

Versatile Multi-Modal Pre-Training for Human-Centric Perception

Fangzhou Hong¹, Liang Pan¹, Zhongang Cai^{1,2,3}, Ziwei Liu¹✉

¹S-Lab, Nanyang Technological University ²SenseTime Research ³Shanghai AI Laboratory

{fangzhou001, liang.pan, ziwei.liu}@ntu.edu.sg caizhongang@sensetime.com

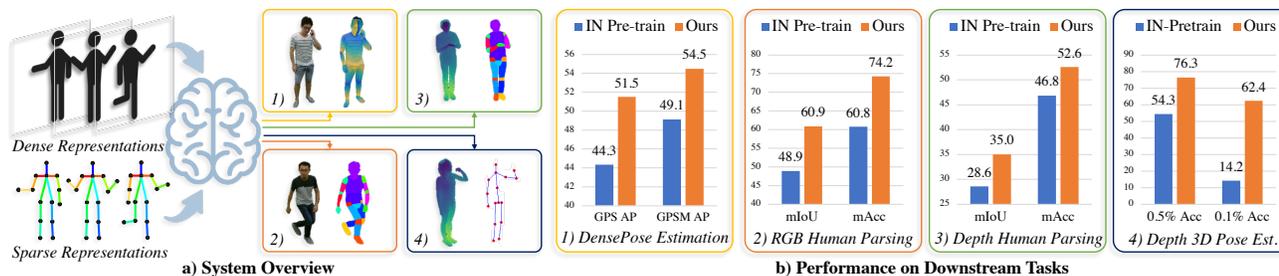


Figure 1. **An Overview of HCMoCo.** a) We present HCMoCo, a versatile multi-modal pre-training framework that takes multi-modal observations of human body as input for human-centric perception. The pre-train models can be transferred to various human-centric downstream tasks with different modalities. b) Our HCMoCo shows superior performance on all four downstream tasks, especially for data-efficient settings (10% DensePose, 20% RGB/depth human parsing, 0.5/0.1% 3D pose estimation). ‘IN’ stands for ImageNet.

Abstract

Human-centric perception plays a vital role in vision and graphics. But their data annotations are prohibitively expensive. Therefore, it is desirable to have a versatile pre-train model that serves as a foundation for data-efficient downstream tasks transfer. To this end, we propose the **Human-Centric Multi-Modal Contrastive Learning** framework **HCMoCo** that leverages the multi-modal nature of human data (e.g. RGB, depth, 2D keypoints) for effective representation learning. The objective comes with two main challenges: dense pre-train for multi-modality data, efficient usage of sparse human priors. To tackle the challenges, we design the novel *Dense Intra-sample Contrastive Learning* and *Sparse Structure-aware Contrastive Learning* targets by hierarchically learning a modal-invariant latent space featured with continuous and ordinal feature distribution and structure-aware semantic consistency. HCMoCo provides pre-train for different modalities by combining heterogeneous datasets, which allows efficient usage of existing task-specific human data. Extensive experiments on four downstream tasks of different modalities demonstrate the effectiveness of HCMoCo, especially under data-efficient settings (7.16% and 12% improvement on DensePose Estimation and Human Parsing). Moreover, we demonstrate the versatility of HCMoCo by exploring cross-modality supervision and missing-modality inference, validating its strong ability in cross-modal association and reasoning. Codes

are available at <https://github.com/hongfz16/HCMoCo>.

1. Introduction

As a long-standing problem, human-centric perception has been studied for decades, ranging from sparse prediction tasks, such as human action recognition [8, 27, 42, 50], 2D keypoints detection [2, 26, 43, 48] and 3D pose estimation [22, 31, 40], to dense prediction tasks, such as human parsing [7, 11, 12, 25] and DensePose prediction [14]. Unfortunately, to train a model with reasonable generalizability and robustness, an enormous amount of labeled real data is necessary, which is extremely expensive to collect and annotate. Therefore, it is desirable to have a versatile pre-train model that can serve as a foundation for all the aforementioned human-centric perception tasks.

With the development of sensors, the human body can be more conveniently perceived and represented in **multi-modalities**, such as RGB, depth, and infrared. In this work, we argue that *the multi-modality nature of human-centric data can induce effective representations* that transfer well to various downstream tasks, due to three major **advantages**: **1)** Learning a modal-invariant latent space through pre-training helps efficient task-relevant mutual information extraction. **2)** A single versatile pre-train model on multi-modal data facilitates multiple downstream tasks using various modalities. **3)** Our multi-modal pre-train setting bridges heterogeneous human-centric datasets through their common modality, which benefits the generalizability of pre-train models.

✉ Corresponding author

We mainly explore two groups of modalities as shown in Fig. 1 a): dense representations (*e.g.* RGB, depth, infrared) and sparse representations (*e.g.* 2D keypoints, 3D pose). Dense representations can provide rich texture and/or 3D geometry information. But they are mostly low-level and noisy. On the contrary, sparse representations obtained by off-the-shelf tools [4, 9] are semantic and structured. But the sparsity results in insufficient details. We highlight that it is non-trivial to integrate these heterogeneous modalities into a unified pre-training framework for the following two main **challenges**: **1)** learning representations suitable for dense prediction tasks in the multi-modality setting; **2)** using weak priors from sparse representations effectively for pre-training.

Challenge 1: Dense Targets. Existing methods [21, 30] perform contrastive learning densely on pixel-level features to achieve view-invariance for dense prediction tasks. However, those methods require multiple views of a static 3D scene [10], which is inapplicable for human-centric applications with only single view. Furthermore, it is preferable to learn representations that are continuously and orderly distributed over the human body. In light of this, we generalize the widely used InfoNCE [33] and propose a dense intra-sample contrastive learning objective that applies a soft pixel-level contrastive target, which can facilitate learning ordinal and continuous dense feature distributions.

Challenge 2: Sparse Priors. To employ priors in contrastive learning, previous works [3, 23, 46] mainly use the supervision to generate semantically positive pairs. However, these methods only focus on the sample-level contrastive learning, which means each sample is encoded to a global embedding. It is not optimal for human dense prediction tasks. To this end, we propose a sparse structure-aware contrastive learning target, which uses semantic correspondences across samples as positive pairs to complement positive intra-sample pairs. Particularly, leveraging sparse human priors leads to an embedding space where semantically corresponding parts are aligned more closely.

