Domain Adaptive Hand Keypoint and Pixel Localization in the Wild

Takehiko Ohkawa^{1,2}, Yu-Jhe Li², Qichen Fu², Ryosuke Furuta¹, Kris M. Kitani², and Yoichi Sato¹

¹ The University of Tokyo, Tokyo, Japan ² Carnegie Mellon University, PA, USA {ohkawa-t,furuta,ysato}@iis.u-tokyo.ac.jp, {yujheli,qichenf,kkitani}@cs.cmu.edu Project: https://tkhkaeio.github.io/projects/22-hand-ps-da/

Abstract. We aim to improve the performance of regressing hand keypoints and segmenting pixel-level hand masks under new imaging conditions (e.g., outdoors) when we only have labeled images taken under very different conditions (e.g., indoors). In the real world, it is important that the model trained for both tasks works under various imaging conditions. However, their variation covered by existing labeled hand datasets is limited. Thus, it is necessary to adapt the model trained on the labeled images (source) to unlabeled images (target) with unseen imaging conditions. While self-training domain adaptation methods (*i.e.*, learning from the unlabeled target images in a self-supervised manner) have been developed for both tasks, their training may degrade performance when the predictions on the target images are noisy. To avoid this, it is crucial to assign a low importance (confidence) weight to the noisy predictions during self-training. In this paper, we propose to utilize the divergence of two predictions to estimate the confidence of the target image for both tasks. These predictions are given from two separate networks, and their divergence helps identify the noisy predictions. To integrate our proposed confidence estimation into self-training, we propose a teacher-student framework where the two networks (teachers) provide supervision to a network (student) for self-training, and the teachers are learned from the student by knowledge distillation. Our experiments show its superiority over state-of-the-art methods in adaptation settings with different lighting, grasping objects, backgrounds, and camera viewpoints. Our method improves by 4% the multi-task score on HO3D compared to the latest adversarial adaptation method. We also validate our method on Ego4D, egocentric videos with rapid changes in imaging conditions outdoors.

1 Introduction

In the real world, hand keypoint regression and hand segmentation are considered important to work under broad imaging conditions for various computer vision applications, such as egocentric video understanding [28,17], hand-object interaction analysis [12,21], AR/VR [39,70], and assistive technology [40,37]. For



Fig. 1: We aim to adapt the model of localizing hand keypoints and pixel-level hand masks to new imaging conditions without annotation.

building models for both tasks, several labeled hand datasets have been proposed in laboratory settings, such as multi-camera studios [34,46,13,80] and attaching sensors to hands [24,75,26]. However, their imaging conditions do not adequately cover real-world imaging conditions [51], consisting of various lighting, hand-held objects, backgrounds, and camera viewpoints. In addition, the annotation of keypoints and pixel-level masks are not always available in real-world environments because they are labor-intensive to acquire. As shown in Fig. 1, when localizing hand keypoints and pixels in real-world egocentric videos [28] (*e.g.*, outdoors), we may only have access to a hand dataset [13] taken under completely different imaging conditions (*e.g.*, indoors). Given these limitations, we need methods that can robustly adapt the models trained on the available labeled images (source) to unlabeled images (target) with new imaging conditions.

To enable such adaptation, the approach of self-training domain adaptation has been developed for both tasks. This approach aims to learn unlabeled target images by optimizing a self-supervised task, which exhibits effectiveness in various domain adaptation tasks [15,18,67,77,7]. For keypoint estimation, consistency training, a method that regularizes keypoint predictions to be consistent under geometric transformations, has been proposed [73,66,77]. As for hand segmentation, prior studies use pseudo-labeling [53,7], which produces hard labels by thresholding a predicted class probability for updating a network. However, these self-training methods for both tasks perform well only when the predictions are reasonably correct. When the predictions become noisy due to the gap in imaging conditions, the trained network will cause over-fitting to the noisy predictions, resulting in poor performance in the target domain.

To avoid this, it is crucial to assign a low importance (confidence) weight to the loss of self-training with noisy predictions. This confidence weighting can mitigate the distractions from the noisy predictions. To this end, we propose self-training domain adaptation with confidence estimation for hand keypoint regression and hand segmentation. Our proposed method consists of (i) confidence estimation based on the divergence of two networks' predictions and (ii) an update rule that integrates a training network for self-training and the two networks for confidence estimation. To (i) estimate confidence, we utilize the predictions of two different networks. While class probability can be used as the confidence in classification tasks, it is not trivial to obtain such a measure in keypoint regression. Thus, we newly focus on the divergence of the two networks' predictions for each target image. We design their networks to have an identical architecture but have different learning parameters. We observe that when the divergence measure is high, the predictions of both networks are noisy and should be avoided in self-training.

To (ii) integrate the estimated confidence into self-training, inspired by the single-teacher-single-student update [64.54], we develop mutual training with self-training based on consistency training for a training network (student) and distillation-based update for the two networks (teachers). For training the student network, we build a unified self-training framework that can work favorably for the two tasks. Motivated by supervised or weakly-supervised learning for jointly estimating both tasks [68,76,27,49,16], we expect that jointly adapting both tasks will allow one task to provide useful cues to the other task even in the unlabeled target domain. Specifically, we enforce the student network to generate consistent predictions for both tasks under geometric augmentation. We weight the loss of the consistency training using the confidence estimated from the divergence of the teachers' predictions. This can reduce the weight of the noisy predictions during the consistency training. To learn the two teacher networks differently, we train the teachers independently from different mini-batches by knowledge distillation, which matches the teacher-student predictions in the output level. This framework enables the teachers to update more carefully than the student and prevent over-fitting to the noisy predictions. Such stable teachers provide reliable confidence estimation for the student's training.

