

TempCLR: Reconstructing Hands via Time-Coherent Contrastive Learning

Andrea Ziani^{1*} Zicong Fan^{1,2*} Muhammed Kocabas^{1,2} Sammy Christen¹ Otmar Hilliges¹
¹ETH Zürich, Switzerland ²Max Planck Institute for Intelligent Systems, Tübingen




Figure 1: State-of-the-art hand reconstruction methods such as [7] (middle), fail to keep coherent hand representations through time. We exploit the underlying temporal constraint in unlabelled videos and train a model in a time-contrastive manner. Our method (TempCLR) keeps embeddings of the same sequence closer in the latent space and achieves better generalization on unseen videos, reconstructing more coherent hands through time.

Abstract

We introduce TempCLR, a new time-coherent contrastive learning approach for the structured regression task of 3D hand reconstruction. Unlike previous time-contrastive methods for hand pose estimation, our framework considers temporal consistency in its augmentation scheme, and accounts for the differences of hand poses along the temporal direction. Our data-driven method leverages unlabelled videos and a standard CNN, without relying on synthetic data, pseudo-labels, or specialized architectures. Our approach improves the performance of fully-supervised hand reconstruction methods by 15.9% and 7.6% in PA-V2V on the HO-3D and FreiHAND datasets respectively, thus establishing new state-of-the-art performance. Finally, we demonstrate that our approach produces smoother hand reconstructions through time, and is more robust to heavy occlusions compared to the previous state-of-the-art which we show quantitatively and qualitatively. Our code and models will be available at <https://eth-ait.github.io/tempclr>.

1. Introduction

Methods for hand pose and shape reconstruction have many applications in human-computer interaction, aug-

mented reality, virtual reality, robotics, and motion generation [40, 39, 8]. Recent research demonstrates impressive results on the task of *supervised* 3D hand reconstruction from monocular RGB images (e.g. [30, 43, 15]). However, generalizing to in-the-wild settings, with fully unconstrained and uncontrollable environmental conditions, would require large amounts of training data captured under the same conditions. As of today, accurate 3D keypoint annotation of in-the-wild data is an open research problem and, therefore, no large-scale in-the-wild dataset with accurate 3D annotations exists. For these reasons, techniques that leverage sparsely annotated data [14] or weakly labelled data [25, 4, 23] have seen much interest. However, such methods rely on pseudo 2D or 3D annotations, which in turn require human effort for acquisition, or may introduce label noise that bounds model performance [4, 23]. Therefore, a promising solution to avoid pseudo-labels entirely, is to make use of *unlabelled* data, for example via contrastive learning [34, 45]. In the context of sequence data, we observe that existing methods often struggle with heavy occlusions, for instance brought on by hand-object interaction. Consider the example from Fig. 1: while the hand pose throughout the grasp is quasi-static, the images change drastically from frame to frame, which causes existing methods to output incorrect hand poses. In this paper, we explore how to learn better representations that capture human motion’s inherent temporal consistency, improving

*Equal contribution

the hand reconstruction stability through time. We do so by leveraging *single-view unlabelled* videos of hands grasping objects to improve 3D hand reconstruction in the most challenging setting of heavy occlusions.

Unlike single images, videos contain temporal information that can help to predict coherent hand reconstructions through time by learning correlations between time-adjacent frames. Combining this idea with the recent progress of contrastive representation learning methods [31, 6, 16], we introduce a *time-coherent contrastive learning* pipeline, dubbed TempCLR. Our approach consists of two stages, as shown in Figure 2: a pre-training stage where we perform *time-coherent contrastive learning* on unlabelled videos and a second stage, where the pre-trained encoder is fine-tuned on the 3D hand reconstruction task using labelled data. In particular, TempCLR contributes two key ideas: 1) a time-coherent augmentation method to impose strong spatial augmentations on each frame of a video while maintaining temporal integrity; and 2) a probabilistic sampling strategy that accounts for the differences in frames along the temporal dimension. In contrast to a vanilla time-contrastive learning approach [45], which repels any non-neighboring frame in a sequence, our sampling strategy takes into consideration that temporally-closer frames often represent more similar hand poses in range of motion. Based on this insight, TempCLR gives more attention to attracting temporally close frames and only repels temporally distant frames. Figure 1 shows that our approach is able to produce smoother hand reconstructions along time, where a state-of-the-art approach [7] fails to do so.

We evaluate TempCLR in different settings and on different datasets. First, we demonstrate that our pre-training improves the performance over a fully-supervised baseline [7] by 15.9% and 7.6% in 3D mesh error on the HO-3D and FreiHAND datasets (*c.f.* Tab. 1 and Tab. 2). Next, we show that our *single-view* time-contrastive method improves over a vanilla time-contrastive approach [45] on FreiHAND. Through cross-dataset evaluation and in-the-wild qualitative results, we show improvements in generalization capabilities. Finally, we demonstrate that our method yields smoother hand reconstructions along the temporal dimension compared to other SotA approaches.

