha} \cdot \mathbf{k}_{\beta}/\tau) + \sum_{n=1}^{N-1} \exp(\mathbf{q}_{\alpha} \cdot \mathbf{k}_{\beta,\,\mathrm{n}}/\tau)} \tag{2}$$

The vector \mathbf{q}_{α} refers to the global feature from the α encoder and \mathbf{k}_{β} the global feature from the β momentum encoder. Since \mathbf{q}_{α} and \mathbf{k}_{β} are generated from the same data sample, they make a positive pair. Features $\mathbf{k}_{\beta,n}$ correspond to other samples and are read from a memory bank with the size N. We use $N = 2^{15}$ in this work. Following previous works, we make our loss symmetric to α and β . The total loss can be formulated as

$$L = L_{\rm l} + L_{\rm g} = 0.25 \cdot \left(L_{\rm l,\,\alpha\beta} + L_{\rm l,\,\beta\alpha} + L_{\rm g,\,\alpha\beta} + L_{\rm g,\,\beta\alpha} \right). \tag{3}$$

Principally, L can be a weighted sum of L_1 and L_g and the weighting factors can be tuned. But we empirically found that the simple arithmetic average already generates good results.

Data Augmentation. We randomly crop C_1 and C_2 . Also, we randomly drop a square area in each crop. We apply random rotation, scaling, and flipping to the point clouds and voxels. We randomly rotate depth maps around principal points and set 20% pixels on the depth map to zero. For RGB images, we apply random color jitter, grayscale, and Gaussian blur.

Dataset. We use ScanNet [11] for the pre-training, following previous works [22, 31, 56, 62]. ScanNet is a large-scale indoor dataset, which contains about 1500 scans reconstructed from 2.5M RGB-D frames. We follow the official train/val split and sample 78K frames (one in every 25 frames) from the train set.

Training. We pre-train the encoders for 120 epochs. We use SGD optimizer with momentum of 0.9 and an initial learning rate of 0.03. We train models on two NVIDIA Tesla V100 GPUs with a total of 64 GB memory and use as large batch size as it fits. The batch size of different strategies varies from 32 to 64. The learning rate is decayed with a cosine schedule. The pre-training takes from two to four days using PyTorch with Distributed Data Parallel.

More technical details can be found in the supplementary material.

4 Experiments and Results

In this section, we first briefly introduce the experimental setups. Then, we compare and analyze different contrastive learning strategies in detail under our unified framework, to clarify the contribution of the invariances. Then, we compare our method (DPCo) with state-of-the-art methods in the point cloud object detection task. At last, we show transfer learning results of our methods on voxels and depth maps. More experimental results can be found in the supplementary material.

4.1 Invariances in 3D Self-supervised Pre-training

In this subsection, we focus on the performance of transfer learning on point cloud-based 3D detection task, since we believe the 3D detection reflects the encoder's capability of capturing both semantic (*i.e.*, objects classification) and geometric (*i.e.*, bounding box regression) information and is thus representative. Also, 3D detection using raw points is well studied in previous works [6, 40–42]. In this work, we fine-tune a VoteNet [41] with a PointNet++ backbone on SUN RGB-D [48] and ScanNet [11] object detection benchmark. The evaluation metrics are the mean Average Precision over the representative classes with the threshold of 25% and 50% 3D-IoU (*i.e.*, AP25 and AP50).

Comparison under the Unified Framework. In this experiment, we compare various contrastive learning strategies under our unified framework. As shown in Tab. 1, all pre-training methods deliver better results than training from scratch in both 3D detection benchmarks. Note that ScanNet benchmark uses point clouds reconstructed from multiple views. Our unified framework, which assumes that the pre-training data are independent single depth maps or RGB-D images, still significantly improves the detection results on this dataset. It implies that the weights pre-trained on single-view data generalize well on multi-view data.

The baseline strategy PPCo utilizes solely the invariance against data augmentation. However, it surpasses PointContrast, which relies on extrinsic parameters, in two out of four metrics. It implies that with a proper design (in our case, the local dense contrast and the MoCo-style instance discrimination), the perspective-invariance is unnecessary in pre-training. A similar observation is also reported in [62]. We hypothesize that in the instance discrimination subproblem, the network has to distinguish inputs from very similar view angles,

as we extract training data from continuous RGB-D videos. This can be interpreted as hard example mining, which forces the network to focus on perspectiverelevant details. Thus, with the help of global correspondence in pre-training, the encoders implicitly learn perspective-relevant information, but not necessarily the invariance in this case.