In our experiments, we validate our proposed method in adaptation settings where lighting, grasping objects, backgrounds, camera viewpoints, etc., vary between labeled source images and unlabeled target images. We use a large-scale hand dataset captured in a multi-camera system [13] as the source dataset (see Fig. 1). For the target dataset, we use HO3D [29] with different environments, HanCo [78] with multiple viewpoints and diverse backgrounds, and FPHA [24] with a novel first-person camera viewpoint. We also apply our method to inthe-wild egocentric video Ego4D [28] (see Fig. 1), including diverse indoor and outdoor activities worldwide. Our method improves the average score of the two tasks by 14.4%, 14.9%, and 18.0% on HO3D, HanCo, and FPHA, respectively, compared to a unadapted baseline. Our method further exhibits distinct improvements compared to the latest adversarial adaption method [33] and consistency training baselines with uncertainty estimation [7], confident instance selection [53], and the teacher-student scheme [64]. We finally confirm that our method also performs qualitatively well on the Ego4D videos.

Our contributions are summarized as follows:

 We propose a novel confidence estimation method based on the divergence of the predictions from two teacher networks for self-training domain adaptation of hand keypoint regression and hand segmentation.

- 4 T. Ohkawa et al.
- To integrate our proposed confidence estimation into self-training, we propose mutual training using knowledge distillation with a student network for self-training and two teacher networks for confidence estimation.
- Our proposed framework outperforms state-of-the-art methods under three adaptation settings across different imaging conditions. It also shows improved qualitative performance on in-the-wild egocentric videos.

2 Related Work

Hand keypoint regression is the task of regressing the positions of hand joint keypoints from a cropped hand image. 2D hand keypoint regression is trained by optimizing keypoint heatmaps [69,50,79] or directly predicting keypoint coordinates [60]. The 2D keypoints are informative for estimating 3D hand poses [61,74,47,5]. To build an accurate keypoint regressor, collecting massive hand keypoint annotations is required but laborious. While early works annotate the keypoints manually from a single view [56, 62, 48], recent studies have collected the annotation more densely and efficiently using synthetic hand models [30,47,79,48], hand sensors [24,75,63,26], or multi-camera setups [34,46,13,6,80,29,43]. However, these methods suffer the gap in imaging conditions with real-world images in deployment [51]. For instance, the synthetic hand models and hand sensors induce different lighting conditions from actual human hands. The multi-camera setup lacks a variety of lighting, grasping objects, and backgrounds. To tackle these problems, domain adaptation is a promising solution that can transfer the knowledge of the network trained on source data to unlabeled target data. Jiang et al. proposed an adversarial domain adaptation for human and hand keypoint regression, optimizing the discrepancy between regressors [33]. Additionally, self-training adaptation methods have been studied in the keypoint regression of animals [11], humans [66], and objects [77]. Unlike these prior works, we incorporate confidence estimation into a self-training method based on consistency training for keypoint regression.

Hand segmentation is the task of segmenting pixel-level hand masks in a given image. CNN-based segmentation networks [65,3,35] are popularly used. The task can be jointly trained with hand keypoint regression because detecting hand regions guides to improve keypoint localization [68,76,27,49,16]. Since hand mask annotation is laborious as hand keypoint regression, a few domain adaptation methods with pseudo-labeling have been explored [7,53]. To reduce the effect of highly noisy pseudo-labels in the target domain, Cai *et al.* incorporate the uncertainty of pseudo-labels in model adaptation [7], and Ohkawa *et al.* select confident pseudo-labels by the overlap of two predicted hand masks [53]. Unlike [7], we estimate the target confidence using two networks. Instead of using the estimated confidence for instance selection [53], we assign the confidence to weight the loss of consistency training.

Domain adaptation via self-training aims to learn unlabeled target data in a self-supervised learning manner. This approach can be divided into three categories. (i) Pseudo-labeling [15,59,81,53,7] learns unlabeled data with hard labels assigned by confidence thresholding from the output of a network. (ii) Entropy minimization [42,67,55] regularizes the conditional entropy of unlabeled data and increases the confidence of class probability. (iii) Consistency regularization [71,14,20] enforces regularization so that the prediction on unlabeled data is invariant under data perturbation. We choose to leverage this consistencybased method for our task because it works for various tasks [45,41,52] and the first two approaches cannot be directly applied. Similar to our work, Yang *et al.* [73] enforce the consistency for two different views and modalities in hand keypoint regression. Mean teacher [64] provides teacher-student training with consistency regularization, which regularizes a teacher network by a student's weights and avoids over-fitting to incorrect predictions. Unlike [73], we propose to integrate confidence estimation into the consistency training and adopt the teacher-student scheme with two networks. To encourage the two networks to have different representations, we propose a distillation-based update rule instead of updating the teacher with the exponential moving average [64].

3 Proposed Method

In this section, we present our proposed self-training domain adaptation with confidence estimation for adapting hand keypoint regression and hand segmentation. We first present our problem formulation and network initialization with supervised learning from source data. We then introduce our proposed modules: (1) geometric augmentation consistency, (2) confidence weighting by using two networks, and (3) teacher-student update via knowledge distillation. As shown in Fig. 2, our adaptation is done with two different networks (teachers) for confidence estimation and another network (student) for self-training of both tasks.

