

Contrastive Predictive Autoencoders for Dynamic Point Cloud Self-Supervised Learning

Xiaoxiao Sheng*, Zhiqiang Shen*, Gang Xiao†

Shanghai Jiao Tong University
{shengxiaoxiao, shenzhiqiang, xiaogang}@sjtu.edu.cn

Abstract

We present a new self-supervised paradigm on point cloud sequence understanding. Inspired by the discriminative and generative self-supervised methods, we design two tasks, namely point cloud sequence based Contrastive Prediction and Reconstruction (CPR), to collaboratively learn more comprehensive spatiotemporal representations. Specifically, dense point cloud segments are first input into an encoder to extract embeddings. All but the last ones are then aggregated by a context-aware autoregressor to make predictions for the last target segment. Towards the goal of modeling multi-granularity structures, local and global contrastive learning are performed between predictions and targets. To further improve the generalization of representations, the predictions are also utilized to reconstruct raw point cloud sequences by a decoder, where point cloud colorization is employed to discriminate against different frames. By combining classic contrast and reconstruction paradigms, it makes the learned representations with both global discrimination and local perception. We conduct experiments on four point cloud sequence benchmarks, and report the results on action recognition and gesture recognition under multiple experimental settings. The performances are comparable with supervised methods and show powerful transferability.

Introduction

The development of depth sensors facilitates the acquisition of dynamic point clouds and makes it widely applied in many scenarios, such as autonomous vehicles and robots. Currently, processing these real-time point clouds for surroundings perception still relies on supervised methods. However, labeling plenty of dynamic point clouds is quite labor-intensive and error-prone. Meanwhile, self-supervised methods in images make great success, significantly alleviating annotation requirement and even outperforming supervised pretraining methods in many tasks (Zbontar et al. 2021; He et al. 2020). Inspired by this, this paper focuses on learning dynamic point cloud representations in a self-supervised manner.

Classical self-supervised methods in images always rely on contrasting the samples with strong data augmentations

*These authors contributed equally.

†Gang Xiao is the corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

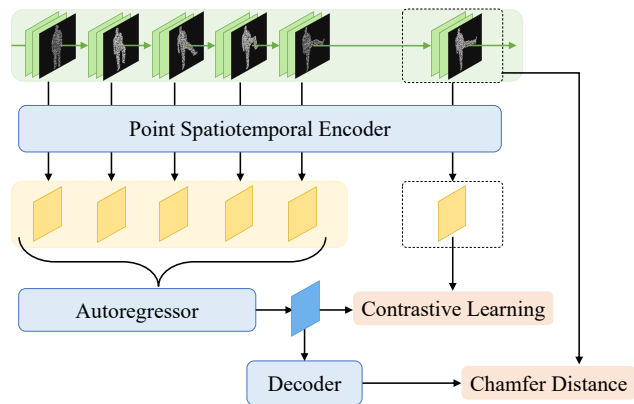


Figure 1: Illustration of our main idea. We combine contrastive predictive coding and reconstruction to establish a unified self-supervised framework for dynamic point clouds.

or reconstructing the randomly masked patches based on visible parts to obtain high-level semantics (Zbontar et al. 2021; Bao et al. 2021). For discriminative methods, rich and diverse sample pairs are usually created. However, it is likely that the semantic consistency of the augmented dynamic point cloud pairs is difficult to be guaranteed. Furthermore, experimental results in (Xie et al. 2020) also demonstrate that real multi-camera views are more effective than hand-crafted views for point clouds. For generative self-supervised methods, the original point coordinates need to be recovered. However, the corresponding coordinates will be added to masked tokens as positional embeddings. This results in information leakage and makes the task less challenging. Based on the above analysis, it is natural to ask a question: *how can we leverage the advantages of discriminative and generative methods for self-supervised learning on dynamic point clouds?*

Considering that point cloud sequences are constantly evolving with geometric dynamics, we can explore temporal-related self-supervised signals lying in the sequence itself, rather than relying on massive augmented data. Contrastive predictive coding is to predict about future based on recent past of the sequences, which has been successfully validated in speeches and videos (Oord, Li, and

Vinyals 2018; Han, Xie, and Zisserman 2020a). Generally, this paradigm aims to use powerful encoder and autoregressor to learn high-level representations with a contrastive pre-text task. In addition, to introduce more task-independent properties, we explore to maximize the mutual information between representations and inputs. Therefore, we integrate point cloud sequence reconstruction into contrastive predictive coding for learning more generalized representations. Our main idea is illustrated in Figure 1.

Specifically, we first sample multiple segments and then send them into an encoder to extract spatiotemporal features. The features of all but the last segments are regarded as tokens to feed into a transformer autoregressor to predict future states. We perform contrastive tasks between the predictions and target segment features. However, due to the unordered and irregular properties of point cloud, the points attached with predicted spatiotemporal features are not aligned to those target ones. We design an interpolation based method to achieve points alignment and further update predicted features correspondingly. Then, local contrastive learning is performed, and we also explore hard negatives to enhance the discriminability of the model. Meanwhile, global contrastive learning is conducted to compare the embeddings of class tokens with sequence-level features of entire current and predicted segments. In addition, raw point cloud sequences of the target segment are reconstructed based on the predictions. Point cloud colorization is applied to the target segment for further discriminating frames. Overall, contrastive prediction task explores more on the discrimination between diverse point cloud sequences, while generative self-supervised task perceives internal structures of point cloud itself. By combining these two self-supervised tasks adapted to dynamic point cloud, it enables the learned representations more comprehensive with both powerful instance discriminability and regional context-awareness.