To sum up, we propose **HCMoCo**, a **Human-Centric multi-Modal Contrastive learning** framework for versatile multi-modal pre-training. To fully leverage multi-modal observations, HCMoCo effectively utilizes both dense measurements and sparse priors using the following three-levels hierarchical contrastive learning objectives: **1)** sample-level modality-invariant representation learning; **2)** dense intra-sample contrastive learning; **3)** sparse structure-aware contrastive learning. As an effort towards establishing a comprehensive multi-modal human parsing benchmark dataset, we label human segments for RGB-D images from NTU RGB+D dataset [42], and contribute the **NTURGBD-Parsing-4K** dataset. To evaluate HCMoCo, we transfer our pre-train model to four human-centric downstream tasks using different modalities, including DensePose es-

timation (RGB) [14], human parsing using RGB [22] or depth frames, and 3D pose estimation (depth) [16]. Under full set and data-efficient training settings, HCMoCo constantly achieves better performance than training from scratch or pre-train on ImageNet. To name a few, as shown in Fig. 1 b), we achieve **7.16%** improvement in terms of GPS AP on 10% training data of DensePose estimation; **12%** improvement in terms of mIoU on 20% training data of Human3.6M human parsing. Moreover, we evaluate the modal-invariance of the latent space learned by HCMoCo for dense prediction on NTURGBD-Parsing-4K with two settings: cross-modality supervision and missing-modality inference. Compared against conventional contrastive learning targets, our method improves the segmentation mIoU by **29%** and **24%** for the two settings, respectively. To the best of our knowledge, we are the first to study multi-modal pre-training for human-centric perception.

The main **contributions** are summarized below: **1)** As the first endeavor, we provide an in-depth analysis for human-centric pre-training, which is formulated as a challenging multi-modal contrastive learning problem. **2)** Together with the novel hierarchical contrastive learning objectives, a comprehensive framework HCMoCo is proposed for effective pre-training for human-centric tasks. **3)** Through extensive experiments, HCMoCo achieves superior performance than existing methods, and meanwhile shows promising modal-invariance properties. **4)** To benefit multi-modal human-centric perception, we contribute an RGB-D human parsing dataset, NTURGBD-Parsing-4K.

2. Related Work

Human-Centric Perception. Many efforts have been put into human-centric perception in decades. Lots of work in 2D keypoint detection [2, 26, 43, 48] has achieved robust and accurate performance. 3D pose estimation has long been a challenging problem and is approached from two aspects, lifting from 2D keypoints [22, 31, 40] and predicting from depth map [16, 49]. Human parsing can be defined in two ways. The first one parses garments together with visible body parts [11, 12, 25]. The second one only focuses on parsing human parts [7, 20, 22]. In this work, we focus on the second setting because the depth and 2D keypoints do not contain the texture information needed for garment parsing. There are a few works [19, 32] about human parsing on depth maps. However, the data and annotations are too coarse or unavailable. To further push the accuracy of human-centric perception, DensePose [14, 44] is proposed to densely model each human body surface point. The cost of DensePose annotation is enormous. Therefore, we also explore data-efficient learning of DensePose.

Multi-Modal Contrastive Learning. Multi-modality naturally provides different views of the same sample which fits well into the contrastive learning framework. CMC [45]

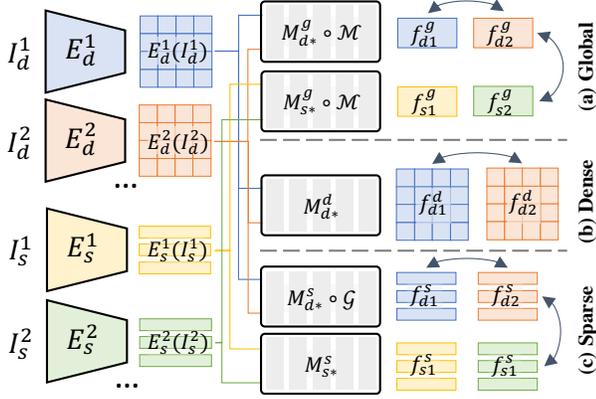


Figure 2. **Illustration of the general paradigm of HCMoCo.** We group modalities of human data into dense I_d^* and sparse representations I_s^* . Three levels of embeddings are extracted (3.1). Combining the nature of human data and tasks (3.2), we present contrastive learning targets for each level of embedding (3.3).

proposes the first multi-view contrastive learning paradigm which takes any number of views. CLIP [39] learns a joint latent space from large-scale paired image-language dataset. Extensive studies [1, 15, 17, 34, 35, 41] focus on video-audio contrastive learning. Recently, 2D-3D contrastive learning [21, 29, 30] has also been studied with the development in 3D computer vision. In this work, aside from commonly used modalities, we also explore the potential of 2D keypoints in human-centric contrastive learning.

3. Our Approach

In this section, we first introduce the general paradigm of HCMoCo (3.1). Following the design principles (3.2), hierarchical contrastive learning targets are formally introduced (3.3). Next, an instantiation of HCMoCo is introduced (3.4). Finally, we propose two applications of HCMoCo to show the versatility (3.5).

3.1. HCMoCo

As shown in Fig. 2, HCMoCo takes multiple modalities of perceived human body as input. The target is to learn human-centric representations, which can be transferred to downstream tasks. The input modalities can be categorized into dense and sparse representations. Dense representations I_d^* are the direct output of imaging sensors, *e.g.* RGB, depth, infrared. They typically contain rich information but are low-level and noisy. Sparse representations are structured abstractions of the human body, *e.g.* 2D keypoints, 3D pose, which can be formulated as graph $I_s^* = G(V, E)$. Different representations of the same view of a human should be spatially aligned, which means intra-sample correspondences can be obtained for dense contrastive learning. HCMoCo aims to pre-train multiple encoders E_d^* and E_s^* that produce embeddings of dense representations and sparse

representations for downstream tasks transfer.

