

Cross-view and Cross-pose Completion for 3D Human Understanding

Matthieu Armando
Vincent Leroy

Salma Galaaoui
Romain Brégier

Fabien Baradel
Philippe Weinzaepfel

Thomas Lucas
Grégory Rogez

NAVER LABS Europe

<https://europe.naverlabs.com/ComputerVision/CroCoMan>

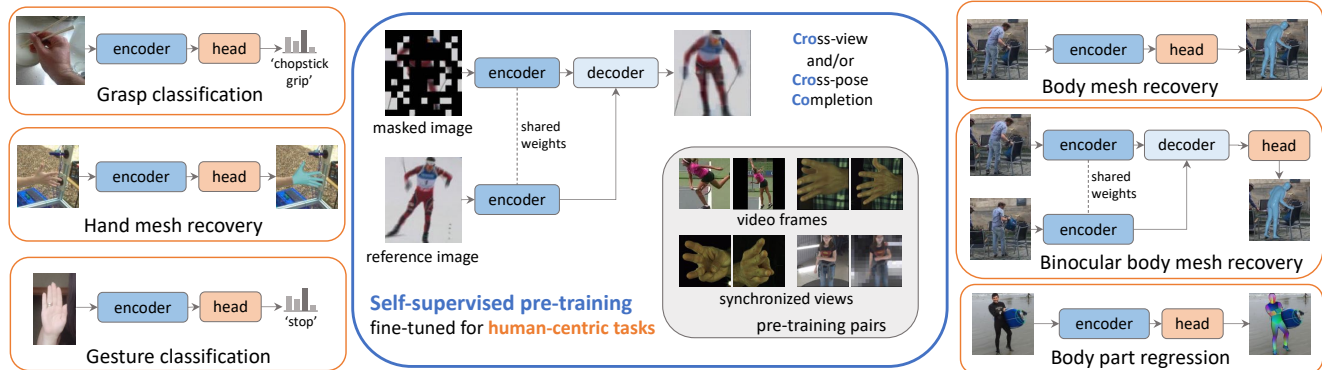


Figure 1. **Human-centric pre-training.** We pre-train a model for cross-view and cross-pose completion on body and hands image pairs (middle). This model serves as initialization for fine-tuning on several downstream tasks for both hands (left) and bodies (right). Our model, based on a generic transformer architecture, achieves competitive performance on these tasks without bells and whistles.

Abstract

Human perception and understanding is a major domain of computer vision which, like many other vision subdomains recently, stands to gain from the use of large models pre-trained on large datasets. We hypothesize that the most common pre-training strategy of relying on general purpose, object-centric image datasets such as ImageNet, is limited by an important domain shift. On the other hand, collecting domain-specific ground truth such as 2D or 3D labels does not scale well. Therefore, we propose a pre-training approach based on self-supervised learning that works on human-centric data using only images. Our method uses pairs of images of humans: the first is partially masked and the model is trained to reconstruct the masked parts given the visible ones and a second image. It relies on both stereoscopic (cross-view) pairs, and temporal (cross-pose) pairs taken from videos, in order to learn priors about 3D as well as human motion. We pre-train a model for body-centric tasks and one for hand-centric tasks. With a generic transformer architecture, these models outperform existing self-supervised pre-training methods on a wide set of human-centric downstream tasks, and obtain state-of-the-art performance for instance when fine-tuning for model-based and model-free human mesh recovery.

1. Introduction

The main catalyst of performance improvement in computer vision tasks in the last decade has arguably been the training of large models on large datasets [11, 23, 45, 49, 51, 64, 80]. For human-centric vision tasks, the standard approach is to pre-train models on ImageNet classification tasks and then fine-tune them on downstream tasks with specific datasets [31, 35, 55, 74]. This has at least three drawbacks: a) the size of the pre-training dataset is limited by label acquisition, b) there can be a large domain shift between ImageNet and downstream images, c) object-centric classification is different from human understanding, which may hinder the relevance of pre-trained features. Collecting large annotated datasets for human centric vision tasks is hard: target signal is costly and hard to acquire in the wild, e.g. relying on motion capture systems to obtain 3D pose.

To leverage large amounts of data and scale to large models, self-supervised pre-training methods such as contrastive learning [8, 10, 14, 30, 77] and masked signal modeling [24, 33] – have been developed. In these paradigms, a pre-text task is constructed from the data itself so that no manual labeling is required. The epitome of this philosophy is that of foundation models, such as GPT [7], trained on a large corpus of data with no manual annotation and fine-tuned to perform efficiently various downstream tasks. In computer vision, most self-supervised learning methods have




Figure 2. Examples of pre-training pairs taken from the different pre-training datasets. 📷 denote multi-view datasets, 🎬 video datasets and ◻ synthetic data.

been developed on ImageNet and evaluated chiefly on image classification tasks. Pioneering work in [79] provided a study of masked image modeling (MIM) pre-training on data dedicated to human-centric tasks, but found limited success. In contrast, we design a pre-training method to capture prior knowledge about the human body, which we use with human-centric data – see Figure 1 for an overview. It is inspired by masked image modeling [24, 33], where parts of an image are hidden and the network is trained to reconstruct them using the visible ones as input.

Unlike MAE [33] which operates on individual images, in our case pairs of images of human bodies are leveraged. These pairs are constructed in two ways: a) by taking two views of the same pose – plenty of those are available but the variability in pose and background is limited [57, 71, 81] – and b) by taking two poses in a motion sequence at different time steps – for instance from videos showing a person in movement, which are also plentiful [36, 54, 74, 82, 84]; see examples in Figure 2. Once a pair is constructed, parts of the first pose are masked, and the network is trained to reconstruct them from the unmasked regions as well as from the second image. This approach is close in spirit to cross-view completion (CroCo), proposed in [76], which performs masked image modeling with pairs of images. However, CroCo works with a different type of data (typically buildings or interiors [66]) with relatively small view-point changes. In their case, the objects are rigid, and therefore

CroCo only requires static pairs of the first type. To capture the deformable nature of the human body, we generalize their approach to the second type of pairs, — two different poses of a dynamic motion – and we refer to the corresponding task as cross-pose completion. We also include static pairs with extreme viewpoint changes, such as front and back views, to acquire a broad 3D understanding of the human body beyond stereo reconstruction. We propose a procedure to build suitable pairs of both types. These are taken from two human pose and motion modalities: views of full human bodies, and closeups of hands. This covers a wider input domain and captures information about human motion at two levels. Indeed, one of the specificity of general human-centric vision is the wide range of expected accuracy depending on the task. For instance, images of humans from afar may be sufficient for body pose estimation, but the millimetric accuracy expected for hand pose estimation requires close-ups of these parts.