Our contributions can be summarized as follows:

1. A novel *single-view* time-contrastive learning approach for 3D hand reconstruction. Our method leverages time-coherent augmentations and a probabilistic sampling strategy to capture long-range dependencies.
2. We experimentally show that by leveraging in-the-wild unlabelled monocular videos, TempCLR outperforms existing methods across different metrics.
3. We provide empirical evidence that our method leads to smoother hand poses estimated over time.

2. Related Work

Fully-supervised 3D hand reconstruction: Reconstructing hands in 3D from images has received increased attention in recent years [38, 42]. Existing methods [12, 19, 27, 9, 33, 46, 2, 43, 26, 15, 35, 41, 18, 11] often leverage full supervision from in-the-lab datasets. For instance, Zimmermann *et al.* [46] propose the first convolutional network to detect 2D hand joints and lift them into the 3D space with an articulation prior. Iqbal *et al.* [19] introduce a 2.5D representation for 3D hand pose estimation. Boukhayma *et al.* [2] and Choutas *et al.* [7] estimate MANO [30] and SMPL-X [28] parameters using a weak perspective camera model. Lin *et al.* [24] introduce a transformer architecture to estimate vertices of the MANO mesh. In contrast to these approaches, we focus on leveraging additional supervision from unlabelled videos to improve 3D hand reconstruction.

Reconstructing hands from limited supervision: Recently, several datasets for 3D hand pose and shape estimation have been introduced [13, 5, 10, 47, 26, 20]. However, capturing 3D hand annotation is difficult to scale: 1) Magnetic trackers [10] provide 3D annotation for hands and objects but they are intrusive and introduce noise in RGB images. 2) Multi-view setups [13, 5, 26, 3] are marker-less, but the labels are obtained by either manual 2D annotation with triangulation [26, 5] or from noisy multi-kinect systems [13]; the quantity of 3D labelled data is still limited, and the background is not diverse. 3) Synthetic data provides perfect ground-truth but lacks photorealism [15, 27].

To allow methods to generalize to unconstrained settings, recently, there has been attention on reducing the reliance on 3D annotation [1, 43, 4, 2, 14, 25, 34, 45, 37, 36]. For example, Hasson *et al.* [14] leverage sparsely annotated data by introducing a photometric loss formulation to learn from partially labelled sequences. Liu *et al.* [25] propose a specialized transformer-based architecture used to collect pseudo labels from in-the-wild videos. These pseudo labels are then used to train the same architecture. Zimmermann *et al.* [45] explore the benefits of multi-view and single-view time-contrastive learning applied on the hand reconstruction task. Spurr *et al.* [34] introduce an equivariant contrastive objective formulation where geometric transformations applied on the image are reversed in the latent space. In this paper, we introduce a self-supervised approach to leverage supervision on unlabelled monocular videos in the wild.

The most relevant methods to us are [14, 25, 34, 45], which leverage unlabelled or partially labelled data. Compared to [14, 25], our method requires neither human intervention for tuning pseudo-labels [25], nor sparsely annotated videos [14]. Similarly to ours, the methods in [34, 45] use a contrastive formulation. However, [34] relies on unlabelled in-the-wild still images while we rely on unlabelled in-the-wild videos. In addition, [45] leverage a multi-view

time-contrastive formulation while our approach is based on monocular videos. Furthermore, to go beyond [45], we introduce a simple-yet-effective time-coherent augmentation method and sampling strategy that reflects the differences in frames along time. Experiments show that this novel combination is crucial for time-contrastive learning.

3. Method

Figure 2 shows a schematic of our method, TempCLR, which consists of two stages: a pre-training stage, and a fine-tuning stage. In the pre-training stage, we leverage a time-contrastive objective to train the image encoder on unlabelled videos. This stage is to obtain additional supervision for the encoder from diverse in-the-wild videos of hand in motion. In the second stage, we train the whole hand reconstruction architecture through supervised fine-tuning. In Section 3.1, we describe our time-contrastive pre-training, motivating the importance of our data augmentation and probabilistic sampling technique. Then, in Sec. 3.2 we present our hand reconstruction model.

3.1. Time-contrastive Learning

We build our self-supervised time-contrastive learning framework as illustrated in Fig. 2A. The core of our framework is an NT-Xent loss [6] applied on features extracted from augmented frames of a sequence (the augmentation module is described below). We denote a video as $X = \{x_1, x_2, \dots, x_n\}$, where x_t is the t -th frame of the sequence. Around a reference frame x_i , we define the temporal window $T_i = \{x_{i-k}, \dots, x_{i-1}, x_{i+1}, \dots, x_{i+k}\}$ with size $2k$. Frames inside this temporal range correspond to the candidate positive pairs of frame x_i , while all the other frames of the same video correspond to candidate negative pairs. We use z_i to denote the encoded representation of x_i .