Method	Invariance	Correspond.	SUN I	RGB-D	Scar	nNet
			AP25	AP50	AP25	AP50
From Scratch	-	-	58.4	33.3	60.0	37.6
PPCo	Augmentation	Local+Global	58.6	34.9	62.6	39.5
PointContrast	Perspective	Local	59.6	34.1	62.8	38.1
PVCo	Format (3D-3D)	Local+Global	59.3	34.9	62.8	39.5
IPCo	Modality	Local+Global	60.2	35.5	63.9	40.9
DPCo (Ours)	Format (2D-3D)	Local+Global	59.8	35.6	64.2	41.5

Table 1: VoteNet fine-tuning performance of self-supervised pre-training strategies with different invariances. We reproduce the results without pre-training using the open-source code of [41], which are slightly better than the original publication.

Moreover, PVCo, which contrasts features from point clouds and voxels, brings slightly better though very similar results as PPCo. It's probably due to the nature of point clouds and voxels as they both represent 3D coordinates directly. Also, PointNet++ is similar to 3D ConvNets, as it conducts convolutionlike local aggregation, uses shared weights in a sliding window manner, and has a hierarchical topology with sub- and up-sampling. Thus, jointly pre-training voxel and point cloud encoders bring limited benefits to the point cloud encoder, compared to pre-training it alone. In this case, incorporating voxel features can be interpreted as a strong data augmentation to point clouds.

However, IPCo and DPCo, which contrast a 2D data format (*e.g.*, color images or depth maps) and a 3D format (*e.g.*, point clouds) achieve significantly better results than PPCo and PVCo, which utilize only 3D formats. It confirms our intuition that the 2D data format is complementary to the 3D format and the correspondence between them can provide strong contrast in self-supervised pre-training. More interestingly, our proposed method DPCo, which uses solely the geometrical information, reaches on-par or better performance as the one using both geometrical and color inputs (IPCo). This is an important advantage in practice, as our method is applicable even if the RGB images are not available or hard to align with depth maps. It also implies that the performance gains of IPCo come probably not from the color information, but from other factors *e.g.*, different resolutions and perspective fields of 2D and 3D networks. Another advantage of DPCo is that it trains faster than PPCo and PVCo, thanks to the efficiency of 2D CNN. Local and Global Correspondence. Our unified framework supports both the local and global correspondence of 3D data in the pre-training. In the following experiments, we investigate the contribution of each type of correspondence separately. As shown in Tab. 2 and Tab. 3, using local and global correspondence alone in the pre-training improves the performance of encoders. Also, comparing with Tab. 1, it's clear that combining them can bring further improvement, which is also observed in 2D pre-training, as discussed in [53]. Moreover, Tab. 2 and Tab. 3 show similar trends as Tab. 1, where IPCo and DPCo show superior performance over others. Interestingly, in Tab. 2 IPCo and DPCo achieve better results than PointContrast even without the global correspondence.

Contrast	SUN RGB-D		ScanNet		
	AP25	AP50	AP25	AP50	
w/o	58.4	33.3	60.0	37.6	
PPCo	58.7	34.8	62.2	38.8	
PVCo	59.1	34.6	62.2	39.0	
PointCo.	59.6	34.1	62.8	38.1	
IPCo	60.1	35.6	62.5	39.4	
DPCo	59.6	35.1	64.2	40.5	

Contrast	SUN RGB-D		$\mathbf{ScanNet}$		
	AP25	AP50	AP25	AP50	
w/o	58.4	33.3	60.0	37.6	
PPCo	59.3	35.1	62.7	39.3	
PVCo	59.0	35.3	62.5	39.6	
IPCo	59.4	34.5	63.3	40.2	
DPCo	59.4	34.9	63.8	41.0	

Table 2: Different choices of local correspondence in pre-training.

Table 3: Different choices of global correspondence in pre-training.

Summary. Our observations concerning with the invariances can be summarized as follows:

- 1. Explicit perspective-invariance in 3D self-supervised learning is unnecessary.
- 2. Format-invariance between 3D formats (*e.g.*, point clouds and voxels) improves the performance but the gains are marginal.
- 3. Format-invariance between depth map and a 3D formats (*e.g.*, depth maps and point clouds) significantly improves the performance, which is slightly better than modality-invariance between point clouds and RGB-images but has fewer requirements on the training data.