Empirically, we pre-train a transformer architecture using the cross-pose and cross-view completion pretext tasks, on the two data modalities, *i.e.*, human bodies and hands. We fine-tune our model on the human-centric downstream tasks of pose and mesh recovery for bodies and hands [39, 44], and dense pose estimation [31, 58]. Using a generic transformer-based architecture, we demonstrate that our pre-training method transfers better than supervised ImageNet pre-training or existing self-supervised pre-training

methods when fine-tuning on downstream tasks. This allows our models to achieve competitive performance on these downstream tasks without requiring task-specific designs. We release two pre-trained models for human-centric vision: CroCo-Body and CroCo-Hand, specialized either for body or hand related tasks respectively. We demonstrate that these can easily be transferred to a variety of coarse- and fine-grained downstream tasks. For instance, our pre-training allows reaching state-of-the-art performance for model-based and model-free human mesh recovery.

2. Related work

We first discuss self-supervised learning in computer vision and then their application to human-centric tasks.

Self-supervised learning. Popular methods for self-supervised pre-training can be coarsely separated into two categories. The first one is based on instance discrimination [9, 10, 14, 30, 77] where various data augmentations are applied to an image to obtain different variants of it, and features extracted from the model are trained to be similar for these variants while being different from features extracted from other images. These methods produce models that transfer remarkably well to image-level tasks such as image classification and have achieved state-of-the-art performance on several vision benchmarks. However, recent studies suggest that the *object-centric* and the *balanced* nature of ImageNet [2, 59, 65] play a big role in this success, and indeed transfer performance degrades when tasks of a very different nature are considered [72].

The second paradigm relies on masked image modeling (MIM). Inspired by BERT [24] in natural language processing, MIM aims at reconstructing masked information from an input image either in pixel space [4, 5, 13, 25, 33, 78] or in feature space [3, 6, 75]. These methods have obtained some success on denser tasks such as object detection [33] and human pose estimation [79], which is of particular interest to our work. In [76], a MIM method that works with pairs of scene images has been proposed based on the cross-view completion task (CroCo). A second view of a scene is added to the MIM framework: this view can be used to improve reconstruction accuracy of the masked parts, provided that the model can exploit scene geometry. This approach allows the design of pretext tasks through a careful choice of image pairs. While CroCo is designed for static problems, we design a pre-training procedure intended to capture the articulated nature of the human body, by building pairs taking different poses from a dynamic motion, referred to as cross-pose completion, which we use alongside cross-view completion. MIM has also been applied on videos [26, 70]: information is also shared between frames, though in this case the goal is to learn from context rather than geometry.

Human-centric pre-training. A large percentage of images/videos captured in our daily life contains humans.

However collecting annotations related to human pose in an image is cumbersome, especially when 3D annotations are required [32, 74], and in practice supervised data is scarce. This motivates the development of self-supervised methods tailored to human-centric data. In recent works [50, 79, 83], existing self-supervised learning strategies were adapted to human-centric data. A popular approach is to leverage the fact that humans have a well-defined structure, for instance by learning human pose in an unsupervised manner [37, 38, 61, 67]. Jakab *et al.* [37] extract 2D landmark of faces and bodies in an unsupervised way by learning to disentangle pose from appearance through conditional image generation from image pairs. However, the learned 2D landmarks do not have structure and do not correspond to any predefined human skeleton. In their following work [38], this issue was solved by leveraging Motion Capture (MoCap) data to ensure that the 2D landmarks are matching a certain pose distribution. However, this is quite limiting, especially in the presence of rare or unseen poses. Following a different pre-training strategy, HCMoCo [34] learn a representation on multimodal data including RGB, depth or 2D keypoints, using contrastive learning. It obtains convincing results on multiple downstream tasks such as part segmentation or 3D pose estimation. However, acquiring such multimodal data outside of studio environments is difficult and prevents the use of this method in the wild. UniHCP [21] and PATH [69] followed this direction by both proposing to learn a model using multitask training. They required annotations on several tasks however, and mainly focused on 2D downstream tasks. Most recently, LiftedCL [16] leverages contrastive learning to obtain a representation invariant to data augmentations. They encourage the pose representation to be realistic via the use of an adversarial strategy on 3D skeleton data. More recently, SOLIDER [15] focused on disentangling semantics and appearances in an unsupervised manner. They generate pseudo semantic labels using DINO [10]. While reaching good performances on 2D tasks such as human parsing or attribute recognition, they do not present results on 3D tasks.

Contrary to these, we do not impose any constraint on the learned human representation. We instead cast our pre-training as a conditional image generation where image pairs are selected from different viewpoints or time steps.

3. Method

We first present our pre-training objective (Section 3.1), then our data collection process to collect cross-view (Section 3.2) and cross-pose (Section 3.3) data, and finally the fine-tuning on downstream tasks (Section 3.4).

3.1. Multi-view masked image modeling

Masked image modeling as proposed in [33] proceeds by dividing an image x into N non-overlapping patches $p =$

$\{\mathbf{p}^1, \dots, \mathbf{p}^N\}$. A subset of $n = \lfloor rN \rfloor$ tokens (*i.e.*, encoded patches) is randomly masked, with $r \in [0, 1]$ an hyper-parameter controlling the masking ratio, and the remaining set of visible patches is denoted $\tilde{\mathbf{p}} = \{\mathbf{p}^i | m_i = 0\}$, with $m_i = 0$ indicating that the patch \mathbf{p}_1^i is not masked ($m_i = 1$ otherwise). A model is trained to predict masked tokens given the visible ones. While there is an inherent ambiguity on the masked content, empirically [33] has shown that such an approach of minimizing a deterministic L_2 loss between prediction and target is a successful pre-training objective.