Multi-granularity representations can be learned by establishing our unified contrastive prediction and reconstruction (CPR) self-supervised framework. We evaluate our method on four dynamic point cloud benchmarks, including MSRAction3D (Li, Zhang, and Liu 2010), NTU-RGBD 60 (Shahroudy et al. 2016), NvGesture (Molchanov et al. 2016), and SHREC'17 (De Smedt et al. 2017). Our method achieves excellent performance comparable with state-of-the-art supervised methods. Ablation studies are performed to investigate the design of self-supervised tasks. The main contributions of this paper are as follows:

- We propose a new contrastive prediction and reconstruction self-supervised framework for dynamic point cloud understanding.
- We conduct local and global contrastive predictions to learn comprehensive representations with both context-awareness and holistic discrimination.
- We reconstruct raw point cloud sequences to promote the generalization of learned representations, and apply point cloud colorization to discriminate different frames.
- We perform extensive experiments on multiple dynamic point cloud benchmarks. We also show detailed ablation studies to analyze how to design self-supervised tasks.

Related Work

Self-supervised Methods for Images

Benefiting from massive data, self-supervised methods in images achieve remarkable performance. This not only alleviates the demands for labeling, but further reveals that more generalized representations are obtained in a self-supervised manner without introducing label ambiguity. Discriminative self-supervised methods compare two views of one sample processed by diverse data augmentations to learn high-level semantics, and utilize plenty of negative samples for contrast to promote instance discriminability (Chen et al. 2020). More effective techniques are also explored to enhance the generalization of pretrained encoder for downstream tasks, such as dynamic queues (He et al. 2020), momentum updating (He et al. 2020), stop-gradients (Chen and He 2021), and clustering (Caron et al. 2020). Generative self-supervised methods aim to learn the generalized representation by mask image modeling (Bao et al. 2021; Xie et al. 2022; He et al. 2022). The masked and unmasked tokens can be simultaneously processed by an encoder, and an extra decoder is necessary to reconstruct raw input or low-level features based on the latent representations (Bao et al. 2021). An asymmetric network structure is also explored to only process the visible patches by the encoder, which significantly improves computation efficiency (He et al. 2022).

Self-supervised Methods for Videos

Self-supervised methods of videos further consider the properties of sequence data itself, such as temporal continuity and redundancy. Similarly, discriminative methods focus on how to build sample pairs from diverse video clips (Pan et al. 2021; Qian et al. 2021). Generative methods randomly mask the spatiotemporal tubes with a higher ratio to prevent information leakage (Tong et al. 2022; Wang et al. 2022b). Besides, many works have been proposed to model motions, such as co-training of two models based RGB and optical flow separately (Han, Xie, and Zisserman 2020b), or building decoupled static and dynamic concepts (Qian et al. 2022). Clearly, contrastive learning based methods emphasize the discrimination between diverse samples, while reconstruction based methods focus on the perception and understanding of samples themselves. They are both limited in learning multi-granularity representations.

Self-supervised Methods for Static Point Clouds

Inspired by the success of self-supervised methods in images, many works focus on learning point cloud representations in a self-supervised manner. PointContrast (Xie et al. 2020) is proposed to contrast real-world multi-view point clouds with various data augmentations. Unfortunately, multi-camera viewpoints are not available in some scenarios. As a BERT-style self-supervised framework, PointBERT (Yu et al. 2022) predicts latent representations learned by an offline tokenizer, which causes two-stage pretraining. In addition, directly recovering raw point cloud easily leads to information leakage due to the positional encoding of masked tokens. To alleviate the above problem, PointMAE (Pang et al. 2022) is proposed to only input unmasked

tokens into encoder and add the masked tokens to the decoder for reconstructing. Alternatively, MaskPoint (Liu, Cai, and Lee 2022) is proposed to randomly select visible points into encoder and train a decoder to discriminate between masked points and noises points. It combines the discriminative method and the masking task to achieve excellent performance.

Modeling Dynamic Point Clouds

Currently, modeling dynamic point cloud is still dominated by traditional supervised methods (Zhong et al. 2022; Fan et al. 2021b; Wei et al. 2022; You and Jiang 2019; Wen et al. 2022; Fan, Yang, and Kankanhalli 2022). MeteorNet (Liu, Yan, and Bohg 2019) gradually learns aggregated features for each point by constructing point-wise spatiotemporal neighborhoods. 3DV (Wang et al. 2020) encodes motion information utilizing regular voxels and then abstracts these dynamic voxels into point sets to model spatiotemporal features. PSTNet (Fan et al. 2021a) hierarchically extracts features of raw point cloud sequences with spatial and temporal decoupled convolutions. P4Transformer (Fan, Yang, and Kankanhalli 2021) first utilizes the point tube convolution to build tokens and further aggregates these tokens by a transformer encoder. Wang *et al.* (Wang et al. 2021) proposes an order prediction self-supervised task on shuffled point cloud clips to learn dynamic point cloud features. However, this method only mines temporal information, while ignoring spatiotemporal modeling. Therefore, the spatiotemporal discriminability and local context awareness of representations learned by this manner are limited.

Unlike the previous methods, we propose a new self-supervised paradigm for modeling point cloud videos. We combine the advantages of contrastive learning and generative methods to design contrastive prediction and reconstruction tasks for dynamic point clouds, jointly facilitating the generalization performance for downstream tasks.