Suppose that we sample M frames per mini-batch, possibly from different videos; for each frame x_i we sample $P_i \subseteq T_i$ (positive pairs), and $N_i \subseteq X \setminus T_i$ (negative pairs). $|P_i|$ and $|N_i|$ are fixed. The NT-Xent loss is defined as:

$$L = \frac{1}{M} \sum_{i=1}^M L_i, \quad (1)$$

$$L_i = - \sum_{x_j \in P_i} \log \frac{\exp(\text{sim}(z_i, z_j)/\tau)}{\sum_{x_k \in N_i} \exp(\text{sim}(z_i, z_k)/\tau)}. \quad (2)$$

Here, $\tau > 0$ is a temperature parameter and $\text{sim}(u, v) = u^T v / \|u\| \|v\|$ is the cosine similarity between z_i and z_j . Hence, the loss encourages embeddings of similar, neighboring frames (positive pairs) to be mutually attracted while those of dissimilar frames in the same sequence (negative pairs) are kept far apart.

Time-coherent geometric transformations: Data augmentations are extensively used in contrastive training for

computer vision tasks [6, 16, 29]. Although a common optimal augmentation procedure does not exist, in a temporal setting a natural approach is to employ existing augmentation methods to the frames of the video one by one. Image augmentation methods often include geometric transformations such as random cropping, rotation, translation. In sequences, however, such transformations could break the inherent motion cues between consecutive frames, negatively affecting representation learning along the temporal dimension. Inspired by Qian *et al.* [29], we apply consistent augmentations through time by applying the same random geometric transformations (*i.e.* rotation, scale, and translation) across frames of the same sequence, while applying independent appearance transformation for each frame (see Figure 3). In this way, the encoder better captures temporal features in the pre-training stage.

Probabilistic pair sampling: Existing method on time-contrastive learning for hand pose estimation [45] defines two immediate neighbouring frames as positive pairs and any couple of non-neighbouring frames as negatives pairs. In the grasping scenario, however, the hand pose has a limited range of movement caused by the interaction between the hand and the object. This means that several consecutive frames could represent similar hand poses and a trivial pair selection may not be beneficial. To address this problem, our key insight is that two images from the same video represent more diverse hand poses when their temporal distance is large. To this end, we use a sampling strategy to account for the temporal changes (see Fig. 3). In particular, given a frame x_i sampled uniformly at random from a sequence, we first define a temporal window T_i , as described in the previous section. Then, from the temporal window, we sample P_i positive pairs with a probability distribution that monotonically decreases with the distance from x_i . Likewise, we sample N_i negative pairs, lying outside the temporal window, with a probability directly proportional to the distance from x_i . Following our sampling strategy, the contrastive training will focus more on attracting temporally closer frames and repelling temporally more distant frames, while reducing the attention that is given to the grey zone of frames representing hand poses with uncertain similarity to x_i .

To summarize our pre-training approach, first, each frame of a sequence is augmented by the same geometric transformation. Then, each frame is augmented independently via random (potentially different) appearance augmentations. After that, the sampling strategy chooses the positive and negative frames. See **SupMat** for more details.

3.2. Hand Reconstruction

Figure 2B shows our hand reconstruction network. Following [7, 14, 2], we use an encoder-decoder formulation. In particular, our method consists of our pre-trained encoder




Figure 2: **Overview of TempCLR:** A) An encoder is trained with a time-contrastive learning approach on unlabelled videos of hands grasping objects. B) The pre-trained encoder is fine-tuned using labelled data.




Figure 3: **Overview of the time-coherent augmentations (top) and the probabilistic sampling step (bottom).**

to obtain an image feature vector and a hand decoder to predict the MANO pose and shape parameters and the weak perspective camera parameters (scale and translation). Formally, given an image, the network predicts the MANO pose vector $\theta = [\theta^{\text{wrist}}; \theta^{\text{fingers}}] \in \mathbb{R}^{16 \times D}$, shape parameters $\beta \in \mathbb{R}^{10}$ and the weak perspective camera parameters (t, s) . The MANO parameters are fed into the MANO differentiable layer to retrieve the 3D hand mesh. The weak perspective camera model aligns the mesh with the image. Following [7], we use the rotation representation from Zhou *et al.* [44] for our MANO pose parameters ($D = 6$).

We train our model using 2D re-projection loss, 3D joint errors, and pose and shape parameter loss $L = \lambda_{2D} L_{2D} + \lambda_{3D} L_{3D} + \lambda_{\Theta} L_{\Theta}$, where $L_{2D} = \|J^{2D} - \hat{J}^{2D}\|_1$, $L_{3D} = \|J^{3D} - \hat{J}^{3D}\|_1$, and $L_{\Theta} = \|\{\theta, \beta\} - \{\hat{\theta}, \hat{\beta}\}\|_2^2$. All variables with a hat denote predictions and $J^{2D} \in \mathbb{R}^{21 \times 2}$ and $J^{3D} \in \mathbb{R}^{21 \times 3}$ represent the 21 keypoints in 2D and 3D.