The framework proposed in CroCo [76] extends this concept to pairs of images. In our case, we build two types of image pairs: a) cross-pose; where the pair comes from two different time-steps but the same camera or – b) cross-view; where the pair depicts the same pose of a person seen from two different viewpoints, see Figure 2 for examples of pairs. Let x_t^v be an image viewed from viewpoint v and taken at timestep t . Images in a pair are divided into sets of patches, denoted $(\mathbf{p}_t^v, \mathbf{p}_t^w)$, patches from the target are randomly masked, let $\tilde{\mathbf{p}}_t^v$ denote the set of remaining patches, both views are encoded separately with an encoder E_θ sharing the weights between the two images, and decoded together with a decoder D_ϕ :

$$\hat{\mathbf{p}}_{t,\phi}^v = D_\phi(E_\theta(\tilde{\mathbf{p}}_t^v); E_\theta(\mathbf{p}_t^w)). \quad (1)$$

While CroCo considered *static* pairs of images, denoted $(\mathbf{x}^v, \mathbf{x}^w)$ in this work, the human body is a deformable object; therefore we also consider *dynamic* pairs, representing different poses taken at different steps of a human motion, denoted $(\mathbf{x}_t, \mathbf{x}_{t'})$. We denote $\mathcal{D}_{\text{pose}}$ the set of cross-pose (dynamic) pairs and $\mathcal{D}_{\text{view}}$ the set of cross-view (static) image pairs. The model is trained by optimizing:

$$\min_{\theta, \phi} \underbrace{\sum_{(\mathbf{p}_t, \mathbf{p}_{t'}) \in \mathcal{D}_{\text{pose}}} \|\hat{\mathbf{p}}_{t,\phi} - \mathbf{p}_t\|^2}_{\text{(a) cross-pose completion}} + \underbrace{\sum_{(\mathbf{p}^v, \mathbf{p}^w) \in \mathcal{D}_{\text{view}}} \|\hat{\mathbf{p}}_{\theta, \phi}^v - \mathbf{p}^v\|^2}_{\text{(b) cross-view completion}}. \quad (2)$$

In the case where only (b) is used, the objective boils down to the same as in [76], with different data.

Human-centric masking strategy. Following standard MIM practice, we mask a subset of whole tokens (*i.e.*, squared patches in the image). While standard MIM methods mask each token with equal probability, we leverage the prior knowledge that pixels belonging to the humans are the relevant ones. Therefore, when segmentations can be easily obtained, we explore non-uniform masking strategies. Specifically, the set of N image patches is separated into N_H *human* and N_B *background* patches ($N = N_H + N_B$). The masking ratio is applied to N_H only, *i.e.* we keep visible $n = N_H - \lfloor rN_H \rfloor$ patches. Since N_H varies between images, so does n : for each training batch, we add to the visible set of patches randomly chosen background patches, up to a fixed sequence length set to the maximum of N_H over

the batch. This has two consequences: a) token sequences of fixed length can be used as input within a batch, which speeds up training and b) the model sees distracting background patches in the first input. Experiments on masking ratio (ranging from 60% to 90%) did not yield any significant downstream performance difference. We thus use a 75% masking ratio, in line with [33].

We now detail how to obtain pairs of images of full-bodied human data suitable for pre-training. To generate pairs for hands, we use the InterHand2.6M [57] dataset, from which we sample both cross-view and cross-pose pairs in equal number.

3.2. Cross-view pair construction

While CroCo [76] required pairs with rather small camera baselines due to unconstrained scenes, in our case, we work with the strong assumption of a known object – a person – and therefore any viewpoint change is admissible. We rely on existing multi-view datasets for two data modalities, namely human *bodies* and *hands*, where subjects are captured from multiple viewpoints by synchronous cameras, selected to have diversity in identities, appearances, poses and backgrounds. We use the HUMBI [81] and AIST [71] datasets as well as synthetic data. HUMBI contains more than 300 subjects, with a wide range of age, body-shapes, and clothing variety but a restricted set of poses, while AIST sequences are captured from only 40 subjects with 9 different camera viewpoints, with plain clothing, but contain a great diversity of poses, with *e.g.* dance moves from professional performers. We gather 1510 sequences with more than $5M$ images in total.

Obtaining information for masking. For HUMBI, we run an off-the-shelf human parser, the PGN method from CIHP [29], to generate silhouette information needed for human masking. Note that only a rough estimate of the silhouette is needed for patch-level-guided masking.

Increased diversity with synthetic data. Taken together, HUMBI and AIST contain a wide variety of identities and poses, but lack diversity in terms of environment, as they were captured in an indoor studio setting. To remedy this, we create a synthetic dataset where we generate multiple-view renderings of SMPL meshes following SUR-REAL [73], with diverse lighting conditions and camera viewpoint, overlaid on top of distracting background images from COCO [49] or rendered inside the Habitat simulator [66]. Body parameters and poses are sampled from AMASS [53] to generate images with fully-visible persons.

3.3. Cross-pose pairs construction

As the human body is non-rigid, going beyond the static setting proposed in CroCo [76] can enable the model to gain some understanding of how body-parts interact and move w.r.t. one another. Dynamic pairs can be constructed from

multi-view data		video data	Body Mesh Recovery		DensePose	
📷	📷	📷	PA-MPJPE↓	MPVPE↓	AUC ₃₀ ↑	AUC ₁₀ ↑
✓			58.6	112.3	0.655	0.379
	✓		58.1	111.8	0.645	0.367
✓	✓		58.4	111.4	0.651	0.372
✓		✓	57.5	111.5	0.692	0.416

(a) **Pre-training data for CroCo-body.** 📷 denotes cross-pose images and 📷 indicates cross-view images.

multi-view data		Hand Mesh Recovery		Gesture
📷	📷	PA-MPJPE↓	MPVPE↓	Acc(%)↑
✓		7.97	16.02	97.48
	✓	8.40	16.96	97.13
✓	✓	8.21	16.69	97.36

(b) **Pre-training data for CroCo-Hand.** 📷 denotes cross-pose images and 📷 indicates cross-view images.