Contrastive Prediction and Reconstruction

We present the overall framework of our method in Figure 2, which includes three main modules. *Point Spatiotemporal Encoder* is first utilized to aggregate dense point cloud segments. Then, we input the aggregated embeddings into a *Transformer Autoregressor* to make predictions for the last target segment. An MLP head is introduced to further project the representations into latent space for contrast. We perform local and global contrastive learning between predictions and targets to progressively capture multi-granularity features. In addition, raw point cloud sequences are reconstructed by a *Decoder* based on the predictions. In this section, we first briefly introduce the encoder and autoregressor, then present contrastive prediction and reconstruction self-supervised tasks in detail.

Point Spatiotemporal Encoder

A point cloud sequence is denoted as $\mathbf{P} \in \mathbb{R}^{T \times N \times 3}$, where T is the sequence length and N is point number in each frame. PSTNet (Fan et al. 2021a) is a typical spatial and temporal decoupled feature extractor for dynamic point clouds.

Specifically, the farthest point sampling is first utilized to sample N' anchor points within each frame. Then, anchors are mapped into adjacent frames to construct spatiotemporal point tubes. Spatial convolution is built to extract local structures within the neighborhoods grouped by k -nearest neighbor algorithm. 1D convolution is utilized to capture temporal features on these tubes. These aggregated spatiotemporal features are denoted as $\{\mathbf{x}_i\}_{i=1}^m$, where m is the number of aggregated super-points. In this paper, we focus more on how to design self-supervised tasks for learning point spatiotemporal representations. Therefore, we directly adopt the backbone of PSTNet as our encoder.

Transformer Autoregressor

The autoregressor is to predict future representations based on spatiotemporal features extracted by the encoder. In classic contrastive prediction based methods (Oord, Li, and Vinyals 2018; Han, Xie, and Zisserman 2020a), LSTMs are always used as autoregressors to predict future states of sequence data. However, when dealing with long sequences, LSTMs are prone to catastrophic forgetting, causing inferior predictions. Considering that 3D action recognition strongly relies on long-term temporal information, we introduce the powerful transformer as the autoregressor. We adopt the standard transformer which consists of multi-head self-attentions and feed-forward networks. We treat the above $\{\mathbf{x}_i\}_{i=1}^m$ as tokens. Positional embeddings $\{\mathbf{e}_i\}_{i=1}^m$ are obtained by projecting the spatiotemporal coordinates (x, y, z, t) of super-points with a linear function. $\mathbf{E}[s]$ is the class token and \mathbf{e}_0 is its positional embedding obtained by projecting (x_0, y_0, z_0, t_0) , where (x_0, y_0, z_0) is mean value of target segment points and t_0 is the timestamp of target segment. Finally, the token sequence $[\mathbf{E}[s] + \mathbf{e}_0, \mathbf{x}_1 + \mathbf{e}_1, \dots, \mathbf{x}_m + \mathbf{e}_m]$ is sent into the transformer to model global spatiotemporal relations and make predictions.

Contrastive Prediction Task

We divide the whole input sequence $\mathbf{P} \in \mathbb{R}^{T \times N \times 3}$ into S segments equally. Each segment is $\mathbf{P}' \in \mathbb{R}^{M \times N \times 3}$, and $T = S \times M$. We utilize the point spatiotemporal encoder to obtain target embeddings $\mathbf{Z} \in \mathbb{R}^{l \times r \times c}$ of the S -th target segment, where l is the frame number after aggregation, r is the super-point number, and c is embedding channel. The former $S-1$ segments are input into the encoder and autoregressor to predict target embeddings. Because of temporal adjacency with target segment, the updated embeddings of the $(S-1)$ th segment are taken as predictions, denoted as $\mathbf{Q} \in \mathbb{R}^{l \times r \times c}$. However, due to the unordered and irregular properties of point clouds, it can not directly align predictions and target embeddings. To alleviate this dilemma, we interpolate the predicted features using k -nearest neighbor algorithm by centering at target super-points, where $k = 3$. We denote the interpolated predictions as $\hat{\mathbf{Q}} \in \mathbb{R}^{l \times r \times c}$. The aligned embeddings are regarded as positive pairs $(\mathbf{z}_i, \hat{\mathbf{q}}_+)$, and the remainings are negative samples. Moreover, the embeddings belonging to former $S-2$ segments are explored as hard negatives to improve discriminability of the model. Then, local contrastive learning is conducted to obtain fine-grained features. We utilize local Info Noise Contrastive Es-