4.2 Comparison with SOTA Methods

In the previous subsection, our proposed method DPCo shows the best performance among all variants. In this subsection, we compare it with other SOTA self-supervised pre-training methods. Still, we use the fine-tuning performance in point cloud object detection task as the metric. To obtain a strong supervised baseline, we follow the setup in [26] and generate bounding box annotations for single frames in ScanNet. Then, we pre-train a VoteNet with full supervision. For a fair comparison, the supervised baseline and other self-supervised methods use exactly the same number of frames for pre-training.

In Tab. 4, we compare our methods with PointContrast [56], DepthContrast [62], pixel-to-point [29] and the method of Hou *et al.* [22], which are already discussed in Sec. 2 and Sec. 3.1. As Tab. 4 shows, our method outperforms other self-supervised pipelines in three metrics out of four. It even outperforms the fully supervised baseline on ScanNet AP25 and AP50. Also, our method has on-par performance on SUN RGB-D AP50 und ScanNet AP25 with the upscaled version of DepthContrast [62], which uses a 3-times larger network and is pre-trained with 5-times more data. This result implies that the contribution of format-invariance between point clouds and depth maps is comparable with scaling up the model capacity and the data amount. Also note that besides depth maps (or the equivalence *e.g.*, range images) and camera intrinsic, which are available in almost all 3D datasets, our method doesn't require any extra data, *e.g.*, color images and camera extrinsic, while a lot of SOTA methods do [22, 29, 56].

Pre-training	SUN	RGB-D	Scar	\mathbf{Net}
	AP25	AP50	AP25	AP50
From Scatch	58.4	33.3	60.0	37.6
PointContrast [56]	-	34.8	-	38.0
PointContrast (Ours)	59.5	34.0	61.6	38.2
Hou <i>et al.</i> [22]	-	-	-	39.2
Pixel-to-point [29]	57.2	33.9	59.7	38.9
Pixel-to-point (Ours)	60.1	35.6	62.5	39.4
DepthContrast [62]	60.4	-	61.3	-
DPCo (Ours)	59.8	35.6	64.2	41.5
DepthContrast [†] [62]	61.6	35.5	64.0	42.9
Supervised	62.0	36.3	61.9	38.6

Table 4: Fine-tuning results of VoteNet on SUN RGB-D and ScanNet (scanlevel) object detection benchmark with different pre-training methods. The absent values are not reported in original publications. We report the results of PointContrast and pixel-to-point with our own implementations, as the original papers use a voxel-based backbone instead of a PointNet++. Grayed methods refer to results with extra data or annotations. Specifically, DepthContrast [62] uses a scaled PointNet++ backbone with $3 \times$ more parameters and is pre-trained on both ScanNet and Redwood indoor RGB-D scan dataset [38].

4.3 Data Efficiency

One important goal of pre-training is to transfer the features to very small datasets. To simulate this scenario, we randomly sample a small partition from the downstream datasets (*e.g.*, 5%, 10%) and fine-tune a VoteNet with the backbone pre-trained by DPCo. Experiments with the same percentage share the

same data samples. The validation set is not sampled. As shown in Fig. 4 and Fig. 5, the pre-training brings more improvement, when less fine-tuning data are available. The trend is more obvious on ScanNet, as it contains fewer training samples than SUN RGB-D (1.2K vs. 5K in total). Especially, the DPCo pre-training boosts the AP25 on ScanNet from 13.3% to 36.5% and the AP50 from 2.4% to 14.4%, when only 5% of training data are used.



Fig. 4: Detection results on SUN RGB-D with reduced amount of data.

Fig. 5: Detection results on ScanNet with reduced amount of data.

4.4 Transfer on Depth Map and Voxel Encoders

Till now, we only showed the transfer learning results of point cloud encoders (PointNet++). In this subsection, we investigate the generalization of our methods (DPCo and DVCo) on the depth map and voxel encoders.