Init.	Body Mesh Recovery		DensePose	
	PA-MPJPE↓	MPVPE↓	AUC ₃₀ ↑	AUC ₁₀ ↑
CroCo [76]	57.5	111.5	0.692	0.416
Random	60.5	117.5	0.620	0.340

(c) **Weights initialization for CroCo-Body.** Starting from pre-trained CroCo weights performs better than a random initialization.

Masking	Body Mesh Recovery		DensePose	
	PA-MPJPE↓	MPVPE↓	AUC ₃₀ ↑	AUC ₁₀ ↑
Human	57.5	111.5	0.692	0.416
Uniform	57.7	112.0	0.654	0.374

(d) **Masking strategy.** Performance on downstream tasks improves when the masking is focused on patches where humans are visible.

Init.	Hand Mesh Recovery		Gesture
	PA-MPJPE↓	MPVPE↓	Acc(%)↑
CroCo [76]	8.21	16.69	97.36
Random	8.63	17.97	97.18

(e) **Weights initialization for CroCo-Hand.** Hand mesh recovery also benefits from starting from pre-trained CroCo weights.

Table 1. **Ablations for CroCo-Body (left) and CroCo-Hand (right) pre-training strategies.** Default settings are highlighted in grey.

a monocular RGB video capturing a human motion. Such data typically also provide more variety in terms of appearance and background, as they can be captured from a single camera, which is easier to collect in the wild. We tested various strategies for temporal sampling, using predefined time intervals. No significant difference was observed on downstream performance. We thus keep a simple random sampling strategy.

Extracting pairs from video datasets. We rely on a mix of diverse human-centric datasets namely 3DPW [74], Posetrack [36], PennAction [82], JRDB [54], MARS [84] and AIST [71]. 3DPW includes video footage taken from a moving phone camera of persons performing common actions such as walking or playing guitar. We use videos from the training set only, leading to more than 22k frames from 34 sequences. Posetrack is a large-scale dataset for multi-person pose estimation and tracking in videos that contains more than 97k images from 3719 sequences. PennAction contains 2326 video sequences of 15 different actions. JRDB contains 310k images from 2509 sequences collected with a social mobile manipulator, JackRabbit. MARS is a CCTV dataset containing video tracks of more than 1200 identities and 20k video sequences.

3.4. Fine-tuning on downstream tasks

Our model can be fine-tuned on a variety of downstream tasks which can be either image-level (*e.g.* as mesh recovery) or dense, *i.e.*, that require pixelwise predictions (*e.g.* dense vertex coordinates regression). Interestingly, our method can tackle monocular and/or binocular tasks. For monocular task, the ViT-Base encoder is used alone, while binocular tasks benefit from the pre-training of both the encoder and the decoder. For fine-tuning, we replace the linear head from the pre-training model with a randomly initialized prediction head. This prediction head can be of two

types depending on the aforementioned type of task. For image-level tasks, we perform an average pooling on the output tokens from the last transformer block and we use an MLP to regress or classify. For dense tasks, we leverage DPT [60], that assembles tokens from various stages into a dense output with a prediction for each pixel.

4. Experiments


We evaluate our pre-training strategies on monocular human-centric tasks listed in Section 4.1. We then provide extensive ablations (Section 4.2) and comparison with prior works (Section 4.3). Finally we show results in binocular settings in Section 4.4. We provide qualitative examples for several downstream tasks in the supplementary material.

4.1. Downstream tasks

Body mesh recovery. We evaluate CroCo-Body on the body mesh recovery task on the 3DPW test set [74] composed of 35k human crops following prior works [39, 40, 44]. We use different type of training data for fine-tuning: a) COCO-Part-EFT [39] which consists of 28k images from COCO with pseudo-ground-truth of SMPL parameters [52], b) Mix: a mix of 2D/3D datasets commonly used by prior works [39, 40, 44] composed of COCO-All-EFT [39], MPII [1], Human3.6M [35] and MPI-INF-3DHP [55]), and c) Mix*: the mix of datasets described in b) plus 3DPW train set when comparing against the state of the art. When not specified, we use COCO-Part-EFT for fine-tuning on this task. We experiment with the two main types of heads used in the literature: a) a model-based method using a vanilla ‘iterative regressor’ [39, 40, 44] prediction head, which consists in regressing the SMPL parameters with an MLP in an iterative manner starting from the mean parameters, and b) a model-free method which directly regresses the 3D location of the vertices using the METRO predic-

Initialization	Body Mesh Recovery		Hand Mesh Recovery		DensePose	
	PA-MPJPE↓	MPVPE↓	PA-MPJPE↓	MPVPE↓	AUC_{30} ↑	AUC_{10} ↑
Random	80.2	180.0	12.92	40.06	0.430	0.206
CroCo [76]	61.5	119.0	9.34	19.24	0.590	0.320
MAE-Body/Hand	59.1	115.1	8.75	17.93	0.625	0.351
MAE-IN1k [33]	59.6	113.0	8.64	16.84	0.678	0.392
CroCo-Body/Hand	57.4	111.5	7.97	16.02	0.692	0.416

(a) Comparison with other pre-training strategies



(b) Data efficiency

Figure 3. **Comparison with other pre-training methods** on different downstream tasks (a) or under different fine-tuning data regimes (b), *i.e.*, when varying the number of annotated training samples from COCO_{part} for fine-tuning on the body mesh recovery task from 10% to 100%. MAE-Body/Hand means that we pre-train MAE on the same data as CroCo-Body/Hand.

tion head [47]. We report standard metrics: MPJPE (Mean Per Joint Error), MPVPE (Mean Per Vertex Error) and PA-MPJPE (Procrustes-Aligned MPJPE) in millimeters.

DensePose. We also evaluate on the DensePose task [31], which consists in mapping every human pixel in an image to a point on the surface of the body. It involves dense classification (of body parts) and dense regression (UV coordinates). Our output and losses are similar to DensePose: we learn a 25-way classification unit, and 24 2D regression functions. We do not target detection, so we train and evaluate on single-person human-centered crops only. For evaluation, we compute the ratio of correct points $RCP(t)$ among annotated pixels, where a correspondence is considered correct if the geodesic distance (on the surface) between estimated pixel value and annotation is below a threshold t . Next, taking all annotated points on the evaluation set, we consider the area under the curve (AUC), defined as $AUC_a = \frac{1}{a} \int_0^a RCP(t) dt$, for $a = 30$ and $a = 10$.