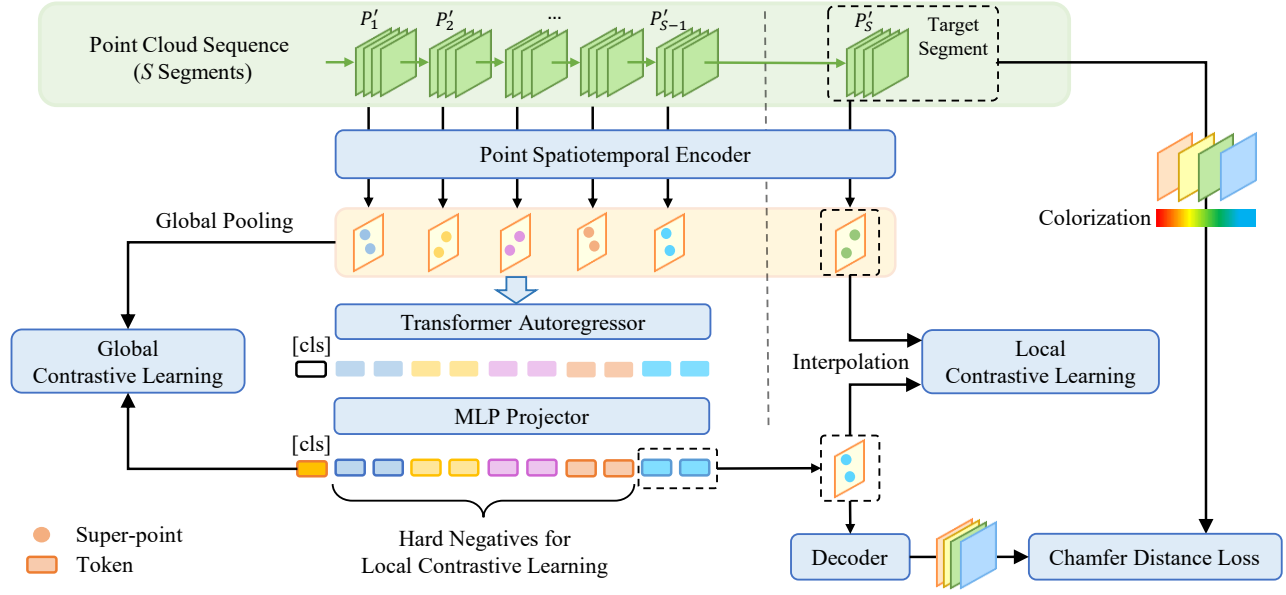


Figure 2: The framework of our approach. It contains three main components, *i.e.*, point spatiotemporal encoder, transformer autoregressor, and decoder. By organically combining these components, we construct contrastive prediction and reconstruction self-supervised tasks.

timization (InfoNCE) loss as follows:

$$\mathcal{L}_l = - \sum_{z_i \in \mathcal{Z}} \log \frac{\exp(z_i^T \hat{q}_+ / \tau)}{\exp(z_i^T \hat{q}_+ / \tau) + \sum_{q_j \in \Psi} \exp(z_i^T q_j / \tau)}, \quad (1)$$

where Ψ is a set that contains the negatives, and τ is temperature hyper-parameter.

To explore holistic semantics, we obtain the global representations of whole input sequence, denoted as $\mathbf{H} \in \mathbb{R}^{B \times c}$, by performing max-pooling on the embeddings of S segments passed through the encoder. Then, we perform global contrastive learning between the representation of class token and \mathbf{H} . The corresponding sample pairs are positives, otherwise negatives. The global InfoNCE loss is as follows:

$$\mathcal{L}_g = - \sum_{h_i \in \mathbf{H}} \log \frac{\exp(h_i^T \hat{g}_+ / \tau)}{\exp(h_i^T \hat{g}_+ / \tau) + \sum_{g_j \in \Theta} \exp(h_i^T g_j / \tau)}, \quad (2)$$

where h_i represents the embedding of i -th input sequence, \hat{g}_+ is the embedding of its class token, and Θ contains all negatives.

By combining local and global contrastive learning, our method effectively captures multi-granularity features. Moreover, we explore more meaningful hard negatives and present their effectiveness in ablation studies.

Point Cloud Sequences Reconstruction Task

Theoretical analysis in (Wang et al. 2022a) points out that reconstructing raw inputs maximizes mutual information between representations and inputs to promote generalizability. Inspired by this, we design sequences reconstruction using the predictions. Rather than recovering all points of target segment, we perform spatial downsampling by the farthest point sampling. This avoids paying more attention

to low-level details and reducing the computational burden. Specifically, we colorize each frame of downsampled target segment with specific RGB values to distinguish different frames. With the increase of the temporal index, the corresponding color changes from red to green to blue (Yang et al. 2021). We denote the colorized target points as $\hat{\mathbf{P}}'_t \in \mathbb{R}^{M \times N' \times 6}$, where N' is the number of downsampled points in each frame and 6 means (x, y, z, r, g, b) .

We utilize average pooling on predictions \mathbf{Q} to obtain global semantics $\mathbf{q}_g \in \mathbb{R}^c$, and then duplicates it to $\mathbf{Q}_g \in \mathbb{R}^{M \times N' \times c}$. Moreover, we add 1D cosine positional encodings to \mathbf{Q}_g to provide temporal clues, and the points in each frame share the same cosine encoding. Finally, we put these updated embeddings into the decoder for spatiotemporal reconstruction. We adopt FoldingNet (Yang et al. 2018) as the decoder and exploit chamfer distance loss to optimize this self-supervised task as follows:

$$d(\mathbf{R}, \hat{\mathbf{P}}'_t) = \frac{1}{|\mathbf{R}|} \sum_{\hat{x} \in \mathbf{R}} \min_{p \in \hat{\mathbf{P}}'_t} \|\hat{x} - p\|_2^2 + \frac{1}{|\hat{\mathbf{P}}'_t|} \sum_{\hat{x} \in \mathbf{R}} \min_{p \in \hat{\mathbf{P}}'_t} \|p - \hat{x}\|_2^2, \quad (3)$$

where $\mathbf{R} \in \mathbb{R}^{M \times N' \times 6}$ is reconstructed sequence.

Overall, the total loss with a regularized parameter λ of our self-supervised framework consists of three parts:

$$\mathcal{L}_{total} = \mathcal{L}_l + \mathcal{L}_g + \lambda d(\mathbf{R}, \hat{\mathbf{P}}'_t). \quad (4)$$

Experiments

Datasets

We perform 3D action recognition on MSRAction3D and NTU-RGBD 60, and gesture recognition on NvGesture and SHREC'17 datasets.