Depth Map Encoders. We fine-tune a 2.5D-VoteNet [26], which is a variant of VoteNet with a depth map-based backbone, by initializing its backbone with the pre-trained weights. In order to clarify the contribution of format-invariance, we also pre-train the depth map encoder with solely depth map input. This strategy is similar to PPCo in Fig. 3 and we name it DDCo (Depth-Depth Contrast). Since 2.5D-VoteNet doesn't support multi-view input, we only fine-tune it on SUN RGB-D dataset. One surprising result in Tab. 5 is that the pre-training using DDCo degrades the performance. As a depth map is an indirect representation of 3D coordinates, we hypothesize that DDCo makes the depth map encoder focus on the 2D textures instead of the true 3D geometry, which can be interpreted as cheating in the pre-training. It also implies that the pre-training of depth map encoders is non-trivial and requires a careful design. However, our proposed methods DPCo and DVCo consistently improve the detection results. Since the point cloud and voxel encoders are able to capture 3D geometrical information by their nature, they can provide guidance to the depth map encoder and prevent the depth map encoder from paying too much attention to 2D patterns. Also, combined with the results in Tab. 4, it's worth noticing that DPCo improves

the 3D and 2D encoders at the same time. It proves that the principle of our methods is different from knowledge distillation (KD), which uses a stronger model as a teacher to improve a weaker student model.

$\mathbf{Pre-training} \mathbf{AP25} \mathbf{AP50}$			
From Scratch	60.8	36.9	
DDCo	56.0	31.2	
DVCo	61.0	39.3	
DPCo	61.4	38.8	

Table 5: Fine-tuning results of 2.5D-VoteNet on SUN RGB-D dataset with different contrasting strategies.

Pre-training	S3DIS	ScanNet
From Scratch	66.1	69.6
PVCo DVCo	66.6 67.2	70.3 70.5

Table 6: Fine-tuning results of Sparse 3D ResNet in semantic segmentation tasks. The evaluation metric is mean IoU over classes (mIoU).

Voxel Transfer. To evaluate our methods on voxel-based networks, we use DVCo to pre-train a voxel encoder and fine-tune it for semantic segmentation on ScanNet [11] and S3DIS [2] dataset. We compare the performance with the not pre-trained baseline and PVCo. As shown in Tab. 6, DVCo brings significant improvement to the baseline on both segmentation tasks. The performance is also higher than PVCo, which is consistent with the transfer learning results of point cloud encoders.

5 Conclusion and Future Works

In this work, we establish a unified framework to fairly compare the contribution of perspective-, format- and modality-invariance in 3D self-supervised learning. With the help of our framework, we find out that contrasting a 3D data format (e.g. point clouds and voxels) with a 2D data format (e.g. images and depth maps) is especially beneficial. Moreover, we propose to contrast point clouds or voxels with depth maps instead of RGB images, which brings better performance and has fewer requirements on the training data than previous methods. Experimental results show that our methods improve all types of encoders in 3D vision, including point cloud, voxel, and depth map encoders.

Furthermore, some concerns deserve more research effort. For instance, in our framework, we jointly pre-train two different encoders. Although they both gain performance boost in downstream tasks, it's still unclear, whether each encoder has reached the optimum in the pre-training. In future work, we intend to investigate the optimization and convergence of the joint pre-training.

Acknowledgement. Parts of this work were financed by Baden-Württemberg Stiftung gGmbH within the project KOMO3D.