To support dense downstream tasks, other than the usual sample-level global embeddings used in [5, 6, 13, 18, 28, 45], we propose to consider different levels of embeddings *i.e.* global embeddings f^g , sparse embeddings f^s and dense embeddings f^d ¹, which are defined as follows: **1)** For dense representations I_d , the global embedding is obtained by applying a mapper network M_d^g to the mean pooling \mathcal{M} of the corresponding feature map, which is formulated as $f_{d1}^g = M_{d1}^g \circ \mathcal{M} \circ E_d(I_{d1})$. Similarly, for sparse representations I_s , the global embedding is defined as $f_{s1}^g = M_{s1}^g \circ \mathcal{M} \circ E_s(I_{s1})$. **2)** Sparse embeddings have the same size as that of sparse representations. Formally, for sparse representations $I_s = G(V, E)$, where $V \in \mathbb{R}^{J \times K}$, the corresponding sparse embedding is defined as $f_{s1}^s = M_{s1}^s \circ E_s(I_{s1})$, where $f_{s1}^s \in \mathbb{R}^{J \times K'}$, M_{s1}^s is a mapper network. For dense representations, the corresponding sparse features are pooled from the dense feature map using the correspondences \mathcal{G} . Then the sparse features are mapped to sparse embeddings as $f_{d1}^s = M_{d1}^s \circ \mathcal{G} \circ E_d(I_{d1})$. **3)** Dense embeddings are only defined on dense representations, which is formulated as $f_{d1}^d = M_{d1}^d \circ E_d(I_{d1})$. With three levels of embeddings defined, we formulate the overall learning objective as

$$\mathcal{L} = \lambda_g \mathcal{L}_g(f^g) + \lambda_d \mathcal{L}_d(f^d) + \lambda_s \mathcal{L}_s(f^s), \quad (1)$$

which is analyzed and explained as follows.

3.2. Principles of Learning Targets Design

In this subsection, we analyze the intuitions when designing learning targets, which makes the following three principles. **1) Mutual Information Maximization:** Inspired by [36, 47], we propose to maximize the lower bound on mutual information, which has been proved by many previous works [5, 6, 18, 45] to be able to produce strong pre-train models. **2) Continuous and Ordinal Feature Distribution:** Inspired by the property of human-centric perception, it is desirable for the feature maps of the human body to be continuous and ordinal. The human body is a structural and continuous surface. The dense predictions, *e.g.* human parsing [11, 12, 25], DensePose [14], are also continuous. Therefore, such property should also be reflected in the learned representations. Besides, for an anchor point on human surfaces, closer points have higher probabilities of sharing similar semantics with the anchor point than that of far away points. Therefore, the learned dense representations should also align with such ordinal relationship. **3) Structure-Aware Semantic Consistency:** Sparse representations are abstractions of the human body, which contains valuable structural semantics about the human body. Instead of identity information, the human pose and struc-

¹For easier understanding of the notations, the superscripts of f and M stand for the kind of embeddings. The subscripts stand for the kind of representations ('g' for 'global'; 'd' for 'dense'; 's' for 'sparse').

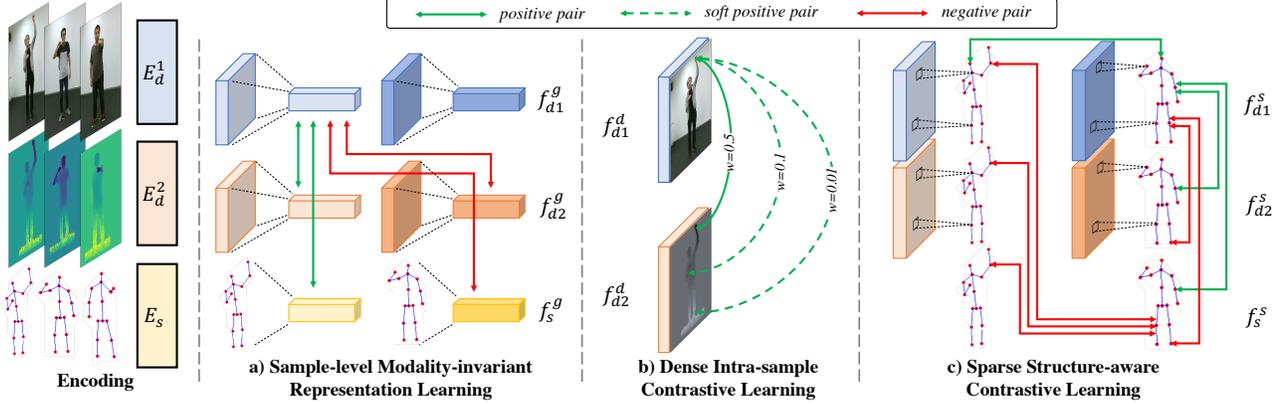


Figure 3. **Our Proposed Instantiation of HCMoCo.** For dense representations, we choose to use RGB and depth. For sparse representations, 2D keypoints are used for its convenience to obtain. a) At sample-level, the global embeddings are used for modality-invariant representation learning. b) Between paired dense embeddings, soft contrastive learning target is proposed for continuous and ordinal feature learning. c) Using human prior provided by sparse representations, intra- and inter-sample contrastive learning targets are proposed.

ture understanding are the keys to our target downstream tasks. Therefore, it is reasonable to eliminate the identity information and enhance the structure information by enforcing structure-aware semantic consistency where semantically close embeddings (*e.g.* embeddings of left hands from different samples) are pulled close and vice versa.

3.3. Hierarchical Contrastive Learning Targets

Based on the above three principles, we formally define hierarchical contrastive learning targets in this subsection.

Sample-level modality-invariant representation learning aims at learning a joint latent space at the sample level using global embeddings, which fulfills the first principle. Inspired by [45], the learning target can be formulated as

$$\mathcal{L}_g = - \mathbb{E}_{\substack{F_1^g, F_2^g \in S_g \\ f_1^g \in F_1^g}} \left[\log \frac{\exp(f_1^g \cdot \bar{f}_2^g / \tau)}{\sum_{f_2^g \in F_2^g} \exp(f_1^g \cdot f_2^g / \tau)} \right], \quad (2)$$

where F_*^g is a set of global embeddings of one modality, S_g is the set of F_*^g of all modalities, \bar{f}_2^g is the embedding of the paired view of that of f_1^g , τ is the temperature. It should be noticed that f_1^g can be sampled from the global embeddings of either dense or sparse representations.