Methods	Input	Accuracy				
Vieira <i>et al.</i> (Vieira et al. 2012)	depth map	78.20	81.14	86.53	88.21	88.50
Kläseret <i>et al.</i> (Klaser, Marszałek, and Schmid 2008)	depth map	81.43	83.50	87.88	89.90	91.20
Actionlet (Wang et al. 2012)	skeleton	88.21	83.17	87.54	89.56	90.94
Frames:		4	8	12	16	24
MeteorNet (Liu, Yan, and Bohg 2019)	point	78.11	81.14	86.53	88.21	88.50
PSTNet (Fan et al. 2021a)	point	81.14	83.50	87.88	89.90	91.20
P4Transformer (Fan, Yang, and Kankanhalli 2021)	point	80.13	83.17	87.54	89.56	90.94
PSTNet++ (Fan et al. 2021b)	point	81.53	83.50	88.15	90.24	92.68
PST-Transformer (Fan, Yang, and Kankanhalli 2022)	point	81.14	83.97	88.15	91.98	93.73
PST ² (Wei et al. 2022)	point	81.14	86.53	88.55	89.22	-
Kinet (Zhong et al. 2022)	point	79.80	83.84	88.53	91.92	93.27
PPTr (Wen et al. 2022)	point	80.97	84.02	89.89	90.31	92.33
4D MinkNet + ROP Pretraining (Wang et al. 2021)	point	-	86.31	-	-	-
MeteorNet + ROP Pretraining (Wang et al. 2021)	point	-	85.40	-	-	-
CPR (Ours)	point	82.50	86.53	91.00	92.15	93.03

Table 1: Action recognition accuracy (%) on the MSRAction3D dataset.

MSRAction3D (Li, Zhang, and Liu 2010) includes 567 videos collected by Kinect, with a total of 23K frames and 20 categories performed by 10 subjects. We use the same training and test splits as (Liu, Yan, and Bohg 2019).

NTU-RGBD 60 (Shahroudy et al. 2016) is collected by three cameras with different angles, containing 60 categories and 56880 videos performed by 40 subjects. Cross-subject and cross-view evaluations are adopted.

NvGesture (Molchanov et al. 2016) contains 1532 gesture videos focusing on touchless driver controlling, with a total of 25 classes. We follow the previous work to split this dataset, where 1050 videos are used for training and 482 videos are for test (Min et al. 2020).

SHREC’17 (De Smedt et al. 2017) collects 2800 videos performed by 28 subjects in two ways, *i.e.*, using one finger or the whole hand. These short videos generally contain dozens of frames and involve both coarse and fine gestures. We adopt the same splits of training and test data as previous work (Min et al. 2020).

Experimental Setting

We perform pretraining on NTU-RGBD 60. Specifically, we consecutively sample 24 frames with stride 2, and each frame contains 1024 points. This sequence is then divided into 6 segments. Following (Fan et al. 2021a), the spatial search radius is 0.1, neighbors for the ball query is 9, and random scaling is adopted as data augmentation for our encoder. Three transformer layers are utilized as our autoregressor. We pretrain 200 epochs and set the batchsize to 88. We use Adam optimizer and cosine annealing scheduler with the initial learning rate 0.0008.

Without special instructions, we adopt the pretrained point spatiotemporal encoder for downstream tasks, and add two linear layers with BN and ReLu for finetuning or one linear layer for linear evaluation. We utilize SGD optimizer with momentum 0.9 and cosine scheduler with warmup 10

epochs. The 16 batchsize corresponds to the 0.01 learning rate, and we follow the scale up rule.

Comparison with State-of-the-art

We perform extensive experiments and compare the performances with state-of-the-art supervised methods and other self-supervised methods.

Transfer to MSRAction3D. The finetune results are presented in Table 1, compared with skeleton-based, depth-based, and point-based methods. We follow the settings of previous work (Liu, Yan, and Bohg 2019) and test the performance under variable lengths. Although the skeleton-based method Actionlet simultaneously feeds all frames, it is still exceeded by our CPR with 12 frames. This shows that point cloud sequences contain richer spatiotemporal information. Moreover, since 3D action recognition relies on temporal information, CPR obtains higher accuracy when inputting longer sequences. Compared with other point-based supervised methods which train from scratch, CPR achieves improvements under diverse inputs, except that it is slightly lower than PST-Transformer on 24 frames. This indicates that high-level semantics can be obtained by pretrained encoder and as a good initialization for transfer learning. Compared with the self-supervised method that performs Recurrent Order Prediction task (ROP) (Choy, Gwak, and Savarese 2019), CPR achieves higher accuracy. ROP utilizes RGB values for pretraining and only focuses on clip-level temporal information, while our method utilizes local/global contrastive learning and reconstruction to explore richer spatiotemporal information just with points.

Transfer to NvGesture and SHREC’17. We perform self-supervised pretraining on human action datasets. To further verify the generalization of our method, we also transfer the pretrained encoder to two gesture datasets. During fine-tuning, the spatial search radius is 0.1, the neighbors of ball query is 9, and the frame number is 32. CPR is compared

Methods	Input	NvG	S17
Human(Molchanov et al. 2016)	RGB	88.4	-
DG-STA (Chen et al. 2019)	skeleton	-	90.7
FlickerNet (Min et al. 2019)	point	86.3	-
PLSTM-base (Min et al. 2020)	point	85.9	87.6
PLSTM-early (Min et al. 2020)	point	87.9	93.5
PLSTM-PSS (Min et al. 2020)	point	87.3	93.1
PLSTM-middle(Min et al. 2020)	point	86.9	94.7
PLSTM-late (Min et al. 2020)	point	87.5	93.5
Kinet(Zhong et al. 2022)	point	89.1	95.2
PSTNet(Fan et al. 2021a)	point	88.2	92.0
CPR (Ours)	point	88.7	93.1

Table 2: Gesture recognition accuracy (%) on NvGesture (NvG) and SHREC’17 (S17).

with skeleton-based, RGB-based and other point-based supervised methods. The results are presented in Table 2. By finetuning with our self-supervised encoder, CPR facilitates the baseline PSTNet to produce comparable results. This shows that the pretrained encoder has strong generalization ability, and exhibits powerful spatiotemporal modeling ability under diverse domains.