4. Experiments

In Section 4.1, we first introduce experiment details such as the datasets, the evaluation metrics, and the implementation details. In Sec. 4.2, we compare our method to state-of-the-art approaches on both hand-grasping-objects and hand-only settings. In Sec. 4.3, we ablate TempCLR and provide qualitative results. Also, we show the effectiveness of TempCLR when 3D annotations are scarce. Finally, in Sec. 4.4, we perform cross-dataset evaluation to demonstrate generalization under domain shifts.

4.1. Datasets, Metrics, and Implementation Details

HO-3D [13]: The dataset provides 3D hand-object annotations during interaction for markerless RGB images. The ground-truth annotations are obtained by fitting a hand model to multi-view RGB-D evidence. We present results on HO-3D v2. The evaluation is performed online; hence we do not have access to the ground truth of the test set.

FreiHAND [47], HanCo [45]: FreiHAND (FH) consists of 130k training and 4k evaluation samples captured with a green screen background in the training set, as well as real backgrounds in the test set. Both 3D and 2D annotations are provided. HanCo does not contain 3D annotations. It only contains short video clips recorded with a calibrated and time-synchronized multi-view camera capture setup. In total, there are 107k time-steps recorded by eight cameras, which results in 860k RGB images. As these datasets are composed of both hand-only and hand-grasping-object sequences, we used HanCo in the time-contrastive pre-training and FH in supervised fine-tuning.

100 Days Of Hands [32]: This is a large-scale and in-the-wild dataset of hand-object interaction footage. The dataset does not provide any hand annotation besides the bounding boxes of the hands in the scene. In some of our experiments, we used a subset of 10 videos collected from this dataset (86k images) exclusively as additional unlabelled frames for time-contrastive pre-training. We show that this pre-training improves hand reconstruction.

Evaluation metrics: We report the End-Point-Error (EPE)

Method	PA-V2V (mm) ↓	PA-EPE (mm) ↓	F@5 mm ↑	F@15 mm ↑
Baseline [7]	12.6	12.7	0.389	0.905
Hasson <i>et al.</i> [14]	11.4	11.4	0.428	0.932
Hasson <i>et al.</i> [15]	11.2	11.1	0.464	0.939
PeCLR [34]	10.8	10.8	0.47	0.936
TempCLR (ours)	10.6	10.6	0.481	0.937
PeCLR [†] [34]	11.0	11.0	0.46	0.934
TempCLR[†] (ours)	10.0	10.1	0.505	0.947
Liu <i>et al.</i> [25]	9.5	9.9	0.526	0.955

Table 1: **Comparison with SotA on HO-3D [13].** TempCLR outperforms the baselines on all reported metrics. The employment of additional in-the-wild data for contrastive pre-training, denoted by [†], further improves the model’s accuracy. Results below the dashed line employ weak supervision.

and the Vertex-to-Vertex End-Point-Error (V2V). The former denotes the average L2 distance between the ground-truth and predicted keypoints, while the latter denotes the average L2 distance between the ground-truth and mesh vertices. We prefix the metrics with PA, RA and STA to denote procrustes alignment, root alignment, and scale-and-translation alignment. We include the F-scores defined as the harmonic mean between recall and precision between two meshes given a distance threshold. Following [22], to measure the temporal stability of the reconstruction, we compute an acceleration error by measuring the difference in acceleration between the 3D GT and the predictions.

Implementation details: For the pre-training we use ResNet [17] as a backbone, which takes monocular RGB images of size 224×224 as input. We employ Adam [21] as the optimizer with a batch size of 2048 and a learning rate of $4.5e^{-3}$ for 50 epochs. The fine-tuning is performed until convergence based on the performance on the validation set. During fine-tuning, we use RGB images of size 224×224 as input. As optimizer, we use Adam with a learning rate of $5e^{-4}$ and a batch size of 128. Further details can be found in **SupMat**. Following [31], we choose the window size to be approximately half of the frame rate for each dataset (15 for HO-3D and 100DOH, 5 for HanCo).

4.2. Comparison with the State-of-the-Art

Here we compare TempCLR with fully-supervised and self-supervised state-of-the-art approaches on HO-3D and FH. Figure 6 shows qualitative results.

Comparison on HO-3D: Table 1 compares TempCLR with the fully-supervised and self-supervised state-of-the-art on HO-3D. First, we pre-train a ResNet18 encoder on unlabelled HO-3D images. Then, we fine-tune the hand reconstruction network with full supervision as described in Sec. 3.2. To show that our self-supervised method can lever-

Method	PA-V2V (mm) ↓	RA-V2V (mm) ↓	F@5 (mm) ↑	F@15 (mm) ↑
Hasson <i>et al.</i> [15]	13.2	-	0.436	0.908
Baseline-18 [7]	11.8	35.96	0.484	0.918
TempCLR-18 (ours)	10.9	25.05	0.513	0.930
Baseline-50 [7]	10.8	31.15	0.518	0.934
MANO CNN [47]	10.7	-	0.529	0.935
HanCo Augm. [45]	10.9	-	0.521	0.934
HanCo Temporal [45]	10.4	-	0.538	0.939
PeCLR-50 [34]	10.6	26.73	0.533	0.940
TempCLR-50 (ours)	10.2	21.68	0.541	0.941
HanCo Multi-view [45]	10.2	-	0.548	0.943

Table 2: **Comparison with SotA on FH [47].** The top-bottom split (solid line) separates results using ResNet18 and ResNet50. The dashed line separates a multi-view temporal approach that is not directly comparable.

age in-the-wild unlabelled data, we repeat the experiment but include additional unlabelled frames from 100DOH, along with the original unlabelled frames in HO-3D, during the contrastive training phase.