Dense intra-sample contrastive learning is operated on the paired dense representations. For any two paired dense embeddings $f_{d1}^d, f_{d2}^d \in \mathbb{R}^{H \times W \times K'}$, to simultaneously satisfy the first and the second principle, the dense intra-sample contrastive learning target between them is defined in a ‘soft’ way as

$$\mathcal{L}_d^{12} = - \mathbb{E}_{\substack{x, y \\ m, n}} \left[\mathcal{W}_{xy}^{mn} \log \frac{\exp(f_{d1}^d(x, y) \cdot f_{d2}^d(m, n) / \tau)}{\sum_{x', y'} \exp(f_{d1}^d(x, y) \cdot f_{d2}^d(x', y') / \tau)} \right], \quad (3)$$

where \mathcal{W}_{xy}^{mn} is the weight, τ is the temperature, $(x, y), (m, n), (x', y')$ are coordinates on the dense repre-

sentation, $1 \leq x, x', m \leq H, 1 \leq y, y', n \leq W$. The above equation is a generalized version of InfoNCE [33]. InfoNCE is a special case when \mathcal{W}_{xy}^{mn} is set to 1 if $x = m$ and $y = n$ else 0. We use the normalized distances as the weights, which is formulated as

$$\mathcal{W}_{xy}^{mn} = \frac{\exp(\sqrt{(x-m)^2 + (y-n)^2})}{\sum_{x', y'} \exp(\sqrt{(x-x')^2 + (y-y')^2})}. \quad (4)$$

For each pair of dense representations, the above learning target is calculated between each pair of dense embeddings. Therefore, the whole learning target is defined as

$$\mathcal{L}_d = \mathbb{E}_{\substack{F_1^d, F_2^d \in S_d \\ f_{d1}^d, f_{d2}^d \in F_1^d, F_2^d}} \mathcal{L}_d^{12}(f_{d1}^d, f_{d2}^d), \quad (5)$$

where F_*^d is a set of dense embeddings of one modality, S_d is the set of all F_*^d , f_{d1}^d and f_{d2}^d are two paired embeddings. It should be noticed that the ‘soft’ learning target cannot guarantee an ordinal feature distribution. Instead, it serves as a computationally efficient relaxation of the requirement of ordinal distribution.

Sparse structure-aware contrastive learning takes two sparse representations f_1^s and f_2^s as inputs. The paired features f_{1j}^s and f_{2j}^s (*i.e.* features of the j -th joint) should be pulled close while unpaired features are pushed away. The two sparse representations can be sampled from the same or different modalities, intra- or inter-sample. The intra-sample alignment satisfies the first principle. The inter-sample alignment follows the third principle. The sparse structure-aware contrastive learning target is formulated as

$$\mathcal{L}_s = - \mathbb{E}_{\substack{F_1^s, F_2^s \in S_s \\ j; f_{1j}^s, f_{2j}^s \in \{F_1^s, F_2^s\}}} \left[\log \frac{\exp(f_{1j}^s \cdot f_{2j}^s / \tau)}{\sum_{j'; f_{ij'}^s \in \{F_1^s, F_2^s\}} \exp(f_{1j}^s \cdot f_{ij'}^s / \tau)} \right], \quad (6)$$

where F_*^s is a set of sparse embeddings of one modality, S_s is the set of F_*^s , τ is the temperature, f_{1j}^s, f_{2j}^s are sam-

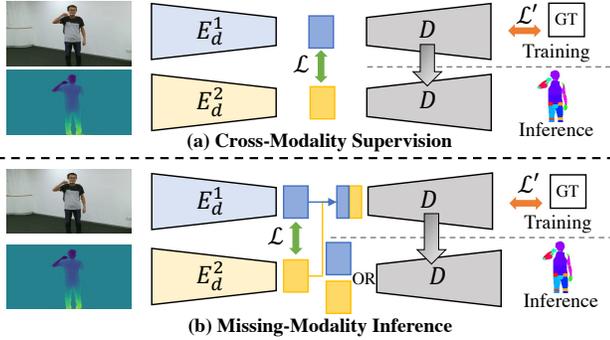


Figure 4. Pipelines of Two Applications of HCMoCo.

pled from the union of F_1^s and F_2^s . To conclude, the overall learning target is formulated as Eq. 1, where λ_* are the weights to balance the targets.

3.4. Instantiation of HCMoCo

In this section, we introduce an instantiation of HCMoCo. As shown in Fig. 3, for dense representations, we use RGB and depth. Large-scale paired human RGB and depth data is easy to obtain with affordable sensors *e.g.* Kinect. These two modalities are the most commonly encountered in human-centric tasks [7, 11, 12, 22, 25]. Moreover, a proper pre-train model for depth is highly desired. Therefore, RGB and depth are reasonable choices of human dense representations, both of which are easy to acquire and important to downstream tasks. For sparse representations, 2D keypoints are used, which provide positions of human body joints in the image coordinate. Off-the-shelf tools [4, 9] are available to quickly and robustly extract human 2D keypoints given RGB images. Using 2D keypoints as the sparse representation is a good balance between the amount of human prior and acquisition difficulty.

For RGB inputs I_d^1 , an image encoder E_d^1 [43] is applied to obtain feature maps $E_d^1(I_d^1)$. Similarly, for depth inputs I_d^2 , an image encoder [43] or 3D encoder [37, 38] E_d^2 can be applied to extract feature maps $E_d^2(I_d^2)$. 2D keypoints I_s are encoded by a GCN-based encoder [51] E_s to produce sparse features $E_s(I_s)$. Mapper networks comprise a single linear layer and a normalization operation.

As for the implementation of contrastive learning targets, we choose to use a memory pool to store all the global embeddings which are updated in a momentum way. Sparse and dense embeddings cannot all fit in memory. Therefore, for the last two types of contrastive learning targets, the negative samples are sampled within a mini-batch.