Problem formulation. Given labeled images from one source domain and unlabeled images from another target domain, we aim to jointly estimate hand keypoint coordinates and pixel-level hand masks on the target domain. We have a source image $\boldsymbol{x}_{\rm s}$ drawn from a set $X_{\rm s} \subset \mathbb{R}^{H \times W \times 3}$, its corresponding labels $(\boldsymbol{y}_{\rm s}^{\rm p}, \boldsymbol{y}_{\rm s}^{\rm m})$, and a target image $\boldsymbol{x}_{\rm t}$ drawn from a set $X_{\rm t} \subset \mathbb{R}^{H \times W \times 3}$. The pose label $\boldsymbol{y}_{\rm s}^{\rm p}$ consists of the 2D keypoint coordinates of 21 hand joints obtained from a set $Y_{\rm s}^{\rm p} \subset \mathbb{R}^{21 \times 2}$, while the mask label $\boldsymbol{y}_{\rm s}^{\rm m}$ denotes a binary mask obtained from $Y_{\rm s}^{\rm m} \subset (0,1)^{H \times W}$. A network parameterized by $\boldsymbol{\theta}$ learns the mappings $f^k(x; \boldsymbol{\theta}) : X \to Y^k$ where $k \in \{p, m\}$ represents the indicator for both tasks.

Initialization with supervised learning. To initialize networks used in our adaptation, we train the network f on the labeled source data following multi-task learning. Given the labeled dataset (X_s, Y_s) and the network θ , a supervised loss function is defined as

$$\mathcal{L}_{\text{task}}\left(\boldsymbol{\theta}, X_{\text{s}}, Y_{\text{s}}\right) = \sum_{k} \lambda^{k} \mathbb{E}_{\left(\boldsymbol{x}_{\text{s}}, \boldsymbol{y}_{\text{s}}^{k}\right) \sim \left(X_{\text{s}}, Y_{\text{s}}^{k}\right)} \left[\mathcal{L}^{k}(\boldsymbol{p}_{s}^{k}, \boldsymbol{y}_{\text{s}}^{k})\right], \tag{1}$$

where $Y_{\rm s} = \{Y_{\rm s}^{\rm p}, Y_{\rm s}^{\rm m}\}$ and $\boldsymbol{p}_{\rm s}^{k} = f^{k}(\boldsymbol{x}_{\rm s}; \boldsymbol{\theta})$. $\mathcal{L}^{k}(\cdot, \cdot) : Y^{k} \times Y^{k} \to \mathbb{R}^{+}$ is a loss function in each task and λ^{k} is a hyperparameter to balance the two tasks. We use a smooth L1 loss [32,58] as $\mathcal{L}^{\rm p}$ and a binary cross-entropy loss as $\mathcal{L}^{\rm m}$.

6 T. Ohkawa et al.

3.1 Geometric Augmentation Consistency

Inspired by semi-supervised learning using hand keypoint consistency [73], we advance a unified training with consistency for both hand keypoint regression and hand segmentation. We expect that joint adaption of both tasks will allow one task to provide useful cues to the other task in consistency training, as studied in supervised or weakly-supervised learning setups [68,76,27,49,16]. We design consistency training by predicting the location of hand keypoints and hand pixels in a given geometrically transformed image, including rotation and transition. This consistency under geometric augmentation encourages the network to learn against positional bias in the target domain, which helps capture the hand structure related to poses and regions. Specifically, given a paired augmentation function $(T_x, T_y^k) \sim \mathcal{T}$ for an image and an label, we generate the prediction on the target images $\mathbf{p}_t^k = f^k(\mathbf{x}_t; \boldsymbol{\theta})$ and the augmented target images $\mathbf{p}_{t,aug}^k = f^k(T_x(\mathbf{x}_t); \boldsymbol{\theta})$. We define the loss function of geometric augmentation consistency (GAC) \mathcal{L}_{gac} between $\mathbf{p}_{t,aug}^k$ and $T_y^k(\mathbf{p}_t)$ as

$$\mathcal{L}_{\text{gac}}\left(\boldsymbol{\theta}, X_{t}, \mathcal{T}\right) = \mathbb{E}_{\boldsymbol{x}_{t}, (T_{x}, T_{y}^{\text{p}}, T_{y}^{\text{m}})} \left[\sum_{k \in \{\text{p,m}\}} \tilde{\lambda}^{k} \tilde{\mathcal{L}}^{k} \left(\boldsymbol{p}_{t, \text{aug}}^{k}, T_{y}^{k}(\boldsymbol{p}_{t}^{k})\right) \right].$$
(2)

To correct the augmented prediction $\boldsymbol{p}_{t,aug}^k$ by $T_y^k(\boldsymbol{p}_t)$, we stop the gradient update for \boldsymbol{p}_t^k , which can be viewed as the supervision to $\boldsymbol{p}_{t,aug}^k$. We use the smooth L1 loss (see Equation 1) as $\tilde{\mathcal{L}}^p$ and a mean squared error as $\tilde{\mathcal{L}}^m$. We introduce $\tilde{\lambda}^k$ as a hyperparameter to control the balance of the two tasks. The augmentation set \mathcal{T} contains the geometric augmentation and photometric augmentation, such as color jitter and blurring. We set $T_y(\cdot)$ to align geometric information to the augmented input $T_x(\boldsymbol{x}_t)$. For example, we apply rotation $T_y(\cdot)$ to the outputs \boldsymbol{p}_t^k with the same degree of rotation $T_x(\cdot)$ to the input \boldsymbol{x}_t .