Top rows of the table show that TempCLR, without employing any in-the-wild data, improves over our fully-supervised baseline [7] (see Baseline on the table) by 15.9% in PA-V2V and PA-EPE. Furthermore, using additional in-the-wild data for time-contrastive pre-training (denoted by [†] in Tab. 1), TempCLR improves further and establishes the new state-of-the-art for self-supervised training. Notably, TempCLR is on par with [25], a weakly-supervised method that uses pseudo-labels. The labels involve manual intervention to generate. TempCLR is self-supervised, so it does not require intervention to train on unlabelled videos.

With additional in-the-wild data, PeCLR pre-training does not further improve. This is consistent to the observation in Fig. 6 of the PeCLR paper – although PeCLR improves hand poses by leveraging additional in-the-wild data (FH+YT3D) compared to fully-supervised training (FH), the improvement is significant in *low data regime*. With more annotation, training with additional in-the-wild data does not lower the error. In contrast, our method consistently improves over the baseline in both low data and high data regime (see Tab. 1 and Fig. 5).

Comparison on FH: Here we use the HanCo dataset alone to perform our contrastive pre-training on ResNet18 and ResNet50 encoders. To show the efficacy of TempCLR, we compare the results produced by our pipeline against fully-supervised methods and state-of-the-art contrastive approaches [45, 34]. Before diving into results, we highlight that we report the RA-V2V scores for the fully-supervised baseline (ExPose [7]) and for our time-contrastive approach only. This is because the FH test set was previously hidden and hosted as competition online, where this metric was not computed. Moreover, we do not have access to the pre-

Method	Accel. (mm/s^2) ↓	RA-EPE (mm) ↓	PA-EPE (mm) ↓
Baseline [7]	54.11	75.42	15.30
TempCLR w/o consist. augm.	45.87	61.28	14.51
TempCLR w/o prob. sampling	47.56	52.31	13.80
TempCLR	45.37	51.72	13.69

Table 3: **Temporal stability evaluation.** Our augmentation strategy improves the hand pose estimation performance, while the sampling strategy helps in temporal stability. The combination of the two leads to the best results.

trained models to reproduce the missing results.

Table 2 shows that TempCLR improves over the ResNet18 fully-supervised baseline by 30.4% in RA-V2V and by 7.6% in PA-V2V, indicating a significant improvement in global orientation and scale. Similarly, with a ResNet50 backbone, TempCLR improves over the baseline by 30.4% in RA-V2V and by 5.5% in PA-V2V. Finally, we establish state-of-the-art performance by improving over the single-view self-supervised approach [45]. Note that the RA-V2V metric is not available for [45]. Our single-view time-contrastive approach is on par with the multi-view time-contrastive approach proposed by Zimmermann *et al.* [45]. We emphasize that monocular videos are more abundant on the Internet and often have very diverse environments in comparison to controlled multi-view setup.

4.3. Ablation Study

Here we ablate our method, and support it with quantitative and qualitative results. First, we analyse the embedding space learned through TempCLR pre-training and compare it to an ImageNet pre-trained encoder. We ablate the importance of time-coherent augmentation and probabilistic sampling for time-contrastive learning, and we provide evidence that time-coherent contrastive learning leads to more stable hand reconstructions through time. Next, we compare different probabilistic sampling strategies. Lastly, we evaluate the efficacy of TempCLR when ground-truth data for fine-tuning is scarce.

Latent space representation: Figure 4 shows a t-SNE plot of two embedding spaces comparing the ImageNet pre-trained backbone and our backbone with a TempCLR pre-training. In particular, ten different sequences from the HanCo dataset have been randomly sampled and augmented. For each image of these sequences, we extract their feature vector and perform a t-SNE clustering. We see that TempCLR leads to better cluster separation and, within the same cluster, similar hand poses are closer in the embedding space. This confirms that our method yields the desired latent spaces we described in Sec. 3.

Effects of time-coherent augmentation and probabilistic sampling: We compare the fully-supervised baseline [7] trained on FH, and our method pre-trained on HanCo and

Augmentation	RA-EPE (mm) ↓	PA-EPE (mm) ↓
Rotation	141.34	18.02
Translation	99.04	15.93
Scale	98.21	15.50
Channel Noise	96.76	15.70
Color Drop	98.19	15.60
Color Jitter	97.36	15.76
Sobel Filter	97.35	15.71

Table 4: **Effects of different data augmentations in pre-training.** We pre-train on HanCo, fine-tune on FH, evaluate on FH test set.