3.5. Versatility of HCMoCo

On top of the pre-train framework HCMoCo, we propose to further extend it on two direct applications: cross-modality supervision and missing-modality inference. The extensions are based on the key design of HCMoCo: dense

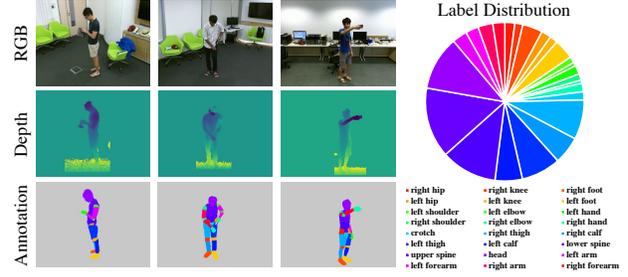


Figure 5. Illustration of the RGB-D human parsing dataset NTURGBD-Parsing-4K.

intra-sample contrastive learning target. With the feature maps of different modalities aligned, it is straightforward to implement the two extensions, which are shown in Fig. 4.

Cross-Modality Supervision is a novel task where we train the network on the source modality, while test on the target modality. This is a practical scenario where people transfer the knowledge of some single modality dataset to other modalities. At training time, an additional downstream task head (*e.g.* segmentation head) D is attached to the backbone of the source modality. The hierarchical contrastive learning targets \mathcal{L} together with downstream task loss \mathcal{L}' are used for end-to-end training. At inference time, D is attached to the backbone of the target modality. The extracted feature maps of the target modality are passed to D for prediction.

Missing-Modality Inference is another novel task where we train the network using multi-modal data and inference on single modality. Multi-modal data collection in practice would inevitably result in data with incomplete modalities, which brings the requirement of missing-modality inference. At training time, the feature maps of multiple modalities are fused using max-pooling and fed to a downstream task head D . Similarly, hierarchical contrastive learning targets \mathcal{L} and downstream task loss \mathcal{L}' are used for co-training. At inference time, the feature map of a single modality is passed to D for missing-modality inference.

4. NTURGBD-Parsing-4K Dataset

Although RGB human parsing has been well studied [7, 11, 12, 25], human parsing on depth [19, 32] or RGB-D data has not been fully addressed due to the lack of labeled data. Therefore, we contribute the first RGB-D human parsing dataset: NTURGBD-Parsing-4K. The RGB and depth are uniformly sampled from NTU RGB+D (60/120) [27, 42]. As shown in Fig. 5, we annotate 24 human parts for paired RGB-D data. The partition protocols follow that of [22]. The train and test set both have 1963 samples. The whole dataset contains 3926 samples. Hopefully, by contributing this dataset, we could promote the development of both human perception and multi-modality learning.

Table 1. DensePose Estimation Results on COCO. * randomly initializes the model before pre-training. † initializes the model by ImageNet pre-train before pre-training. All results in [%].

Method	Pre-train Datasets	Full Data				10% Data			
		BBox AP	GPS AP	GPSM AP	IOU AP	BBox AP	GPS AP	GPSM AP	IoU AP
From Scratch	-	57.27	62.03	63.61	65.88	39.38	35.75	41.62	49.92
CMC* [45]	NTURGBD+MPII	60.33	64.97	65.66	66.96	44.92	43.84	47.94	54.00
MMV* [1]	NTURGBD+MPII	59.89	64.23	65.47	67.03	43.24	41.40	45.99	52.52
Ours*	NTURGBD+MPII	61.33	65.89	66.92	67.66	47.76	48.47	51.65	56.15
IN Pre-train	-	62.66	66.48	67.42	68.63	48.28	44.34	49.11	56.11
CMC† [45]	NTURGBD+MPII	62.76	66.16	67.30	68.06	49.21	48.82	52.57	57.94
MMV† [1]	NTURGBD+MPII	62.97	66.67	67.51	68.29	50.16	50.28	53.54	58.32
Ours†	NTURGBD+MPII	63.11	67.33	68.12	68.72	50.29	51.50	54.47	58.66
CMC† [45]	NTURGBD+COCO	63.58	67.22	67.77	68.46	51.77	53.53	56.18	59.37
Ours†	NTURGBD+COCO	62.95	67.77	68.29	68.63	52.18	54.01	56.64	59.93

Table 2. Human Parsing Results on Human3.6M. * randomly initializes the model before pre-training. † initializes the model by ImageNet pre-train before pre-training. All results in [%].

Method	Full Data			20% Data			10% Data			1% Data		
	mIoU	mAcc	aAcc	mIoU	mAcc	aAcc	mIoU	mAcc	aAcc	mIoU	mAcc	aAcc
From Scratch	44.13	58.88	98.82	42.41	56.25	98.81	32.61	43.76	98.52	7.27	10.97	97.45
CMC* [45]	54.33	68.01	99.09	52.10	65.65	99.03	48.37	61.18	98.95	14.61	20.07	98.07
MMV* [1]	52.69	65.82	99.06	50.66	63.55	99.01	46.23	58.52	98.90	12.86	17.10	97.94
Ours*	61.36	75.09	99.25	59.17	73.44	99.19	57.08	71.75	99.13	16.55	22.27	98.18
IN Pre-train	56.90	69.94	99.14	48.86	60.75	98.97	44.55	56.86	98.87	14.65	20.22	98.09
CMC† [45]	58.93	71.70	99.20	57.41	70.13	99.17	54.35	67.47	99.09	17.77	23.77	98.20
MMV† [1]	59.08	71.57	99.20	57.28	69.69	99.17	53.86	66.46	99.08	17.66	23.54	98.20
Ours†	62.50	75.84	99.27	60.85	74.23	99.23	58.28	71.99	99.17	20.78	27.52	98.34

5. Experiments

5.1. Experimental Setup

Implementation Details. The default RGB and depth encoders are HRNet-W18 [43]. The default datasets for pre-train are NTU RGB+D [27] and MPII [2]. The former provides paired indoor human RGB, depth, and 2D keypoints, The latter provides in-the-wild human RGB and 2D keypoints. Mixing human data from different domains helps our pre-train models adapt to a wilder domain.