3.2 Confidence Estimation by Two Separate Networks

Since the target predictions are not always reliable, we aim to incorporate the estimated confidence weight for each target instance into the consistency training. In Equation 2, the generated outputs p_t^k that is the supervision to $p_{t,aug}^k$ may be unstable and noisy due to the domain gap between source and target domains. Due to that, the network trained with the consistency readily overfits to the incorrect supervision p_t^k , which is known as confirmation bias [2,64]. To reduce the bias, it is crucial to assign a low importance (confidence) weight to the consistency training with the incorrect supervision. This enables the network to learn primarily from reliable supervision while avoiding being biased to such erroneous predictions. In classification tasks, predicted class probability can serve as the confidence, while these measures are not trivially defined and available in regression tasks. To estimate the confidence of keypoint predictions, Yang *et al.* [73] measure the confidence of 3D hand keypoints by the distance to the fitted 3D hand template, but the hand template fitting is an ill-posed problem



Fig. 2: Method overview. Left: Student training with confidence-aware geometric augmentation consistency. The student learns from the consistency between its prediction and the two teachers' predictions. The training is weighted by the target confidence computed by the divergence of both teachers. **Right**: Teacher training with knowledge distillation. Each teacher independently learns to match the student's predictions. The task index k is omitted for simplicity.

for 2D hands and is not applicable to hand segmentation. Dropout [22,7,8] is a generic way of estimating uncertainty (confidence), calculated by the variance of multiple stochastic forwards. However, the estimated confidence is biased to the current state of the training network because the training and confidence estimation are done by a single network. When the training network works poorly, the confidence estimation becomes readily unreliable.

To perform reliable confidence estimation for both tasks, we propose a confidence measure by computing the divergence of two predictions. Specifically, we introduce two networks (*a.k.a.*, teachers) for the confidence estimation and the estimated confidence is used to train another network (*a.k.a.*, student) for the consistency training. The architecture of the teachers is identical, yet they have different learning parameters. We observe that when the divergence of the two predictions from the teachers for a target instance is high, the predictions of both networks become unstable. In contrast, a lower divergence indicates that the two teacher networks predict stably and agree on their predictions. Thus, we use the divergence for representing the target confidence. Given the teachers θ^{tch1} , θ^{tch2} , we define a disagreement measure $\ell_{disagree}$ to compute the divergence as

$$\ell_{\text{disagree}}\left(\boldsymbol{\theta}^{\text{tch1}}, \boldsymbol{\theta}^{\text{tch2}}, \boldsymbol{x}_{\text{t}}\right) = \sum_{k \in \{\text{p,m}\}} \tilde{\lambda}^{k} \tilde{\mathcal{L}}^{k}(\boldsymbol{p}_{\text{t1}}^{k}, \boldsymbol{p}_{\text{t2}}^{k}), \quad (3)$$

where $p_{t1}^k = f^k(\boldsymbol{x}_t; \boldsymbol{\theta}^{tch1})$ and $p_{t2}^k = f^k(\boldsymbol{x}_t; \boldsymbol{\theta}^{tch2})$.

As a proof of concept, we visualize the correlation between the disagreement measure and a validation score averaged over evaluation metrics of the two tasks (PCK and IoU) in Fig. 3. We compute the score between the ensemble of the teachers' predictions $p_{ens}^k = (p_{t1}^k + p_{t2}^k)/2$ and its ground truth in the validation set on HO3D [29]. The instances with a small disagreement measure tend to

8 T. Ohkawa et al.



Fig. 3: The correlation between a disagreement measure and task scores. Target instances with smaller disagreement values between the two teacher networks tend to have higher task scores.

have high validation scores. In contrast, the instances with a high disagreement measure entail false predictions, e.g., detecting the hand-held object as a hand joint and hand class. When the disagreement measure was high at the bottom of Fig. 3, we found that both predictions were particularly unstable on the keypoints of the ring finger (yellow). This study shows that the disagreement measure can represent the correctness of the target predictions.

With the disagreement measure ℓ_{disagree} , we define a confidence weight $w_t \in [0, 1]$ for assigning importance to the consistency training. We compute the weight w_t as $w_t = 2 \left(1 - \text{sigm} \left(\lambda_d \ell_{\text{disagree}} \left(\boldsymbol{\theta}^{\text{tch1}}, \boldsymbol{\theta}^{\text{tch2}}, \boldsymbol{x}_t\right)\right)\right)$ where w_t is a normalized disagreement measure with sign inversion, $\text{sigm}(\cdot)$ denotes a sigmoid function, and λ_d controls the scale of the measure. With the confidence weight w_t , we enforce the consistency training between the student's prediction on the augmented target images $\boldsymbol{p}_{s,\text{aug}}^k$ and the ensemble of the two teachers' predictions $\boldsymbol{p}_{\text{ens}}^k$. Our proposed loss function of confidence-aware geometric augmentation consistency (C-GAC) $\mathcal{L}_{\text{cgac}}$ for the student $\boldsymbol{\theta}^{\text{stu}}$ is formulated as