References

- Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3d point clouds. In: International Conference on Machine Learning (ICML) (2018)
- Armeni, I., Sax, S., Zamir, A.R., Savarese, S.: Joint 2d-3d-semantic data for indoor scene understanding. CoRR abs/1702.01105 (2017)
- Bachman, P., Hjelm, R.D., Buchwalter, W.: Learning representations by maximizing mutual information across views. In: Neural Information Processing Systems (2019)
- Bewley, A., Sun, P., Mensink, T., Anguelov, D., Sminchisescu, C.: Range conditioned dilated convolutions for scale invariant 3d object detection. CoRR abs/2005.09927 (2020), https://arxiv.org/abs/2005.09927
- 5. Chaitanya, K., Erdil, E., Karani, N., Konukoglu, E.: Contrastive learning of global and local features for medical image segmentation with limited annotations. In: Neural Information Processing Systems (2020)
- Chen, J., Lei, B., Song, Q., Ying, H., Chen, D.Z., Wu, J.: A hierarchical graph network for 3d object detection on point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.E.: Big selfsupervised models are strong semi-supervised learners. CoRR abs/2006.10029 (2020), https://arxiv.org/abs/2006.10029
- Chen, X., Fan, H., Girshick, R.B., He, K.: Improved baselines with momentum contrastive learning. CoRR abs/2003.04297 (2020), https://arxiv.org/abs/2003.04297
- Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 15750–15758 (June 2021)
- Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3075–3084 (June 2019)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scan-Net: Richly-annotated 3d reconstructions of indoor scenes. In: Proc. Computer Vision and Pattern Recognition (CVPR), IEEE (2017)
- Dai, A., Nießner, M.: 3DMV: Joint 3d-multi-view prediction for 3d semantic scene segmentation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) Computer Vision – ECCV 2018. Springer International Publishing, Cham (2018)
- Deng, J., Dong, W., Socher, R., Li, L., Kai Li, Li Fei-Fei: ImageNet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255 (2009)
- Fan, D.P., Zhai, Y., Borji, A., Yang, J., Shao, L.: BBS-Net: RGB-D salient object detection with a bifurcated backbone strategy network. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) Computer Vision – ECCV 2020. pp. 275–292. Springer International Publishing, Cham (2020)
- Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. CoRR (2018), http://arxiv.org/abs/1803.07728
- Graham, B., Engelcke, M., van der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)

- 16 L. Li and M. Heizmann
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., kavukcuoglu, k., Munos, R., Valko, M.: Bootstrap your own latent - a new approach to selfsupervised learning. In: Advances in Neural Information Processing Systems. vol. 33, pp. 21271–21284 (2020)
- Hassani, K., Haley, M.: Unsupervised multi-task feature learning on point clouds. In: ICCV (2019)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (June 2016)
- Hou, J., Dai, A., Niessner, M.: 3D-SIS: 3d semantic instance segmentation of RGB-D scans. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4421–4430 (June 2019)
- 22. Hou, J., Graham, B., Nießner, M., Xie, S.: Exploring data-efficient 3d scene understanding with contrastive scene contexts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15587–15597 (2021)
- Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: PointPillars: Fast encoders for object detection from point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
- Li, B., Zhang, T., Xia, T.: Vehicle detection from 3d lidar using fully convolutional network. CoRR (2016), http://arxiv.org/abs/1608.07916
- Li, G., Liu, Z., Ye, L., Wang, Y., Ling, H.: Cross-modal weighting network for RGB-D salient object detection. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) Computer Vision – ECCV 2020. pp. 665–681. Springer International Publishing, Cham (2020)
- Li, L., Heizmann, M.: 2.5D-VoteNet: Depth map based 3d object detection for real-time applications. In: Britisch Machine Vision Conference (BMVC) (2021)
- 27. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: PointCNN: Convolution on x-transformed points. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 31, pp. 820–830. Curran Associates, Inc. (2018)
- Liang, Z., Zhang, M., Zhang, Z., Zhao, X., Pu, S.: RangeRCNN: Towards fast and accurate 3d object detection with range image representation. CoRR (2020), https://arxiv.org/abs/2009.00206
- Liu, Y., Huang, Y., Chiang, H., Su, H., Liu, Z.Y., Chen, C., Tseng, C., Hsu, W.H.: Learning from 2d: Pixel-to-point knowledge transfer for 3d pretraining. CoRR abs/2104.04687 (2021), https://arxiv.org/abs/2104.04687
- Liu, Y., Fan, Q., Zhang, S., Dong, H., Funkhouser, T.A., Yi, L.: Contrastive multimodal fusion with tupleinfonce. CoRR abs/2107.02575 (2021), https://arxiv.org/abs/2107.02575
- Liu, Y., Yi, L., Zhang, S., Fan, Q., Funkhouser, T.A., Dong, H.: P4contrast: Contrastive learning with pairs of point-pixel pairs for RGB-D scene understanding. CoRR abs/2012.13089 (2020), https://arxiv.org/abs/2012.13089
- Liu, Z., Qi, X., Fu, C.W.: 3d-to-2d distillation for indoor scene parsing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4464–4474 (June 2021)