Downstream Tasks. We test our pre-train models on four different human-centric downstream tasks, two on RGB images and two on depth. **1)** DensePose estimation on COCO [14]: DensePose aims at mapping pixels of the observed human body to the surface of a 3D human body, which is a highly challenging task. **2)** RGB human parsing on Human3.6M [22]. Human3.6M provides pure human part segmentation, which aligns with our objectives. We uniformly sample 2fps of the video for training and evaluation. **3)** Depth human parsing on NTURGBD-Parsing-4K. **4)** 3D pose estimation from depth maps on ITOP [16] (only side view). For all the above downstream tasks, we use the pre-train backbones for end-to-end fine-tune.

Comparison Methods. Since there are few previous human-centric multi-modal pre-train methods, we propose

to use general multi-modal contrastive learning methods CMC [45] and MMV [1] as the baselines. Although there are other multi-modal contrastive learning works, they either require the multi-view calibration [21] or focus on multi-modal downstream tasks [17, 29] and therefore are not suitable for comparison. In addition, for RGB tasks, we also experiment under two settings, one initializes encoders with supervised ImageNet [24] (IN) pre-train while the other does not.

5.2. Performance on Downstream Tasks

DensePose Estimation. As shown in Tab. 1, we test DensePose estimation [14] under two settings: full and 10% of the training data. The trained models are tested on the full validation set of DensePose. Firstly, if not using IN pre-train, our pre-train model significantly outperforms both ‘From Scratch’ and two baseline methods. Especially under 10% of training data, **12.7%** improvement in terms of GPS AP is observed. And our pre-train model even outperforms that using IN pre-train by **4.13%** in terms of GPS AP. When we use IN pre-train as initialization, which is a common practice for 2D tasks, our method still outperforms all the baselines. Our method surpasses IN pre-train by **7.2%** and **5.4%** in terms of GPS/GPSM AP under 10% setting. To further test the performance of in-domain transfer, we also

Table 3. Ablation Study on Densepose/ Human3.6M/ ITOP/ NTURGBD-Parsing-4K. All results in [%].

Method	DensePose 10%				ITOP 0.1%/ 0.2%		Human3.6M 10%		NTURGBD 20%	
	BBox	GPS	GPSM	IoU	Acc	Acc	mIoU	mAcc	mIoU	mAcc
Sample-level Mod-invariant	49.21	48.82	52.57	57.94	57.73	50.08	54.35	67.47	30.40	51.54
+ Hard Dense Intra-sample	49.40	49.14	52.49	57.30	56.43	54.05	55.36	68.43	31.26	51.54
+ Soft Dense Intra-sample	50.21	50.25	53.42	57.70	62.33	51.50	56.35	69.26	32.20	51.06
+ Sparse Structure-aware	50.29	51.50	54.47	58.66	65.83	62.36	58.28	71.99	35.01	52.55

pre-train models using training sets of NTU RGB+D and COCO. The performance gain under 10% setting further improves to **9.7%** and **7.5%** in terms of GPS/GPSM AP.

RGB Human Parsing. As shown in Tab. 2, we test four settings on Human3.6M [22]: full, 20%, 10% and 1% training data. In all settings, our method outperforms all baselines in all metrics. On full training data, we outperform IN pre-train by **5.6%** in terms of mIoU. The performance gain increases with the amount of training data decreases. It is worth noticing that with only 10% of training data, our method outperforms IN pre-train with full training data.

Table 4. Human Parsing Results on NTURGBD-Parsing-4K [%].

Method	Full Data			20% Data		
	mIoU	mAcc	aAcc	mIoU	mAcc	aAcc
IN Pre-train	37.49	57.52	98.36	28.56	46.81	98.10
CMC [45]	38.20	58.73	98.39	30.40	51.54	98.02
MMV [1]	38.09	58.49	98.37	30.41	50.62	98.07
Ours	39.32	58.79	98.47	35.01	52.55	98.53

Depth Human Parsing. As shown in Tab. 4, we test the pre-train depth backbone on our proposed Dataset NTURGBD-Parsing-4K with all training data and 20% training data. We outperform all baselines on two settings. Especially, only using 20% of training data, we surpass IN pre-train by **6.4%** and MMV [1] by **4.6%** in terms of mIoU.

Table 5. 3D Pose Estimation Results on ITOP. All results in [%].

Method	100%	10%	1%	0.5%	0.2%	0.1%
IN Pre-train	85.19	83.44	77.20	54.31	13.27	14.21
CMC [45]	87.08	85.36	79.49	75.07	57.73	50.08
MMV [1]	86.13	83.49	79.70	71.70	60.83	54.44
Ours	87.19	85.49	78.71	76.34	65.83	62.36

3D Pose Estimation. As shown in Tab. 5, we test the pre-train depth backbone on ITOP [16] with six different ratios of training data. Our pre-train model outperforms all baselines on most settings. With only 10% training data, the accuracy of our method outperforms that of IN pre-train with all training data. It is also worth noticing that 0.1% of training data are 17 samples, which makes this a few-shot learning setting. With such limited training data, IN pre-train barely produce meaningful results, while our method improves the accuracy by **48.2%**.

5.3. Ablation Study

In this subsection, we perform a thorough ablation study on HCMoCo to justify the design choices. As shown in Tab. 3, we firstly report the results of only applying sample-level modality-invariant representation learning. Then we add dense intra-sample contrastive learning and sparse structure-aware contrastive learning in order. To further demonstrate the effect of the ‘soft’ design in dense intra-sample contrastive learning, we also report results of the ‘hard’ learning target, which takes the form of a classic InfoNCE [33]. We report the results of the ablation study on all four downstream tasks under data-efficient settings.

For DensePose estimation, it is important to learn feature maps that are continuously and ordinally distributed, which is the expected result of soft dense intra-sample contrastive learning. The performance gain of the soft learning target over the hard counterpart justifies the observation and the learning target design. The dense intra-sample contrastive learning also shows superiority on three other downstream tasks, which shows the importance of fine-grained contrastive learning targets for dense prediction tasks.

Explicitly injecting human prior into the network through sparse structure-aware contrastive learning also proves its effectiveness by further improving the performance on DensePose. Thanks to the strong hints provided by 2D keypoints, the performance of 3D pose estimation is improved. Moreover, the sparse structure-aware contrastive learning boosts the performance of human parsing both on RGB and depth maps by **1.9%** and **2.8%** respectively in terms of mIoU. Although 2D keypoints are sparse priors, they still provide the rough location of each part of the human body, which facilitate the feature alignment of same body parts. To summarize, the sparse and dense learning targets both contribute to the performance of our methods, which is in line with our analysis.