$$\mathcal{L}_{\text{cgac}}\left(\boldsymbol{\theta}^{\text{stu}}, \boldsymbol{\theta}^{\text{tch1}}, \boldsymbol{\theta}^{\text{tch2}}, X_{\text{t}}, \mathcal{T}\right) = \mathbb{E}_{\boldsymbol{x}_{\text{t}}, (T_{\text{x}}, T_{\text{y}}^{\text{p}}, T_{\text{y}}^{\text{m}})} \left[w_{\text{t}} \sum_{k \in \{\text{p,m}\}} \tilde{\lambda}^{k} \tilde{\mathcal{L}}^{k} \left(\boldsymbol{p}_{\text{s,aug}}^{k}, T_{\text{y}}^{k} (\boldsymbol{p}_{\text{ens}}^{k}) \right) \right]$$

$$(4)$$

where $\boldsymbol{p}_{\mathrm{s,aug}}^k = f^k(T_{\mathrm{x}}(\boldsymbol{x}_{\mathrm{t}}); \boldsymbol{\theta}^{\mathrm{stu}})$. Following [64,54], we design the student prediction $\boldsymbol{p}_{\mathrm{s,aug}}^k$ to be supervised by the teachers. We generate the teachers' prediction by doing ensemble $\boldsymbol{p}_{\mathrm{ens}}^k$, which is better than the prediction of either teacher.

3.3 Teacher-Student Update by Knowledge Distillation

In addition to the student's training, we formulate an update rule for the two teacher networks by using knowledge distillation. Since ℓ_{disagree} would not work

if the two teachers had the same output values, we aim to learn two teachers that have different representations yet keep high task performance as co-training works [4,57,15,59]. In a prior teacher-student update, Tarvainen *et al.* [64] found that the teacher's update by an exponential moving average (EMA), which averages the student's weights iteratively, makes the teacher's learning more slowly and mitigates the confirmation bias as discussed in Section 3.2. While this EMA-based teacher-student framework is widely used in various domain adaptation tasks [19,9,38,72,25], naively applying the EMA rule to the two teachers would produce exactly the same weights for both networks.

To prevent this, we propose independent knowledge distillation for building two different teachers. The distillation matches the teacher-student predictions in the output level. To let both networks have different parameters, we train the teachers from different mini-batches and using stochastic augmentation as

$$\mathcal{L}_{\text{distill}}\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{\text{stu}}, X_{\text{t}}, \mathcal{T}\right) = \mathbb{E}_{\boldsymbol{x}_{\text{t}}, T_{\text{x}}}\left[\sum_{k \in \{\text{p,m}\}} \tilde{\lambda}^{k} \tilde{\mathcal{L}}^{k}(\boldsymbol{p}_{\text{t,aug}}^{k}, \boldsymbol{p}_{\text{s,aug}}^{k})\right], \quad (5)$$

where $\boldsymbol{\theta} \in \{\boldsymbol{\theta}^{\text{tch1}}, \boldsymbol{\theta}^{\text{tch2}}\}, \boldsymbol{p}_{\text{t,aug}}^k = f^k(T_x(\boldsymbol{x}_{\text{t}}); \boldsymbol{\theta}), \text{ and } \boldsymbol{p}_{\text{s,aug}}^k = f^k(T_x(\boldsymbol{x}_{\text{t}}); \boldsymbol{\theta}^{\text{stu}}).$ The distillation loss $\mathcal{L}_{\text{distill}}$ is used for updating the teacher networks only. This helps the teachers to adapt to the target domain more carefully than the student and avoid falling into exactly the same predictions on a target instance.

3.4 Overall Objectives

Overall, the objective of the student's training consists of the supervised loss (Equation 1) from the source domain and the self-training with confidence-aware geometric augmentation consistency (Equation 4) in the target domain as

$$\min_{\boldsymbol{\theta}^{\text{stu}}} \mathcal{L}_{\text{task}} \left(\boldsymbol{\theta}^{\text{stu}}, X_{\text{s}}, Y_{\text{s}} \right) + \mathcal{L}_{\text{cgac}} \left(\boldsymbol{\theta}^{\text{stu}}, \boldsymbol{\theta}^{\text{tch1}}, \boldsymbol{\theta}^{\text{tch2}}, X_{t}, \mathcal{T} \right).$$
(6)

The two teachers are asynchronously trained with the distillation loss (Equation 5) in the target domain, which is formulated as

$$\min_{\boldsymbol{\theta}} \mathcal{L}_{\text{distill}} \left(\boldsymbol{\theta}, \boldsymbol{\theta}^{\text{stu}}, X_{\text{t}}, \mathcal{T} \right), \tag{7}$$

where $\theta \in {\{\theta^{tch1}, \theta^{tch2}\}}$. Since the teachers are updated carefully and can perform better than the student, we use the ensemble of the two teachers' predictions for a final output in inference.

4 Experiments

In this section, we first present our experimental datasets and implementation details and then provide quantitative and qualitative results along with the ablation studies. We analyze our proposed method by comparing it with several existing methods in three different domain adaptation settings. We also show qualitative results by applying our method to in-the-wild egocentric videos.

10 T. Ohkawa et al.

4.1 Experiment Setup

Datasets. We experimented with several hand datasets including a variety of hand-object interactions, the annotation of 2D hand keypoints, and hand masks as follows. We adopted **DexYCB** [13] dataset as our source dataset since it contains a large amount of training images, their corresponding labels, and natural hand-object interactions. We chose to use the following datasets as our target datasets: **HO3D** [29] captured in different environments with the same YCB objects [10] as the source dataset, **HanCo** [78] captured in a multi-camera studio and generated with synthesized backgrounds, and **FPHA** [24] captured by a first-person view. We also used **Ego4D** [28] to verify the effectiveness of our method in real-world scenarios. During training, we used cropped images of the hand regions from the original images as input.