Method	RA-EPE (mm) ↓	PA-EPE (mm) ↓
Baseline [7]	35.96	11.8
TempCLR-Lin	25.05	10.9
TempCLR-Exp	28.91	11.4
TempCLR-Tanh	28.85	11.1

Table 5: **Effects of different sampling strategies in pre-training.** We pre-train on HanCo, fine-tune on FH, and evaluate on FH test set.

fine-tuned on FH. In addition, we investigate the influence on the final performance of each of our contributions by removing our time-coherent geometric augmentation and the probabilistic sampling strategy (see Sec. 3). Since FreiHAND is not a temporal dataset and the HO-3D test set is hidden, we evaluate on the HO-3D training split. Table 3 shows that the greatest improvement in hand pose estimation (RA-EPE and PA-EPE) comes from the augmentation strategy, while the probabilistic sampling strategy contributes more to the temporal stability (see the acceleration metric). These results confirm our insight that when performing time-contrastive learning for images with hands in motion, it is crucial to sample distant frames to ensure the feasibility of the pre-training task. The acceleration metric demonstrates that our pre-training leads to more stable results even using a single-frame model. Moreover, the time-coherent geometric augmentation and the sampling strategy complement each other and the combination of the two leads to the best overall improvement. See **SupMat** for additional qualitative results and failure cases.

Different augmentation strategies: Table 4 shows the impact of different augmentations in the pre-training stage. In particular, we pre-train on HanCo [45], and fine-tune on FreiHAND [47] with a ResNet18 [17] backbone. Similar to [34], the appearance transformations are more beneficial than geometric transformations. This motivates our choice to keep independent appearance transformations for each frame of a sequence while preserving the motion of the video with coherent geometric transformations in time.

Different sampling strategies: Table 5 shows the effects




Figure 4: **Comparison of the 2D t -SNE embeddings** produced by an encoder pre-trained on ImageNet and by our time-contrastive pre-trained encoder. On the right hand side, we see that hand poses close along the temporal dimension are located in proximity to each other. Contrary, on the left hand side hand poses close in time are more distant in the embedding space.




Figure 5: **Self-supervised performance on HO-3D.** TempCLR achieves better PA-EPE (top) and STA-EPE (bottom) performances than the fully-supervised baseline ExPose [7]. Additional in-the-wild unlabelled data improves TempCLR further.

of different sampling strategies. Namely, we compare linear sampling, exponential sampling, and sampling using the absolute value of the hyperbolic tangent function. We see that linear sampling leads to the best performance.

Learning with different amount of supervision: We investigate the impact of our pre-training objective with respective to different amount of human-annotated data and in-the-wild unlabelled data. The ExPose [7] baseline uses an ImageNet pre-trained encoder. For our method, we apply time-contrastive pre-training either using HO-3D only, or HO-3D plus 100DOH to demonstrate the advantage of adding in-the-wild data for self-supervised training. All the hand reconstruction networks are fine-tuned on sparsely annotated sequences from HO-3D. We evaluate the performance of the networks on the HO-3D test set. Fig. 5 summarizes the results in EPE by progressively increasing the percentage of annotated frames from 5% to 40%. We see that, TempCLR consistently improves hand reconstruction by leveraging additional unlabelled data. Moreover, the use of additional in-the-wild unlabelled data (see 100DOH) further improves our performance. Interestingly, only 20% of supervised frames are necessary to reach the performance of more densely annotated data. This behaviour is confirmed by [14] and can be explained by the high correlation between neighboring frames of the HO-3D sequences.

Window size: When trained on HanCo and fine-tuned on FreiHAND, the PA-EPE error of TempCLR with the window sizes 3, 5, 15 are $11.1mm$, $10.9mm$, $11.3mm$, respectively. Future work could leverage optical flow to detect changes in the sequences for an adaptive window size.

4.4. Cross-dataset Evaluation

Cross-dataset generalization is rarely reported in the hand reconstruction literature, perhaps because it is widely assumed to be challenging. Yet, it is clearly important for real-world applications. Given the use of a large amount




Figure 6: **Qualitative results** on HanCo [45] unlabelled sequences, HO-3D [13] test set, and in-the-wild [32] unlabelled sequences. Predictions are produced using models described in Sec. 4.2. Further qualitative results can be found in **SupMat**.