5.4. Performance on HCMoCo Versatility

Cross-Modality Supervision. We test the cross-modality supervision pipeline on the task of human parsing on NTURGBD-Parsing-4K because it has two modalities and respective dense annotations. Two baseline methods are adopted: 1) using CMC [45] contrastive learning target; 2) no contrastive learning target. For a fair comparison, the backbones of all methods are initialized by CMC [45] pre-train. At training time, the target modality of training data

Table 6. Cross-Modality Supervised Human Parsing Results on NTURGBD-Parsing-4K. All results in [%].

Method	RGB → Depth			Depth → RGB		
	mIoU	mAcc	aAcc	mIoU	mAcc	aAcc
No Contrastive	3.94	4.36	92.24	3.71	4.03	91.63
CMC [45]	3.86	5.59	86.81	3.85	4.27	91.75
Ours	33.19	54.38	94.70	26.80	48.80	92.84

is not available. We experiment on two settings where we supervise on RGB, test on depth (RGB → Depth), and vice versa (Depth → RGB). As shown in Tab. 6, our method outperforms both baselines under two settings. Specifically, our method improves the mIoU of both settings by **29.2%** and **23.0%**, respectively. Even compared to methods with direct supervision, we can achieve comparable results.

Table 7. Missing-Modality Human Parsing Results on NTURGBD-Depth. All results in [%].

Method	Only RGB			Only Depth		
	mIoU	mAcc	aAcc	mIoU	mAcc	aAcc
No Contrastive	13.45	14.77	93.35	24.41	30.49	95.27
CMC [45]	19.62	28.19	92.94	16.58	19.83	93.94
Ours	43.88	64.27	96.15	43.98	63.66	96.34

Missing-Modality Inference. For missing-modality inference, we report the experiments on the same dataset and same baselines as above. As shown in Tab. 7, with no pixel-level alignment, the two baseline methods struggle in two missing-modality settings *i.e.* ‘Only RGB’ and ‘Only Depth’. While our method improves the segmentation mIoU by **24.3%** and **19.6%** on two settings.

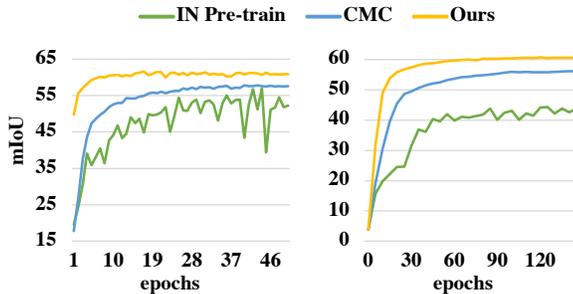


Figure 6. Validation mIoU Changes with Training Epochs Increase. Left: Human3.6M human parsing full training set. Right: Human3.6M human parsing 20% training set.

5.5. Further Analysis

Faster Convergence. One of the advantages of pre-training is the fast convergence speed when transferred to downstream tasks. Our HCMoCo also shows superiority in this feature. We log the validation mIoU of Human3.6M human parsing at different training epochs. As shown in Fig. 6, compared with IN pre-train and CMC [45], our pre-train model is able to converge within a few training epochs in

Table 8. Experiments on changing the backbone. * stands for ‘IN Pre-train’ for DensePose and ‘From Scratch’ for NTURGBD-Parsing-4K. All results in [%].

Method	DensePose 10%				NTURGBD 20%	
	BBox	GPS	GPSM	IoU	mIoU	mAcc
*	55.10	54.60	57.60	61.73	45.36	59.51
CMC [45]	53.88	54.62	57.46	61.14	48.74	62.94
Ours	54.55	55.80	58.36	61.75	49.43	63.52

both the full training data and data-efficient settings.

Changing Backbone. So far our experiments are all performed on HRNet-W18. To further demonstrate HCMoCo’s performance on other backbones, for the 2D backbone, we also experiment with HRNet-W32 [43]. For the depth backbone, we choose to test with PointNet++ [38]. For the RGB pre-train model, we experiment on the 10% DensePose estimation. For the depth pre-train model, we experiment on the 20% NTURGBD-Parsing-4K. As shown in Tab. 8, our method outperforms its pre-train counterparts by a reasonable margin, which is in line with our previous experimental results.

6. Discussion and Conclusion

In this work, we propose the first versatile multi-modal pre-training framework HCMoCo specifically designed for human-centric perception tasks. Hierarchical contrastive learning targets are designed based on the nature of human datasets and the requirements of human-centric downstream tasks. Extensive experiments on four different human downstream tasks of different modalities demonstrated the effectiveness of our pre-training framework. We contribute a new RGB-D human parsing dataset NTURGBD-Parsing-4K to support research of human perception on RGB-D data. Besides downstream task transfer, we also propose two novel applications of HCMoCo to show its versatility and ability in cross-modal reasoning.

Potential Negative Impacts & Limitations. Usage of large amounts of data and long training time might negatively impact the environment. Moreover, even though we did not collect any new human data in this work, human data collection could happen if our framework is used in other applications, which potentially raises privacy concerns. As for the limitations, due to limited resources, we could only experiment with one possible instantiation of HCMoCo. And for the same reason, even though the theoretical possibility exists, we do not have the chance to further scale up the amount of human dataset and network size.