Implementation details. Our teacher-student networks share an identical network architecture, which consists of a unified feature extractor and taskspecific branches for hand keypoint regression and hand segmentation. For training our student network, we used the Adam optimizer [36] with a learning rate of 10^{-5} , while the learning rate of the teacher networks was set to 5×10^{-6} . We set the hyperparameters ($\lambda^{\rm p}(=\tilde{\lambda}^{\rm p}), \lambda^{\rm m}, \tilde{\lambda}^{\rm m}, \lambda_{\rm d}$) to ($10^7, 10^2, 5, 0.5$). Since both taskspecific branches have different training speeds, we began our adaptation with the backbone and keypoint regression branch. We then trained all sub-networks, including the hand segmentation branch. We report the percentage of correct keypoints (PCK) and the mean joint position error (MPE) for hand keypoint regression, and the intersection over union (IoU) for hand segmentation.

Baseline methods. We compared quantitative performance with the following methods. Source only denotes the network trained on the source dataset without any adaptation. To compare with another adaptation approach with adversarial training, we trained DANN [23] that aligns marginal feature distributions between domains, and RegDA [33] with an adversarial regressor that optimizes domain disparity. In addition, we implemented several self-training adaptation methods by replacing pseudo-labeling with the consistency training. GAC is a simple baseline with the consistency training updated by Equation 2. GAC + UMA [7] is a GAC method with confidence estimation by Dropout [22]. GAC + CPL [53] is a GAC method with confident instance selection using the agreement with another network. GAC + MT [64] is a GAC method with the single-teacher-single-student architecture using EMA for the teacher update. Target only indicates the network trained on the target dataset with labels, which shows an empirical performance upper bound.

Our method. We denote our full method as C-GAC introduced in Section 3.4. As an ablation study, we present a variant of the proposed method as GAC-Distill with a teacher-student pair, which is updated by the consistency training (Equation 2) and the distillation loss (Equation 5). GAC-Distill is different from GAC + MT only in the way of the teacher update.

Table 1: **DexYCB** [13] \rightarrow **HO3D** [29]. We report PCK (%) and MPE (px) for hand keypoint regression and IoU (%) for hand segmentation. Each score format of *val* / *test* indicates the validation and test scores. Red and blue letters indicate the best and second best values.

Method	2D Pose		Seg	2D Pose + Seg		
	$PCK \uparrow (\%)$	$\mathrm{MPE}\downarrow(\mathrm{px})$	$\mathrm{IoU}\uparrow(\%)$	Avg. \uparrow (%)		
Source only	42.8/33.5	15.39/19.32	57.9/49.1	50.3/41.3		
DANN [23]	49.0/46.8	12.39/13.39	52.8/54.7	50.9/50.8		
RegDA [33]	48.8/48.2	12.50/12.64	55.7/55.3	52.2/51.7		
GAC	47.6/47.4	12.47/12.54	58.0/56.9	52.8/52.2		
GAC + UMA [7]	47.1/45.3	12.97/13.51	58.0/55.0	52.5/50.2		
GAC + CPL [53]	48.1/48.1	12.74/12.61	57.2/55.6	52.7/51.8		
GAC + MT [64]	45.5/44.4	13.65/14.05	54.8/52.3	50.2/48.3		
GAC-Distill (Ours)	49.9/50.4	11.98/11.51	60.7/60.6	55.3/55.5		
C-GAC (Ours-Full)	50.3/51.1	11.89/11.22	60.9/60.3	55.6/55.7		
Target only	55.1/58.6	11.00/9.29	68.2/66.1	61.7/62.4		

4.2 Quantitative Results

We show the results of three adaptation settings: DexYCB \rightarrow {HO3D, HanCo, FPHA} in Tables 1 and 2. We then provide detailed comparisons of our method.

DexYCB \rightarrow **HO3D**. Table 1 shows the results of the adaptation from DexYCB to HO3D where the grasping objects are overlapped. The baseline of the consistency training (**GAC**) was effective in learning target images in both tasks. Our proposed method (**C-GAC**) improved by 5.3/14.4 in the average task score from the source-only performance. The method also outperformed all comparison methods and achieved close performance to the upper bound.

 $DexYCB \rightarrow HanCo$. Table 2 shows the results of the adaptation from DexYCB to HanCo across laboratory setups. The source-only network less generalized to the target domain because the HanCo has diverse backgrounds, while **GAC** succeeded in adapting up to 47.4/47.9 in the average score. Our method **C-GAC** showed further improved results in hand keypoint regression.

 $\text{DexYCB} \rightarrow \text{FPHA}$. Table 2 also shows the results of the adaptation from DexYCB to FPHA, which captures egocentric users' activities. Since hand markers and in-the-wild target environments cause large appearance gaps, the source-only performance performed the most poorly among the three adaption settings. In this challenging setting, **RegDA** and **GAC** + **UMA** performed well for hand segmentation, while their performance on hand keypoint regression was inferior to the **GAC** baseline. Our method **C-GAC** further improved than the **GAC** method in the MPE and IoU metrics and exhibited stability in adaptation training among the comparison methods.