Method	HO-3D (train), FH (test)	FH (train), HO-3D (test)
	RA-EPE/PA-EPE	STA-EPE/PA-EPE
Baseline	104.5/18.5	66.1/13.9
PeCLR	96.0/17.8	62.2/13.6
TempCLR	84.6/17.0	53.5/13.6

Table 6: **Cross-dataset evaluation.** Methods are trained on HO-3D and evaluated on FH and vice versa. TempCLR generalizes best in both domain shifts. Metrics are in mm.

of unlabelled data for time-contrastive pre-training, we expect our approach to produce features that are beneficial for generalization on unseen scenes. To this end, we verify the effectiveness of the models from Sec. 4.2 in a cross-dataset setting. In particular, we evaluate the performance of the model when trained on FH and evaluated on HO-3D, and vice versa. This reveals how the models perform under a domain shift. Table 6 reports an improvement over the baseline of 19% in both RA-EPE on FreiHAND and STA-EPE on HO-3D. These results show that our pre-training objec-

tive enables better generalization to unseen scene.

5. Conclusion

We introduce, TempCLR, a time-contrastive method for hand pose and shape estimation that yields stable 3D reconstructions through time. We introduce time-coherent augmentations and probabilistic pair sampling to better account for the temporal information provided by unlabelled videos. We thoroughly investigate our method, showing that it better captures temporal features and improves reconstruction stability through time. We demonstrate that our TempCLR achieves state-of-the-art results on the HO-3D and FreiHAND datasets. Finally, by means of cross-dataset evaluation, we show the potential of our method’s generalization capabilities.

Acknowledgement: Muhammed Kocabas is supported by the Max Planck ETH Center for Learning Systems. The authors would like to thank Vassilis Choutas for providing the code of the baseline model adopted for the project.

References

- [1] Seungryul Baek, Kwang In Kim, and Tae-Kyun Kim. Pushing the envelope for RGB-based dense 3D hand pose estimation via neural rendering. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1067–1076, 2019.
- [2] Adnane Boukhayma, Rodrigo Andrade de Bem, and Philip H. S. Torr. 3D hand shape and pose from images in the wild. In *Computer Vision and Pattern Recognition (CVPR)*, pages 10843–10852, 2019.
- [3] Samarth Brahmbhatt, Chengcheng Tang, Christopher D. Twigg, Charles C. Kemp, and James Hays. ContactPose: A dataset of grasps with object contact and hand pose. In *European Conference on Computer Vision (ECCV)*, volume 12358, pages 361–378. Springer, 2020.
- [4] Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *International Conference on Computer Vision (ICCV)*, pages 12397–12406. IEEE, 2021.
- [5] 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 *Computer Vision and Pattern Recognition (CVPR)*, pages 9044–9053, 2021.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 2020.
- [7] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J. Black. Monocular expressive body regression through body-driven attention. In *European Conference on Computer Vision (ECCV)*, volume 12355, pages 20–40. Springer, 2020.
- [8] Sammy Christen, Muhammed Kocabas, Emre Aksan, Jemin Hwangbo, Jie Song, and Otmar Hilliges. D-Grasp: Physically plausible dynamic grasp synthesis for hand-object interactions. In *Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [9] Zicong Fan, Adrian Spurr, Muhammed Kocabas, Siyu Tang, Michael J. Black, and Otmar Hilliges. Learning to disambiguate strongly interacting hands via probabilistic per-pixel part segmentation. In *International Conference on 3D Vision (3DV)*, pages 1–10. IEEE, 2021.
- [10] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with RGB-D videos and 3D hand pose annotations. In *Computer Vision and Pattern Recognition (CVPR)*, pages 409–419. Computer Vision Foundation / IEEE Computer Society, 2018.
- [11] Lihao Ge, Zhou Ren, Yuncheng Li, Zehao Xue, Yingying Wang, Jianfei Cai, and Junsong Yuan. 3D hand shape and pose estimation from a single RGB image. In *Computer Vision and Pattern Recognition (CVPR)*, pages 10833–10842, 2019.
- [12] Patrick Grady, Chengcheng Tang, Christopher D. Twigg, Minh Vo, Samarth Brahmbhatt, and Charles C. Kemp. ContactOpt: Optimizing contact to improve grasps. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1471–1481, 2021.
- [13] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. HOnnotate: A method for 3D annotation of hand and object poses. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3193–3203, 2020.
- [14] Yana Hasson, Bugra Tekin, Federica Bogo, Ivan Laptev, Marc Pollefeys, and Cordelia Schmid. Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 568–577, 2020.
- [15] Yana Hasson, Gül Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11807–11816, 2019.
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735, 2020.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 770–778. IEEE Computer Society, 2016.
- [18] Yihui He, Rui Yan, Katerina Fragkiadaki, and Shoubo Yu. Epipolar transformers. In *Computer Vision and Pattern Recognition (CVPR)*, pages 7776–7785, 2020.
- [19] Umar Iqbal, Pavlo Molchanov, Thomas M. Breuel, Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5D heatmap regression. In *European Conference on Computer Vision (ECCV)*, volume 11215, pages 125–143. Springer, 2018.
- [20] Sheng Jin, Lumin Xu, Jin Xu, Can Wang, Wentao Liu, Chen Qian, Wanli Ouyang, and Ping Luo. Whole-body human pose estimation in the wild. In *European Conference on Computer Vision (ECCV)*, volume 12354, pages 196–214. Springer, 2020.
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- [22] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. VIBE: video inference for human body pose and shape estimation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5252–5262, 2020.
- [23] Dominik Kulon, Riza Alp Güler, Iasonas Kokkinos, Michael M. Bronstein, and Stefanos Zafeiriou. Weakly-supervised mesh-convolutional hand reconstruction in the wild. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4989–4999, 2020.
- [24] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1954–1963, 2021.