Hand mesh recovery. We fine-tune CroCo-Hand on the task of hand mesh recovery. The setup is similar to the body mesh recovery task with the MANO parametric model [63]. We only use the vanilla ‘iterative regressor’ head [39, 40, 44] and we report results on the FreiHand dataset [86].

Hand gesture classification. We fine-tune CroCo-Hand on HaGRID [41], which contains 18 gesture classes with 509,323 images for training and 43,669 for testing. Using the full training set typically leads to very high accuracies (*e.g.* above 99%), likely thanks to the large amount of annotated samples, and in this regime, unsupervised pre-training is thus unnecessary. In the regime where little data is available, however, pre-training is expected to have a noticeable impact. We thus report few-shot accuracy considering 64 samples per class except otherwise stated.

Grasp classification. We fine-tune CroCo-Hand on the Grasp Understanding (GUN-71) dataset [62], which consists of 12k first-person RGB images annotated with 71 hand-object manipulation classes. We follow the evaluation protocol of [62] on the 33 grasp taxonomy defined in [27] with the ‘Best View’ and ‘All’ settings. This dataset is of particular interest as the classes are fine-grained, with some highly-specific grasps.

4.2. Ablation studies

Pair construction. In Table 1 (a) and (b), we construct cross-view and cross-pose pairs only from datasets that can be used to create *both* types of pairs – *i.e.*, for body-centric data, multi-view datasets that also include some temporal sampling of motion, namely HUMBI, AIST and synthetic data, InterHand for hand-centric data. This way, the pair construction method is the only difference and the data used is the same across variants and results are not influenced by the use of additional data. Overall, the best results are obtained with cross-view pairs, and it seems they should be favored when possible. However, the diversity of multi-view data available is limited, and enriching the training data mix with cross-pose pairs from monocular videos leads to a clear performance boost, see last row of Table 1 (a).

Initialization. In Table 1 (c) and (e) we start the training of CroCo-Body and CroCo-Hand from random weights, instead of weights initialized from CroCo pre-training (default setting), and train for the same number of iterations. We observe that this decreases the results on all tasks.

Masking. In Table 1 (d), we investigate the importance of the masking strategy in our pre-training. We compare our human-centric masking strategy to a naive *uniform* masking strategy which consists in randomly sampling tokens from the image. Performance degrades across tasks when using the uniform sampling: we conclude that masking pixels that belong to the person with higher probability is beneficial to learning human-centric features.

Comparison with other pre-training strategies. In Figure 3 (a), we compare our pre-trained CroCo-Body and CroCo-Hand models with the publicly released CroCo [76] and MAE [33] weights. In order to evaluate the respective contributions of pre-training data and pre-training method, we also pre-train MAE on our data. We refer to these models as MAE-Body/MAE-Hand and refer to the original MAE model as MAE-IN1k. Our pre-training (bottom row) outperforms both the CroCo and MAE-IN1k models. Interestingly, MAE-Body/Hand performs worse than MAE-IN1k, thus the gains of our models come from the pre-training objective rather than from the difference in pre-training data. This is partly unexpected, as our pre-training

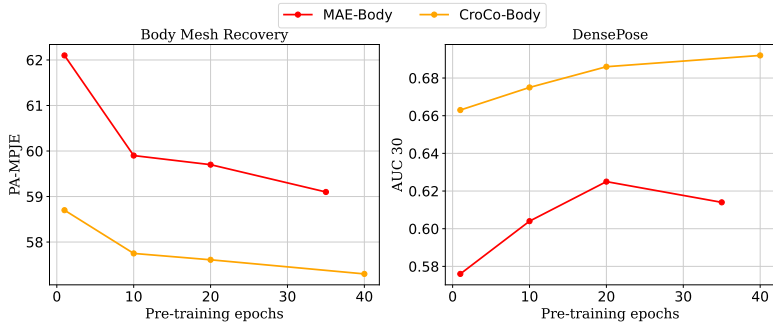


Figure 4. **Impact of the number of pre-training epochs.** CroCo-Body is initialized from CroCo while MAE is initialized from ImageNet.

Pre-training supervision	Body Mesh Recovery	
	PA-MPJPE ↓	MPVPE ↓
2D keypoints (supervised)	54.6	116.7
CroCo-Body (self-supervised)	53.8	112.8
CroCo-Body + 2D keypoints	52.8	111.5

Table 3. **Ablation on pretext task used for pre-training.** All models are pre-trained only on datasets for which we have keypoints annotations, namely 3DPW, PoseTrack, PennAction, AIST, and JRDB. We use `Mix` for fine-tuning.

data is more tailored to the downstream tasks, but may be explained by the fact that ImageNet-1k offers a vast diversity of backgrounds and texture, which may be beneficial for pre-training. Figure 4 shows the evolution of performance on downstream tasks after a varying number of epochs; it shows that CroCo-Body starts off from a better performance, and converges fast compared to MAE-Body.

In Table 2, we compare our CroCo-Hand to MAE-Hand and a random initialization on the hand classification tasks (gesture and grasp). In all cases, we outperform MAE-Hand, which itself outperforms a random initialization.

Fine-tuning data efficiency. One motivation for pre-training is that the fine-tuning stage on a specific task requires less annotated training samples to reach good performance. In Figure 3 (b), we report the performance on body tasks of different pre-training strategies and their efficiency under different fine-tuning data regimes. CroCo-Body achieves better performance compared to other pre-training strategies (MAE-IN1K and Random) with the same number of training samples: it reaches the same performance as MAE-Body with 2 to 3 times less training samples. In particular, CroCo-Body obtains a PA-MPJPE of 65mm with 10% of the training set while MAE-IN1K requires 25% of the train set to reach a similar performance. Table 2 reports the results for grasp recognition and for few-shot gesture classification with 16, 32 and 64 samples per class for a network initialized randomly, with MAE-Hand or with CroCo-Hand. CroCo-Hand outperforms MAE-Hand which itself obtains significantly higher accuracies than a random initialization. Interestingly, CroCo-Hand obtains

Initialization	Gesture↑			Grasp↑	
	16	32	64	BV	All
Random	45.7	68.4	81.2	10.9	8.6
MAE-Hand	85.1	92.7	95.9	33.2	25.8
CroCo-Hand	92.5	95.8	97.4	37.8	29.2

Table 2. **Impact of pre-training method for gesture and grasp classification.** For gesture, we report the accuracy (%) when varying the number of samples per class used for training. For grasp, we report both the ‘Best View’ (BV) and ‘All’ protocols from [62].

above 92% accuracy even with only 16 shots, while MAE-Hand requires 32 samples per class to reach such accuracy.