A Closer Look at Invariances in Self-supervised Pre-training for 3D Vision

- Liu, Z., Tang, H., Lin, Y., Han, S.: Point-Voxel CNN for efficient 3d deep learning. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 32, pp. 965–975. Curran Associates, Inc. (2019)
- Maturana, D., Scherer, S.: VoxNet: A 3d convolutional neural network for realtime object recognition. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 922–928 (2015)
- 35. Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
- Noroozi, M., Favaro, P.: Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) Computer Vision – ECCV 2016. pp. 69–84. Springer International Publishing, Cham (2016)
- van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. CoRR abs/1807.03748 (2018), http://arxiv.org/abs/1807.03748
- Park, J., Zhou, Q.Y., Koltun, V.: Colored point cloud registration revisited. In: ICCV (2017)
- Pinheiro, P.O., Almahairi, A., Benmalek, R.Y., Golemo, F., Courville, A.C.: Unsupervised learning of dense visual representations. In: Neural Information Processing Systems (2020)
- Qi, C.R., Chen, X., Litany, O., Guibas, L.J.: ImVoteNet: Boosting 3d object detection in point clouds with image votes. In: The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4404–4413 (June 2020)
- Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep Hough voting for 3d object detection in point clouds. In: The IEEE International Conference on Computer Vision (ICCV). pp. 9277–9286 (October 2019)
- 42. Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum PointNets for 3d object detection from RGB-D data. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 918–927 (June 2018)
- 43. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep learning on point sets for 3d classification and segmentation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 652–660 (July 2017)
- 44. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep hierarchical feature learning on point sets in a metric space. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems 30, pp. 5099–5108. Curran Associates, Inc. (2017)
- 45. Sanghi, A.: Info3d: Representation learning on 3d objects using mutual information maximization and contrastive learning. In: ECCV (2020)
- 46. Sauder, J., Sievers, B.: Self-supervised deep learning on point clouds by reconstructing space. In: NeurIPS (2019)
- 47. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: PV-RCNN: Point-voxel feature set abstraction for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
- Song, S., Lichtenberg, S.P., Xiao, J.: SUN RGB-D: A RGB-D scene understanding benchmark suite. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 567–576 (June 2015)
- Song, S., Xiao, J.: Deep sliding shapes for amodal 3d object detection in RGB-D images. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 808–816 (June 2016)

- 18 L. Li and M. Heizmann
- Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th international conference on Machine learning. pp. 1096–1103. ICML '08, Association for Computing Machinery, New York, NY, USA (Jul 2008). https://doi.org/10.1145/1390156.1390294
- Wang, S., Suo, S., Ma, W.C., Pokrovsky, A., Urtasun, R.: Deep parametric continuous convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
- Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (December 2015)
- Wang, X., Zhang, R., Shen, C., Kong, T., Li, L.: Dense contrastive learning for self-supervised visual pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3024–3033 (June 2021)
- Wu, W., Qi, Z., Fuxin, L.: PointConv: Deep convolutional networks on 3d point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2015)
- Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: PointContrast: Unsupervised pre-training for 3d point cloud understanding. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) Computer Vision ECCV 2020. pp. 574–591. Springer International Publishing, Cham (2020)
- 57. Xing, Y., Wang, J., Chen, X., Zeng, G.: 2.5D convolution for RGB-D semantic segmentation. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 1410–1414 (2019). https://doi.org/10.1109/ICIP.2019.8803757
- Xing, Y., Wang, J., Zeng, G.: Malleable 2.5d convolution: Learning receptive fields along the depth-axis for RGB-D scene parsing. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) Computer Vision – ECCV 2020. pp. 555–571. Springer International Publishing, Cham (2020)
- Xu, D., Anguelov, D., Jain, A.: PointFusion: Deep sensor fusion for 3d bounding box estimation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 244–253 (June 2018)
- Zhang, R., Isola, P., Efros, A.A.: Colorful Image Colorization. In: Computer Vision – ECCV 2016. pp. 649–666. Lecture Notes in Computer Science, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_40
- Zhang, R., Isola, P., Efros, A.A.: Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
- Zhang, Z., Girdhar, R., Joulin, A., Misra, I.: Self-supervised pretraining of 3d features on any point-cloud. CoRR abs/2101.02691 (2021), https://arxiv.org/abs/2101.02691
- 63. Zhao, J.X., Cao, Y., Fan, D.P., Cheng, M.M., Li, X.Y., Zhang, L.: Contrast prior and fluid pyramid integration for RGBD salient object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
- Zhou, Y., Tuzel, O.: VoxelNet: End-to-end learning for point cloud based 3d object detection. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4490–4499 (June 2018)