Comparison to different confidence estimation methods. We compare the results with existing confidence estimation methods. GAC + UMAand GAC + CPL estimate the confidence of target predictions by computing the variance of multiple stochastic forwards and the task scores between a training network and an auxiliary network, respectively. GAC + UMA performed

Table 2: **DexYCB** $[13] \rightarrow \{\text{HanCo} [78], \text{FPHA} [24]\}$. We report PCK (%) and MPE (px) for hand keypoint regression and IoU (%) for hand segmentation. We show the validation and test results on HanCo and the validation results on FPHA. Red and blue letters indicate the best and second best values.

	$DexYCB \rightarrow HanCo$				$DexYCB \rightarrow FPHA$			
Method	2D	Pose	Seg	[$\overline{]}$ $\overline{2D}$	Pose	Seg	
	$PCK \uparrow (\%)$	$\mathrm{MPE}\downarrow(\mathrm{px})$	$\mathrm{IoU}\uparrow(\%)$	Avg. \uparrow (%)	PCK	MPE	IoU	Avg.
Source only	26.0/27.3	21.82/21.48	41.8/41.4	33.9/34.3	14.0	31.32	24.8	19.4
DANN [23]	32.3/33.0	19.99/19.82	56.3/56.9	44.3/45.0	24.4	25.79	28.4	26.4
RegDA [33]	33.0/33.6	19.51/19.44	57.8/58.4	45.4/46.0	23.7	24.27	41.7	32.7
GAC	36.6/37.1	16.63/16.59	58.1/58.8	47.4/47.9	37.2	17.02	33.3	35.3
GAC + UMA [7]	35.1/35.6	17.51/17.48	57.1/57.7	46.1/46.6	36.8	17.29	39.2	38.0
GAC + CPL [53]	32.7/33.5	19.85/19.62	55.8/56.4	44.2/45.0	25.7	24.99	32.7	29.2
GAC + MT [64]	33.2/33.8	18.93/18.83	54.3/55.1	43.8/44.4	31.3	20.81	38.4	34.9
GAC-Distill (Ours)	38.8/39.5	16.06/15.97	57.5/57.7	48.1/48.6	36.8	15.99	35.5	36.1
C-GAC (Ours-Full)	39.2/39.9	15.83/15.74	58.2/58.6	48.7/49.2	37.2	15.36	37.7	37.4
Target only	76.8/77.3	4.91/4.80	75.9/76.1	76.3/76.7	63.3	8.11	-	-

effectively on DexYCB \rightarrow FPHA, whereas the performance gain was thin in the other settings compared to **GAC**. **GAC** + **CPL** worked well for keypoint regression on DexYCB \rightarrow HO3D, but it cannot address the other settings with a large domain gap well since the prediction of the auxiliary network became unstable. Although these prior methods had different disadvantages depending on the settings, our method **C-GAC** using the divergence of the two teachers for confidence estimation performed stably in the three settings.

Comparison to standard teacher-student update. We compare our teacher update with the update with an exponential moving average (EMA) [64]. The EMA-based update (**GAC-MT**) degraded the performance from the source only in hand segmentation in Table 1. This suggests that the EMA update can be sensitive to the task. In contrast, our method **GAC-Distill** matching the teacher-student predictions in the output level did not produce such performance degeneration and worked more stably.

Comparison to adversarial adaptation methods. We compared our method with another major adaptation approach with adversarial training. In Tables 1 and 2, the performance of **DANN** and **RegDA** was mostly worse than the consistency-based baseline **GAC**. We found that instead of matching features between both domains [23,33], directly learning target images by the consistency training was critical in the adaptation of our tasks.

Comparison to an off-the-shelf hand pose estimator. We tested the generalization ability of an open-source library for pose estimation: **Open-Pose** [31]. It resulted in 15.75/12.72, 18.31/18.42, and 29.02 in the MPE on HO3D, HanCo, and FPHA, respectively. Since it is built on multiple source datasets [34,1,44], the baseline showed higher generalization than the source-only network. However, the performance did not exceed our proposed method in the MPE. This shows that generalizing hand keypoint regression to other



Fig. 4: Qualitative results. We show qualitative examples of the source-only network (top), the Ours-Full method (middle), and ground truth (bottom) on HO3D [29], HanCo [78], FPHA [24], and Ego4D [28] without ground truth.

datasets is still challenging, and our adaptation framework supports improving target performance.

4.3 Qualitative Results

We show the qualitative results of hand keypoint regression and hand segmentation in Fig. 4. When hands are occluded in HO3D and FPHA or the backgrounds are diverse in HanCo, the keypoint prediction of the source only (top) represented infeasible hand poses and hand segmentation was too noisy or missing. However, our method **C-GAC** (middle) corrected the hand keypoint errors and improved to localize hand regions. Hand segmentation in FPHA was still noisy because visible white markers obstructed hand appearance. We can also see distinct improvements in the Ego4D dataset. We provide additional qualitative analysis in adaptation to the Ego4D beyond countries, cultures, ages, indoors/outdoors, and performing tasks with hands in our supplementary material.

4.4 Ablation Studies

Effect of confidence estimation. To confirm the effect of our proposed confidence estimation, we compare our full method C-GAC and our ablation model GAC-Distill without the confidence weighting. In Tables 1 and 2, while GAC-Distill mostly surpassed the comparison methods in most cases, C-GAC showed further performance gain in all three adaptation settings.