Keypoints supervision. We evaluate how the proposed self-supervised pre-training strategy compares to a fully-supervised pre-training. 2D keypoints annotations are available for multiple human-centric datasets, and we thus pre-train a network in a fully-supervised manner to regress body keypoints as dense 2D heatmaps, following usual practice in the field [17, 68]. We fine-tune this model on human-related downstream tasks, and compare performance with the proposed CroCo-Body pre-training on the same data. CroCo-Body pre-training led to better performances than the supervised approach for all downstream tasks considered (see Table 3), confirming the interest of the proposed approach. Yet, by combining 2D keypoints supervision to cross-view completion during pre-training, we were able to obtain a small additional performance boost.

4.3. Comparison to the state of the art

We now compare our performance against state-of-the-art methods, reported in Table 4 and 5.

On the task of Body Mesh Recovery (Table 4 (a)), CroCo-Body achieves state-of-the-art performances using either a model-based or model-free head. CroCo-Body yields consistent gains across choices of head, which indicates that our proposed pre-training stage learns highly transferable features for the task of mesh recovery.

We report in Table 4 (b) the results on the DensePose task. We follow the setup of [31] and observe that CroCo-Body achieves better performance on all metrics.

Next, we report the results on hand mesh recovery in Table 5 (a). CroCo-Hand gets results on par with prior work. Recent methods leverage more complicated prediction heads (see Section 4.1), however, we show that without bells and whistles, we perform on par with these methods.

Finally, we report grasp classification in Table 5 (b). CroCo-Hand is competitive with prior works [18, 22, 62] but does not reach state-of-the-art performances. This can be explained by the extremely limited variety of back-

Method		PA-MPJPE↓	MPJPE↓
model-based	HybrIK [46]	48.8	80.0
	ROMP [68]	47.3	76.7
	PARE [43]	46.5	74.5
	Choi <i>et al.</i> [20]	56.8	-
	CroCo-Body	46.2	<u>76.2</u>
model-free	Pose2Mesh [19]	56.3	89.5
	METRO [47]	47.9	77.1
	MeshGraphormer [48]	45.6	74.7
	FastMETRO [17]	<u>44.6</u>	<u>73.5</u>
	PointHMR [42]	44.9	73.9
	CroCo-Body	44.2	73.3

(a) Body mesh recovery on 3DPW [74] after fine-tuning on Mix*

Method	AUC ₁₀ ↑	AUC ₃₀ ↑
DensePose [31]	0.378	0.614
CroCo-Body	0.416	0.692

(b) DensePose estimation on COCO [31]

Table 4. **CroCo-Body vs. the state of the art.** Using a generic architecture, we achieve competitive performances with recent works thanks to the proposed pre-training.

Method	PA-MPJPE↓	PA-MPVPE↓
PIXIE hand expert [28]	12.0	12.1
MANO-CNN [86]	11.0	10.9
Pose2Mesh [19]	7.7	7.8
Hand4Whole [56]	7.7	<u>7.7</u>
METRO [47]	6.8	6.7
CroCo-Hand	<u>7.4</u>	8.0

(a) Hand mesh recovery on FreiHand [86]

Method	Best View↑	All↑
Rogez <i>et al.</i> [62]	22.7	20.5
Coskun <i>et al.</i> [22]	<u>42.3</u>	N/A
Choi <i>et al.</i> [18]	N/A	<u>39.8</u>
CroCo-Hand	37.8	29.2
CroCo-Hand[†]	49.9	49.6

(b) Grasp classification on GUN-71 [62]

Table 5. **CroCo-Hand vs. the state of the art.** CroCo-Hand[†] represents our method pre-trained with more diverse data.

grounds and the absence of manipulated objects in the pre-training dataset [57]. Pre-training on diverse images seems to be critical for reaching high performance on GUN-71, as suggested by the results of Coskun *et al.* [22] who pre-trained their network on ImageNet. To solve this issue, we pre-train CroCo-Hand using cross-view and cross-pose pairs from more datasets containing hand-object interactions and more diverse backgrounds (InterHand2.6M [57], DexYCB [12], HO3D [32], HanCo [85]). This variant,

Input	Pre-training	Hand Mesh Recovery		Body Mesh Recovery	
		PA-MPJPE↓	PA-MPVPE↓	PA-MPJPE↓	PA-MPVPE↓
Monocular	MAE-IN1k	11.6	11.2	58.9	113.0
	CroCo-Body/Hand	11.0	10.8	57.5	111.2
Binocular	MAE-IN1k	9.9	9.5	60.5	115.0
	CroCo-Body/Hand	9.5	8.8	55.0	108.4

Table 6. **Monocular vs. binocular mesh recovery** performance on the 3DPW test set (body) and on HanCo (hand).

denoted as CroCo-Hand[†], leads to a clear improvement, reaching state-of-the-art accuracy.

4.4. Extension to binocular tasks

One interesting property of our CroCo-Body and CroCo-Hand architecture is that the decoder can also be leveraged in binocular settings. In Table 6 we report performance of CroCo-Body and CroCo-Hand using either a single-image input with the encoder only, or using two inputs images and leveraging both the encoder and decoder. For the body mesh recovery task, we fine-tune on COCO-Part and the 3DPW train set as well since 3DPW contains sequences, so we can feed neighboring frames to our model as well. For hand mesh recovery, we fine-tune on Hanco [85], a multi-view extension of FreiHand. Results for the texture estimation task are reported in the supplementary material.