Finetune on NTU-RGBD 60. We keep the encoder and autoregressor for finetuning. As shown in Table 3, we compare CPR with skeleton-based (Liu et al. 2017; Si et al. 2019; Li et al. 2019; Zhang et al. 2019; Shi et al. 2019b,a), depth-based (Xiao et al. 2019), and point-based supervised methods. Particularly, Kinet builds a hierarchical motion branch, and 3DV explicitly encodes voxel-based motion cues. Instead, without introducing hand-crafted motion features and complicated two-stream design, CPR achieves competitive accuracy under two evaluations. This demonstrates the superiority of our self-supervised approach. By exploiting pre-training to fully mine motion information in the raw data, it can help the network acquire rich semantics without using additional dynamical model. Our performance under cross-view evaluation is consistent with PSTNet++, and we will explore more advanced encoders in the future.

Linear and Semi-Supervised Learning. To test whether the pretrained encoder has learned high-level semantics, we evaluate it on MSRAction3D and NTU-RGBD 60 under linear evaluation and limited training data. Semi-supervised training data consists of randomly selected 30% of the training samples from each class. We conduct self-supervised pretraining on MSRAction3D and NTU-RGBD 60, respectively. The results are shown in Table 4. It is observed that the linear accuracy of *MSR to MSR*, *NTU to NTU*, and *NTU to MSR* is 85.7%, 70.0%, and 79.1%. This shows that our self-supervised pretraining captures rich semantics beneficial for 3D action recognition. After semi-supervision with 30% data, all the performances get promoted compared with linear evaluations. Since the 30% MSRAction3D data is too small, the improvement of finetuning on MSR is limited. In addition, the semi-supervised performance of *NTU to MSR* is higher than that of *MSR to MSR*. This indicates that self-

Methods	Input	NTU-RGBD 60	
		Subject	View
GCA-LSTM	skeleton	74.4	82.8
AGC-LSTM	skeleton	89.2	95.0
AS-GCN	skeleton	86.8	94.2
VA-fusion	skeleton	89.4	95.0
2s-AGCN	skeleton	88.5	95.1
DGNN	skeleton	89.9	96.1
MVDI	depth	84.6	87.3
PointNet++	point	80.1	85.1
3DV (motion)	voxel	84.5	95.4
3DV-PointNet++	voxel + point	88.8	96.3
PSTNet	point	90.5	96.5
P4Transformer	point	90.2	96.4
PSTNet++	point	91.4	96.7
PST-Transformer	point	91.0	96.4
Kinet	point	92.3	96.4
CPR (Ours)	point	91.0	96.7

Table 3: Action recognition accuracy (%) on the NTU-RGBD 60 dataset.

SSL-Data	Linear		Semi-supervision	
	Data	Accuracy	Data	Accuracy
MSR	MSR	85.7	30% MSR	86.4
NTU	NTU	70.0	30% NTU	83.3
NTU	MSR	79.1	30% MSR	87.2

Table 4: Recognition accuracy (%) of linear evaluation and semi-supervised learning on MSRAction3D (MSR) and NTU-RGBD 60 (NTU). SSL-Data is self-supervised data.

supervised pretraining with larger dataset can learn more general semantics while still needs limited data for transfer adaptation.

Ablation Studies

In this section, we show extensive ablation studies to explore that the length of whole input sequence utilized for our self-supervised framework, the design of point cloud sequences reconstruction, the effectiveness of hard negatives, and the influences of diverse self-supervised tasks. These experiments are all pretrained on MSRAction3D and then finetuned on it.

How long of the point cloud sequence? We perform self-supervised pretraining with variable input lengths and then finetune with diverse frames. The results are shown in Table 5. Notably, when sequence length of self-supervised pretraining is longer, it is more beneficial for finetuning under multiple frames. We finally choose 24 frames for self-supervised pretraining to cover as much dynamic information as possible.

How to design reconstruction tasks? For reconstructing branch, we try various experimental settings and the results

SS-f \ Fin-f					
	4	8	12	16	24
12	78.45	82.83	88.82	90.91	91.99
16	79.12	83.84	88.55	91.03	92.33
20	80.13	85.70	89.56	91.92	93.03
24	82.06	86.20	90.42	92.07	93.38

Table 5: Finetune accuracy (%) at different frames. Fin-f and SS-f are the frame number for finetune and self-supervised pretraining.

Reconstruction Target	Points of each frame	Acc.
One frame	1024	91.32
One segment	256	92.68
One segment + Color.	256	93.38
One segment + Color.	512	91.99
One segment + Color.	1024	92.36

Table 6: Finetune accuracy (%) with different reconstruction targets. Color. means colorization.

are presented in Table 6. The accuracy of reconstructing one segment is higher than that of reconstructing one frame under the same number of points. This may be due to the fact that spatiotemporal prediction is more challenging for self-supervised tasks. When applying point cloud sequence colorization, we achieve higher accuracy. Colorization is beneficial to distinguish diverse frames by assigning different timestamps for the target segment. We also try to reconstruct more points in each frame. However, this does not lead to improvements. It is possible that excessive raw points provide more low-level details and even noise, which does not help to promote generalization and discrimination. Finally, we choose to reconstruct one colorized segment with 256 points in each frame.