Acknowledgments This work is supported by NTU NAP, MOE AcRF Tier 2 (T2EP20221-0033), and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Jean-Baptiste Alayrac, Adria Recasens, Rosalia Schneider, Relja Arandjelovic, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *NeurIPS*, 2(6):7, 2020. [3](#), [6](#), [7](#)
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. [1](#), [2](#), [6](#)
- [3] Mahmoud Assran, Nicolas Ballas, Lluís Castrejon, and Michael Rabbat. Supervision accelerates pre-training in contrastive semi-supervised learning of visual representations. *arXiv preprint arXiv:2006.10803*, 2020. [2](#)
- [4] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. [2](#), [5](#)
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [3](#)
- [6] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [3](#)
- [7] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1971–1978, 2014. [1](#), [2](#), [5](#)
- [8] Yuxin Chen, Ziqi Zhang, Chunfeng Yuan, Bing Li, Ying Deng, and Weiming Hu. Channel-wise topology refinement graph convolution for skeleton-based action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13359–13368, 2021. [1](#)
- [9] MMPose Contributors. Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose>, 2020. [2](#), [5](#)
- [10] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. [2](#)
- [11] Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. Instance-level human parsing via part grouping network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 770–785, 2018. [1](#), [2](#), [3](#), [5](#)
- [12] Ke Gong, Xiaodan Liang, Dongyu Zhang, Xiaohui Shen, and Liang Lin. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 932–940, 2017. [1](#), [2](#), [3](#), [5](#)
- [13] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doversch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020. [3](#)
- [14] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018. [1](#), [2](#), [3](#), [6](#)
- [15] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *arXiv preprint arXiv:2010.09709*, 2020. [3](#)
- [16] Albert Haque, Boya Peng, Zelun Luo, Alexandre Alahi, Serena Yeung, and Li Fei-Fei. Towards viewpoint invariant 3d human pose estimation. In *European Conference on Computer Vision*, pages 160–177. Springer, 2016. [2](#), [6](#), [7](#)
- [17] Devamanyu Hazarika, Roger Zimmermann, and Soujanya Poria. Misa: Modality-invariant and-specific representations for multimodal sentiment analysis. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1122–1131, 2020. [3](#), [6](#)
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. [3](#)
- [19] Antonio Hernández-Vela, Nadezhda Zlateva, Alexander Marinov, Miguel Reyes, Petia Radeva, Dimo Dimov, and Sergio Escalera. Graph cuts optimization for multi-limb human segmentation in depth maps. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 726–732. IEEE, 2012. [2](#), [5](#)
- [20] Fangzhou Hong, Liang Pan, Zhongang Cai, and Ziwei Liu. Garment4d: Garment reconstruction from point cloud sequences. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. [2](#)
- [21] Ji Hou, Saining Xie, Benjamin Graham, Angela Dai, and Matthias Nießner. Pri3d: Can 3d priors help 2d representation learning? *arXiv preprint arXiv:2104.11225*, 2021. [2](#), [3](#), [6](#)
- [22] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014. [1](#), [2](#), [5](#), [6](#), [7](#)
- [23] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020. [2](#)
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. [6](#)
- [25] Jianshu Li, Jian Zhao, Yunchao Wei, Congyan Lang, Yidong Li, Terence Sim, Shuicheng Yan, and Jiashi Feng. Multiple-human parsing in the wild. *arXiv preprint arXiv:1705.07206*, 2017. [1](#), [2](#), [3](#), [5](#)

- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1, 2
- [27] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2684–2701, 2019. 1, 5, 6
- [28] Songtao Liu, Zeming Li, and Jian Sun. Self-emd: Self-supervised object detection without imagenet. *arXiv preprint arXiv:2011.13677*, 2020. 3
- [29] Yunze Liu, Qingnan Fan, Shanghang Zhang, Hao Dong, Thomas Funkhouser, and Li Yi. Contrastive multimodal fusion with tupleinforce. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 754–763, 2021. 3, 6
- [30] Yunze Liu, Li Yi, Shanghang Zhang, Qingnan Fan, Thomas Funkhouser, and Hao Dong. P4contrast: Contrastive learning with pairs of point-pixel pairs for rgb-d scene understanding. *arXiv preprint arXiv:2012.13089*, 2020. 2, 3
- [31] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2640–2649, 2017. 1, 2
- [32] Kaichiro Nishi and Jun Miura. Generation of human depth images with body part labels for complex human pose recognition. *Pattern Recognition*, 71:402–413, 2017. 2, 5
- [33] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 4, 7
- [34] Mandela Patrick, Yuki M Asano, Polina Kuznetsova, Ruth Fong, Joao F Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations. *arXiv preprint arXiv:2003.04298*, 2020. 3
- [35] Mandela Patrick, Yuki M Asano, Polina Kuznetsova, Ruth Fong, João F Henriques, Geoffrey Zweig, and Andrea Vedaldi. On compositions of transformations in contrastive self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9577–9587, 2021. 3
- [36] Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180. PMLR, 2019. 3
- [37] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 5
- [38] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017. 5, 8
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. 3
- [40] N Dinesh Reddy, Laurent Guigues, Leonid Pishchulin, Jayan Eledath, and Srinivasa G Narasimhan. Tesseract: End-to-end learnable multi-person articulated 3d pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15190–15200, 2021. 1, 2
- [41] Andrew Rouditchenko, Angie Boggust, David Harwath, Brian Chen, Dhiraj Joshi, Samuel Thomas, Kartik Adhikhasi, Hilde Kuehne, Rameswar Panda, Rogerio Feris, et al. Avlnet: Learning audio-visual language representations from instructional videos. *arXiv preprint arXiv:2006.09199*, 2020. 3
- [42] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019, 2016. 1, 2, 5
- [43] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5693–5703, 2019. 1, 2, 5, 6, 8
- [44] Feitong Tan, Danhang Tang, Mingsong Dou, Kaiwen Guo, Rohit Pandey, Cem Keskin, Ruofei Du, Deqing Sun, Sofien Bouaziz, Sean Fanello, et al. Humangps: Geodesic preserving feature for dense human correspondences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1820–1830, 2021. 2
- [45] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020. 2, 3, 4, 6, 7, 8
- [46] Longhui Wei, Lingxi Xie, Jianzhong He, Jianlong Chang, Xiaopeng Zhang, Wengang Zhou, Houqiang Li, and Qi Tian. Can semantic labels assist self-supervised visual representation learning? *arXiv preprint arXiv:2011.08621*, 2020. 2
- [47] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 3
- [48] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 466–481, 2018. 1, 2
- [49] Fu Xiong, Boshen Zhang, Yang Xiao, Zhiguo Cao, Taidong Yu, Joey Tianyi Zhou, and Junsong Yuan. A2j: Anchor-to-joint regression network for 3d articulated pose estimation from a single depth image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 793–802, 2019. 2
- [50] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action

recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018. 1

- [51] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3425–3435, 2019. 5