- [25] Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3D hand-object poses estimation with interactions in time. In *Computer Vision and Pattern Recognition (CVPR)*, pages 14687–14697, 2021.
- [26] Gyeongsik Moon and Kyoung Mu Lee. I2L-MeshNet: Image-to-voxel prediction network for accurate 3d human pose and mesh estimation from a single RGB image. In *European Conference on Computer Vision (ECCV)*, volume 12352, pages 752–768. Springer, 2020.
- [27] Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. GANerated hands for real-time 3D hand tracking from monocular RGB. In *Computer Vision and Pattern Recognition (CVPR)*, pages 49–59. Computer Vision Foundation / IEEE Computer Society, 2018.
- [28] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019.
- [29] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge J. Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6964–6974, 2021.
- [30] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: modeling and capturing hands and bodies together. *ACM Trans. Graph.*, 36(6):245:1–245:17, 2017.
- [31] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, and Sergey Levine. Time-contrastive networks: Self-supervised learning from video. In *ICRA*, pages 1134–1141. IEEE, 2018.
- [32] Dandan Shan, Jiaqi Geng, Michelle Shu, and David F. Fouhey. Understanding human hands in contact at internet scale. In *Computer Vision and Pattern Recognition (CVPR)*, pages 9866–9875, 2020.
- [33] Tomas Simon, Hanbyul Joo, Iain A. Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4645–4653. IEEE Computer Society, 2017.
- [34] Adrian Spurr, Aneesh Dahiya, Xi Wang, Xucong Zhang, and Otmar Hilliges. Self-supervised 3D hand pose estimation from monocular RGB via contrastive learning. In *International Conference on Computer Vision (ICCV)*, pages 11210–11219. IEEE, 2021.
- [35] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3d hand pose estimation via biomechanical constraints. In *European Conference on Computer Vision (ECCV)*, 2020.
- [36] Adrian Spurr, Pavlo Molchanov, Umar Iqbal, Jan Kautz, and Otmar Hilliges. Adversarial motion modelling helps semi-supervised hand pose estimation. *arXiv preprint arXiv:2106.05954*, 2021.
- [37] Adrian Spurr, Jie Song, Seonwook Park, and Otmar Hilliges. Cross-modal deep variational hand pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 89–98, 2018.
- [38] James Steven Supančič III, Grégory Rogez, Yi Yang, Jamie Shotton, and Deva Ramanan. Depth-based hand pose estimation: Data, methods, and challenges. In *International Conference on Computer Vision (ICCV)*, pages 1868–1876, 2015.
- [39] Omid Taheri, Vasileios Choutas, Michael J. Black, and Dimitrios Tzionas. GOAL: Generating 4D whole-body motion for hand-object grasping. In *CVPR*, pages 13263–13273, 2022.
- [40] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *European Conference on Computer Vision (ECCV)*, pages 581–600, Aug. 2020.
- [41] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+O: unified egocentric recognition of 3D hand-object poses and interactions. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4511–4520, 2019.
- [42] Shanxin Yuan, Guillermo Garcia-Hernando, Björn Stenger, Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee, Pavlo Molchanov, Jan Kautz, Sina Honari, Lihao Ge, Junsong Yuan, Xinghao Chen, Guijin Wang, Fan Yang, Kai Akiyama, Yang Wu, Qingfu Wan, Meysam Madadi, Sergio Escalera, Shile Li, Dongheui Lee, Iason Oikonomidis, Antonis A. Argyros, and Tae-Kyun Kim. Depth-based 3D hand pose estimation: From current achievements to future goals. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2636–2645, 2018.
- [43] Xiong Zhang, Qiang Li, Hong Mo, Wenbo Zhang, and Wen Zheng. End-to-end hand mesh recovery from a monocular RGB image. In *International Conference on Computer Vision (ICCV)*, pages 2354–2364. IEEE, 2019.
- [44] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5745–5753, 2019.
- [45] Christian Zimmermann, Max Argus, and Thomas Brox. Contrastive representation learning for hand shape estimation. In *German Conference on Pattern Recognition (GCPR)*, volume 13024, pages 250–264. Springer, 2021.
- [46] Christian Zimmermann and Thomas Brox. Learning to estimate 3D hand pose from single RGB images. In *International Conference on Computer Vision (ICCV)*, pages 4913–4921. IEEE Computer Society, 2017.
- [47] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan C. Russell, Max J. Argus, and Thomas Brox. FreiHAND: A dataset for markerless capture of hand pose and shape from single RGB images. In *International Conference on Computer Vision (ICCV)*, pages 813–822. IEEE, 2019.