Why utilize hard negative samples? The tokens of the former $S-2$ segments are temporally adjacent to the predictions and targets, and therefore they contain similar spatiotemporal semantics. These tokens are mined as hard negatives to enhance local perception and discriminability. The results in Table 7 show that the accuracy of local contrastive prediction without hard negatives is 91.34%, and the accuracy increases to 92.23% after adding them. This indicates that mining hard negative samples is crucial for the design of self-supervised tasks.

How the effects of each task? We respectively evaluate the effectiveness of local contrastive prediction, global contrastive prediction, and point cloud sequence reconstruction tasks. The results are shown in Table 7. The hard negatives can increase the performance of local contrastive prediction by about 1%. Clearly, by introducing the reconstruction task, the finetune performance has increased by 0.83%, which indicates reconstructing raw inputs helps to learn semantics information suitable for 3D action recognition. Compared with only utilizing local contrastive prediction, the introduction of global contrastive prediction increases the accuracy

Method	Acc.
Local contrastive prediction	91.34
Global contrastive prediction	90.89
Sequence reconstruction	87.10
Local contrastive prediction with hard negatives	92.23
Local and global contrastive prediction	92.55
Contrastive prediction and reconstruction	93.38

Table 7: Ablation studies on architecture design.




Figure 3: Visualization of the t-SNE features. We visualize the feature distributions of our self-supervised encoder (a) after pretraining, (b) after finetuning on NTU-RGBD 60, and (c) after finetuning on MSRAction3D.

by 0.32%. This shows that global discrimination and local perception are all essential for capturing multi-granularity representations. Overall, the most beneficial paradigm is to combine contrastive prediction and point cloud sequence reconstruction.

Visualization

In Figure 3, we visualize the feature distributions of t-SNE after pretraining (a) and after finetuning (b)(c). It can be seen that after pretraining, there are approximate boundaries between 60 categories. This illustrates that self-supervised pretraining can learn certain high-level semantics. From Figure 3(c), it is observed that each cluster has clear outlines, indicating that the representations learned by our self-supervised framework have strong generalization and transferability. Moreover, the proposed framework can guide the encoder to obtain domain-agnostic and general knowledge.

Conclusion

In this work, we propose a unified contrastive prediction and reconstruction self-supervised framework for dynamic point cloud understanding. By integrating discriminative and generative self-supervised tasks, it makes the learned representations with both powerful instance discrimination and local perception. Extensive experiments under linear evaluation, semi-supervised learning, and transfer learning are performed on multiple benchmarks to demonstrate the superiority of our self-supervised framework. In the future, we will explore more dynamic point cloud related downstream tasks.

Acknowledgments

This paper is sponsored by National Natural Science Foundation of China (No.61673270; No.61973212).

References

- Bao, H.; Dong, L.; Piao, S.; and Wei, F. 2021. Beit: bert pre-training of image transformers. In *International Conference on Learning Representations (ICLR)*.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems (NeurIPS)*, 9912–9924.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning (ICML)*, 1597–1607.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 15750–15758.
- Chen, Y.; Zhao, L.; Peng, X.; Yuan, J.; and Metaxas, D. N. 2019. Construct dynamic graphs for hand gesture recognition via spatial-temporal attention. *arXiv preprint arXiv:1907.08871*.
- Choy, C.; Gwak, J.; and Savarese, S. 2019. 4D spatio-temporal convnets: minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3075–3084.
- De Smedt, Q.; Wannous, H.; Vandeborre, J.-P.; Guerry, J.; Le Saux, B.; and Filliat, D. 2017. Shrec’17 track: 3d hand gesture recognition using a depth and skeletal dataset. In *3DOR-10th Eurographics Workshop on 3D Object Retrieval*, 1–6.
- Fan, H.; Yang, Y.; and Kankanhalli, M. 2021. Point 4d transformer networks for spatio-temporal modeling in point cloud videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 14204–14213.
- Fan, H.; Yang, Y.; and Kankanhalli, M. 2022. Point spatio-temporal transformer networks for point cloud video modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Fan, H.; Yu, X.; Ding, Y.; Yang, Y.; and Kankanhalli, M. 2021a. Pstnet: point spatio-temporal convolution on point cloud sequences. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Fan, H.; Yu, X.; Yang, Y.; and Kankanhalli, M. 2021b. Deep hierarchical representation of point cloud videos via spatio-temporal decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Han, T.; Xie, W.; and Zisserman, A. 2020a. Memory-augmented dense predictive coding for video representation learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 312–329.
- Han, T.; Xie, W.; and Zisserman, A. 2020b. Self-supervised co-training for video representation learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 5679–5690.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 16000–16009.
- He, K.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 9729–9738.
- Klaser, A.; Marszałek, M.; and Schmid, C. 2008. A spatio-temporal descriptor based on 3d-gradients. In *2008-19th British Machine Vision Conference (BMVC)*, 995–1004.
- Li, M.; Chen, S.; Chen, X.; Zhang, Y.; Wang, Y.; and Tian, Q. 2019. Actional-structural graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 3595–3603.
- Li, W.; Zhang, Z.; and Liu, Z. 2010. Action recognition based on a bag of 3d points. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 9–14.
- Liu, H.; Cai, M.; and Lee, Y. J. 2022. Masked discrimination for self-supervised learning on point clouds. *arXiv preprint arXiv:2203.11183*.
- Liu, J.; Wang, G.; Hu, P.; Duan, L.-Y.; and Kot, A. C. 2017. Global context-aware attention lstm networks for 3d action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 1647–1656.
- Liu, X.; Yan, M.; and Bohg, J. 2019. Meteornet: deep learning on dynamic 3D point cloud sequences. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 9246–9255.
- Min, Y.; Chai, X.; Zhao, L.; and Chen, X. 2019. FlickerNet: adaptive 3d gesture recognition from sparse point clouds. In *2019-30th British Machine Vision Conference (BMVC)*, 5.
- Min, Y.; Zhang, Y.; Chai, X.; and Chen, X. 2020. An efficient pointlstm for point clouds based gesture recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (ICCV)*, 5761–5770.
- Molchanov, P.; Yang, X.; Gupta, S.; Kim, K.; Tyree, S.; and Kautz, J. 2016. Online detection and classification of dynamic hand gestures with recurrent 3d convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4207–4215.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Pan, T.; Song, Y.; Yang, T.; Jiang, W.; and Liu, W. 2021. Videomoco: contrastive video representation learning with temporally adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 11205–11214.
- Pang, Y.; Wang, W.; Tay, F. E.; Liu, W.; Tian, Y.; and Yuan, L. 2022. Masked autoencoders for point cloud self-supervised learning. *arXiv preprint arXiv:2203.06604*.
- Qian, R.; Ding, S.; Liu, X.; and Lin, D. 2022. Static and dynamic concepts for self-supervised video representation learning. *arXiv preprint arXiv:2207.12795*.