Using both the pre-trained encoder and decoder for fine-tuning in a binocular framework proves to be beneficial for both body and hand mesh recovery. For sanity’s sake, we also fine-tune MAE-IN1k in binocular settings by initializing the decoder from scratch. While it degrades the performance on the body mesh recovery task compared to its monocular counterpart, we observe that it gets reasonable performance on the hand mesh recovery task. This can be explained by the fact that the HanCo training set is large compared to the training set used for the body mesh recovery task so the pre-training has less of an impact. CroCo-Body/Hand still obtains better performance than MAE-IN1k in both monocular and binocular settings.

5. Conclusion

We present a strategy to pre-train a model for various human-centric tasks. Our approach is based on a self-supervised pre-training method that leverages cross-view completion (CroCo), using pairs of images of people taken from different viewpoints or in different poses. Fine-tuning a model pre-trained with this strategy leads to state-of-the-art performance on various tasks, such as human mesh recovery and competitive performance on gesture classification, while using a generic model architecture. The proposed pre-training outperforms other popular pre-training strategies such as MAE [33]. Lastly, we show how this pre-training can be leveraged to perform binocular tasks such as body mesh recovery from multiple views.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014. 5
- [2] Mahmoud Assran, Randall Balestriero, Quentin Duval, Florian Bordes, Ishan The hidden uniform cluster prior in self-supervised learning. In *ICLR*, 2023. 3
- [3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Michael Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *ECCV*, 2022. 3
- [4] Sara Atito, Muhammad Awais, and Josef Kittler. SiT: Self-supervised vision Transformer. *arXiv preprint arXiv:2104.03602*, 2021. 3
- [5] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. MultiMAE: Multi-modal Multi-task Masked Autoencoders. In *ECCV*, 2022. 3
- [6] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language. *arXiv preprint arXiv:2202.03555*, 2022. 3
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, 2020. 1
- [8] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 1
- [9] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In *NeurIPS*, 2020. 3
- [10] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *ICCV*, 2021. 1, 3
- [11] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. 1
- [12] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S. Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, Jan Kautz, and Dieter Fox. Dexycb: A benchmark for capturing hand grasping of objects. In *CVPR*, 2021. 8
- [13] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative Pretraining From Pixels. In *ICML*, 2020. 3
- [14] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 1, 3
- [15] Weihua Chen, Xianzhe Xu, Jian Jia, Hao Luo, Yaohua Wang, Fan Wang, Rong Jin, and Xiuyu Sun. Beyond appearance: a semantic controllable self-supervised learning framework for human-centric visual tasks. In *CVPR*, 2023. 3
- [16] Ziwei Chen, Qiang Li, Xiaofeng Wang, and Wankou Yang. Liftedcl: Lifting contrastive learning for human-centric perception. In *ICLR*, 2023. 3
- [17] Junhyeong Cho, Kim Youwang, and Tae-Hyun Oh. Cross-attention of disentangled modalities for 3d human mesh recovery with transformers. In *ECCV*, 2022. 7, 8
- [18] Chiho Choi, Sang Ho Yoon, Chin-Ning Chen, and Karthik Ramani. Robust hand pose estimation during the interaction with an unknown object. In *ICCV*, 2017. 7, 8
- [19] Hongsuk Choi, Gyeongsik Moon, and Kyoung Mu Lee. Pose2mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *ECCV*, 2020. 8
- [20] Hongsuk Choi, Hyeongjin Nam, Taeryung Lee, Gyeongsik Moon, and Kyoung Mu Lee. Rethinking self-supervised visual representation learning in pre-training for 3d human pose and shape estimation. *ICLR*, 2023. 8
- [21] Yuanzheng Ci, Yizhou Wang, Meilin Chen, Shixiang Tang, Lei Bai, Feng Zhu, Rui Zhao, Fengwei Yu, Donglian Qi, and Wanli Ouyang. Unihcp: A unified model for human-centric perceptions. In *CVPR*, 2023. 3
- [22] Huseyin Coskun, M Zeeshan Zia, Bugra Tekin, Federica Bogo, Nassir Navab, Federico Tombari, and Harpreet Sawhney. Domain-specific priors and meta learning for few-shot first-person action recognition. *IEEE Trans. PAMI*, 2021. 7, 8
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1
- [24] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL HLT*, 2019. 1, 2, 3
- [25] Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jégou, and Edouard Grave. Are Large-scale Datasets Necessary for Self-Supervised Pre-training? *arXiv preprint arXiv:2112.10740*, 2021. 3
- [26] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. In *NeurIPS*, 2022. 3
- [27] T. Feix, R. Pawlik, H. Schmiemayer, J. Romero, and D. Kragic. A comprehensive grasp taxonomy. In

Robotics, Science and Systems: Workshop on Understanding the Human Hand for Advancing Robotic Manipulation, 2009. 6

- [28] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Collaborative regression of expressive bodies using moderation. In *3DV*, 2021. 8
- [29] Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. Instance-level human parsing via part grouping network. In *ECCV*, 2018. 4
- [30] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In *NeurIPS*, 2020. 1, 3
- [31] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *CVPR*, 2018. 1, 2, 6, 7, 8
- [32] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *CVPR*, 2020. 3, 8
- [33] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked Autoencoders are Scalable Vision Learners. In *CVPR*, 2022. 1, 2, 3, 4, 6, 8
- [34] Fangzhou Hong, Liang Pan, Zhongang Cai, and Ziwei Liu. Versatile multi-modal pre-training for human-centric perception. In *CVPR*, 2022. 3
- [35] 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 Trans. PAMI*, 2013. 1, 5
- [36] Umar Iqbal, Anton Milan, and Juergen Gall. Pose-track: Joint multi-person pose estimation and tracking. In *CVPR*, 2017. 2, 5
- [37] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. *NeurIPS*, 2018. 3
- [38] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Self-supervised learning of interpretable keypoints from unlabelled videos. In *CVPR*, 2020. 3
- [39] Hanbyul Joo, Natalia Neverova, and Andrea Vedaldi. Exemplar fine-tuning for 3d human model fitting towards in-the-wild 3d human pose estimation. In *3DV*, 2021. 2, 5, 6
- [40] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 5, 6
- [41] Alexander Kapitanov, Andrey Makhlyarchuk, and Katarina Kvanchiani. Hagrid - hand gesture recognition image dataset. *arXiv preprint arXiv:2206.08219*, 2022. 6
- [42] Jeonghwan Kim, Mi-Gyeong Gwon, Hyunwoo Park, Hyukmin Kwon, Gi-Mun Um, and Wonjun Kim. Sampling is Matter: Point-guided 3d human mesh reconstruction. In *CVPR*, 2023. 8
- [43] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *CVPR*, 2021. 8
- [44] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 2, 5, 6
- [45] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 1
- [46] Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *CVPR*, 2021. 8
- [47] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *CVPR*, 2021. 6, 8
- [48] Kevin Lin, Lijuan Wang, and Zicheng Liu. Mesh graphormer. In *ICCV*, 2021. 8
- [49] 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 *ECCV*, 2014. 1, 4
- [50] Zhenguang Liu, Runyang Feng, Haoming Chen, Shuang Wu, Yixing Gao, Yunjun Gao, and Xiang Wang. Temporal feature alignment and mutual information maximization for video-based human pose estimation. In *CVPR*, 2022. 3
- [51] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. 2015. 1
- [52] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Trans. Graphics*, 2015. 5
- [53] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. Amass: Archive of motion capture as surface shapes. In *ICCV*, 2019. 4
- [54] Roberto Martin-Martin, Mihir Patel, Hamid Rezatofighi, Abhijeet Sheno, JunYoung Gwak, Eric Frankel, Amir Sadeghian, and Silvio Savarese. Jrd: A dataset and benchmark of egocentric robot

- visual perception of humans in built environments. *IEEE Trans. PAMI*, 2021. 2, 5
- [55] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3DV*, 2017. 1, 5
- [56] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Accurate 3d hand pose estimation for whole-body 3d human mesh estimation. In *CVPR*, 2022. 8
- [57] Gyeongsik Moon, Shoou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. Interhand2.6m: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *ECCV*, 2020. 2, 4, 8
- [58] Natalia Neverova, David Novotny, Marc Szafraniec, Vasil Khalidov, Patrick Labatut, and Andrea Vedaldi. Continuous surface embeddings. In *NeurIPS*, 2020. 2
- [59] Senthil Purushwalkam and Abhinav Gupta. Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. In *NeurIPS*, 2020. 3
- [60] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 5
- [61] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised geometry-aware representation for 3d human pose estimation. In *ECCV*, 2018. 3
- [62] Grégory Rogez, James S Supancic, and Deva Ramanan. Understanding everyday hands in action from rgb-d images. In *ICCV*, 2015. 6, 7, 8
- [63] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Trans. Graphics*, 2017. 6
- [64] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Dex: Deep expectation of apparent age from a single image. In *ICCVW*, 2015. 1
- [65] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 3
- [66] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *ICCV*, 2019. 2, 4
- [67] Luca Schmidtko, Athanasios Vlontzos, Simon Ellershaw, Anna Lukens, Tomoki Arichi, and Bernhard Kainz. Unsupervised human pose estimation through transforming shape templates. In *CVPR*, 2021. 3
- [68] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *CVPR*, 2021. 7, 8
- [69] Shixiang Tang, Cheng Chen, Qingsong Xie, Meilin Chen, Yizhou Wang, Yuanzheng Ci, Lei Bai, Feng Zhu, Haiyang Yang, Li Yi, et al. Humanbench: Towards general human-centric perception with projector assisted pretraining. In *CVPR*, 2023. 3
- [70] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *NeurIPS*, 2022. 3
- [71] Shuhei Tsuchida, Satoru Fukayama, Masahiro Hamasaki, and Masataka Goto. Aist dance video database: Multi-genre, multi-dancer, and multi-camera database for dance information processing. In *ISMIR*, 2019. 2, 4, 5
- [72] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *CVPR*, 2018. 3
- [73] Gul Varol, Javier Romero, Xavier Martin, Nareen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *CVPR*, 2017. 4
- [74] Timo von Marcard, Roberto Henschel, Michael J. Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *ECCV*, 2018. 1, 2, 3, 5, 8
- [75] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked Feature Prediction for Self-Supervised Visual Pre-Training. In *CVPR*, 2022. 3
- [76] Weinzaepfel, Philippe and Leroy, Vincent and Lucas, Thomas and Brégier, Romain and Cabon, Yohann and Arora, Vaibhav and Antsfeld, Leonid and Chidlovskii, Boris and Csurka, Gabriela and Revaud Jérôme. CroCo: Self-Supervised Pre-training for 3D Vision Tasks by Cross-View Completion. In *NeurIPS*, 2022. 2, 3, 4, 5, 6
- [77] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance-level discrimination. In *CVPR*, 2018. 1, 3
- [78] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. SimMIM: A Simple Framework for Masked Image Modeling. In *CVPR*, 2022. 3
- [79] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. In *NeurIPS*, 2022. 2, 3
- [80] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 1
- [81] Zhixuan Yu, Jae Shin Yoon, In Kyu Lee, Prashanth

- Venkatesh, Jaesik Park, Jihun Yu, and Hyun Soo Park. Humbi: A large multiview dataset of human body expressions. In *CVPR*, 2020. 2, 4
- [82] Weiyu Zhang, Menglong Zhu, and Konstantinos G. Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *ICCV*, 2013. 2, 5
- [83] Long Zhao, Yuxiao Wang, Jiaping Zhao, Liangzhe Yuan, Jennifer J Sun, Florian Schroff, Hartwig Adam, Xi Peng, Dimitris Metaxas, and Ting Liu. Learning view-disentangled human pose representation by contrastive cross-view mutual information maximization. In *CVPR*, 2021. 3
- [84] Liang Zheng, Zhi Bie, Yifan Sun, Jingdong Wang, Chi Su, Shengjin Wang, and Qi Tian. Mars: A video benchmark for large-scale person re-identification. In *ECCV*, 2016. 2, 5
- [85] Christian Zimmermann, Max Argus, and Thomas Brox. Contrastive representation learning for hand shape estimation. In *GCPH*, 2021. 8
- [86] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan Russell, Max Argus, and Thomas Brox. Freihand: A dataset for markerless capture of hand pose and shape from single rgb images. In *ICCV*, 2019. 6, 8