- Qian, R.; Meng, T.; Gong, B.; Yang, M.-H.; Wang, H.; Belongie, S.; and Cui, Y. 2021. Spatiotemporal contrastive video representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 6964–6974.
- Shahroudy, A.; Liu, J.; Ng, T.-T.; and Wang, G. 2016. Ntu rgb+d: a large scale dataset for 3D human activity analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1010–1019.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019a. Skeleton-based action recognition with directed graph neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 7912–7921.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019b. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 12026–12035.
- Si, C.; Chen, W.; Wang, W.; Wang, L.; and Tan, T. 2019. An attention enhanced graph convolutional lstm network for skeleton-based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 1227–1236.
- Tong, Z.; Song, Y.; Wang, J.; and Wang, L. 2022. Videomae: masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*.
- Vieira, A. W.; Nascimento, E. R.; Oliveira, G. L.; Liu, Z.; and Campos, M. F. 2012. Stop: space-time occupancy patterns for 3d action recognition from depth map sequences. In *Iberoamerican congress on pattern recognition*, 252–259.
- Wang, H.; Guo, X.; Deng, Z.-H.; and Lu, Y. 2022a. Rethinking minimal sufficient representation in contrastive learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 16041–16050.
- Wang, H.; Yang, L.; Rong, X.; Feng, J.; and Tian, Y. 2021. Self-supervised 4d spatio-temporal feature learning via order prediction of sequential point cloud clips. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 3762–3771.
- Wang, J.; Liu, Z.; Wu, Y.; and Yuan, J. 2012. Mining actionlet ensemble for action recognition with depth cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1290–1297.
- Wang, R.; Chen, D.; Wu, Z.; Chen, Y.; Dai, X.; Liu, M.; Jiang, Y.-G.; Zhou, L.; and Yuan, L. 2022b. Bevt: bert pre-training of video transformers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 14733–14743.
- Wang, Y.; Xiao, Y.; Xiong, F.; Jiang, W.; Cao, Z.; Zhou, J. T.; and Yuan, J. 2020. 3DV: 3d dynamic voxel for action recognition in depth video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 511–520.
- Wei, Y.; Liu, H.; Xie, T.; Ke, Q.; and Guo, Y. 2022. Spatial-temporal transformer for 3d point cloud sequences. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1171–1180.
- Wen, H.; Liu, Y.; Huang, J.; Duan, B.; and Yi, L. 2022. Point primitive transformer for long-term 4D point cloud video understanding. *arXiv preprint arXiv:2208.00281*.
- Xiao, Y.; Chen, J.; Wang, Y.; Cao, Z.; Zhou, J. T.; and Bai, X. 2019. Action recognition for depth video using multi-view dynamic images. *Information Sciences*, 480: 287–304.
- Xie, S.; Gu, J.; Guo, D.; Qi, C. R.; Guibas, L.; and Litany, O. 2020. Pointcontrast: unsupervised pre-training for 3d point cloud understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 574–591.
- Xie, Z.; Zhang, Z.; Cao, Y.; Lin, Y.; Bao, J.; Yao, Z.; Dai, Q.; and Hu, H. 2022. Simmim: a simple framework for masked image modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 9653–9663.
- Yang, S.; Liu, J.; Lu, S.; Er, M. H.; and Kot, A. C. 2021. Skeleton cloud colorization for unsupervised 3d action representation learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 13423–13433.
- Yang, Y.; Feng, C.; Shen, Y.; and Tian, D. 2018. Foldingnet: point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 206–215.
- You, Q.; and Jiang, H. 2019. Action4d: online action recognition in the crowd and clutter. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 11857–11866.
- Yu, X.; Tang, L.; Rao, Y.; Huang, T.; Zhou, J.; and Lu, J. 2022. Point-bert: pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 19313–19322.
- Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow twins: self-supervised learning via redundancy reduction. In *International Conference on Machine Learning (ICML)*, 12310–12320.
- Zhang, P.; Lan, C.; Xing, J.; Zeng, W.; Xue, J.; and Zheng, N. 2019. View adaptive neural networks for high performance skeleton-based human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(8): 1963–1978.
- Zhong, J.-X.; Zhou, K.; Hu, Q.; Wang, B.; Trigoni, N.; and Markham, A. 2022. No pain, big gain: classify dynamic point cloud sequences with static models by fitting feature-level space-time surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 8510–8520.