Multi-task vs. single-task adaptation. We studied the effect of our multitask adaptation compared with single-task adaptation on DexYCB \rightarrow HO3D. The single-task adaptation results are 50.1/51.0 in the PCK and 58.2/57.7 in the IoU. Compared to Table 1, our method in the multi-task setting improved by 2.7/2.6 over the single-task adaption in hand segmentation while it provided marginal gain in hand keypoint regression. This shows that the adaptation of hand keypoint regression helps to localize hand regions in the target domain.



Fig. 5: Visualization of bone length distributions. We show the distributions of the bone length between hand joints, namely, Wrist, metacarpophalangeal (MCP), proximal interphalangeal (PIP), distal interphalangeal (DIP), and fingertip (TIP). Using kernel density estimation, we plotted the density of the bone length for the predictions of the source only, the Ours-Full method, and ground truth on test data of HO3D [29].

Bone length distributions. To study our adaptation results in each hand joint, we show the distributions of bone length between hand joints in Fig. 5. In Wrist-MCP, PIP-DIP, and DIP-TIP, the distribution of the source-only prediction on target images (blue) was far from that of the target ground truth (green), whereas our method (orange) improved to approximate the target distribution (green). In MCP-PIP, we could not observe such clear differences because the source-only model already represented the target distribution well. This indicates that our method improved to learn hand structure near the palm and fingertips.

5 Conclusion

In this work, we tackled the problem of joint domain adaptation of hand keypoint regression and hand segmentation. Our proposed method consists of the self-training with geometric augmentation consistency, confidence weighting by the two teacher networks, and the teacher-student update by knowledge distillation. The consistency training under geometric augmentation served to learn the unlabeled target images for both tasks. The divergence of the predictions from two teacher networks could represent the confidence of each target instance, which enables the student network to learn from reliable target predictions. The distillation-based teacher-student update guided the teachers to learn from the student carefully and mitigated over-fitting to the noisy predictions. Our method delivered state-of-the-art performance on the three adaptation setups. It also showed improved qualitative results in the real-world egocentric videos.

Acknowledgments

This work was supported by JST ACT-X Grant Number JPMJAX2007, JSPS Research Fellowships for Young Scientists, JST AIP Acceleration Research Grant Number JPMJCR20U1, and JSPS KAKENHI Grant Number JP20H04205, Japan. This work was also supported in part by a hardware donation from Yu Darvish.

References

- M. Andriluka, L. Pishchulin, P. V. Gehler, and B. Schiele. 2D human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3686–3693, 2014. 12
- E. Arazo, D. Ortego, P. Albert, N. E. O'Connor, and K. McGuinness. Pseudolabeling and confirmation bias in deep semi-supervised learning. In *IEEE Interna*tional Joint Conference on Neural Networks (IJCNN), pages 1–8, 2020. 6
- G. Benitez-Garcia, L. Prudente-Tixteco, L. C. Castro-Madrid, R. Toscano-Medina, J. Olivares-Mercado, G. Sanchez-Perez, and L. J. G. Villalba. Improving real-time hand gesture recognition with semantic segmentation. *Sensors*, 21(2), 2021.
- A. Blum and T. M. Mitchell. Combining labeled and unlabeled data with cotraining. In Proceedings of the ACM Annual Conference on Computational Learning Theory (COLT), pages 92–100, 1998.
- A. Boukhayma, R. A. Bem, and P. H. S. Torr. 3d hand shape and pose from images in the wild. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), pages 10843–10852, 2019.
- S. Brahmbhatt, C. Tang, C. D. Twigg, C. C. Kemp, and J. Hays. ContactPose: A dataset of grasps with object contact and hand pose. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 361–378, 2020.
- M. Cai, E. Lu, and Y. Sato. Generalizing hand segmentation in egocentric videos with uncertainty-guided model adaptation. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), pages 14380–14389, 2020.
 2, 3, 4, 7, 10, 11, 12
- M. Cai, M. Luo, X. Zhong, and H. Chen. Uncertainty-aware model adaptation for unsupervised cross-domain object detection. CoRR, abs/2108.12612, 2021.
- Q. Cai, Y. Pan, C.-W. Ngo, X. Tian, L. Duan, and T. Yao. Exploring object relation in mean teacher for cross-domain detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11457– 11466, 2019.
- B. Çalli, A. Walsman, A. Singh, S. S. Srinivasa, P. Abbeel, and A. M. Dollar. Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set. *IEEE Robotics Automation Magazine*, 22(3):36–52, 2015. 10
- J. Cao, H. Tang, H. Fang, X. Shen, Y.-W. Tai, and C. Lu. Cross-domain adaptation for animal pose estimation. In *Proceedings of the IEEE International Conference* on Computer Vision (ICCV), pages 9497–9506, 2019.
- Z. Cao, I. Radosavovic, A. Kanazawa, and J. Malik. Reconstructing hand-object interactions in the wild. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 12417–12426, 2021.
 Y.-W. Chao, W. Yang, Y. Xiang, P. Molchanov, A. Handa, J. Tremblay, Y. S.
- Y.-W. Chao, W. Yang, Y. Xiang, P. Molchanov, A. Handa, J. Tremblay, Y. S. Narang, K. Van Wyk, U. Iqbal, S. Birchfield, J. Kautz, and D. Fox. DexYCB: A benchmark for capturing hand grasping of objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9044–9053, 2021. 2, 3, 4, 10, 11, 12
- C.-H. Chen, A. Tyagi, A. Agrawal, D. Drover, R. MV, S. Stojanov, and J. M. Rehg. Unsupervised 3d pose estimation with geometric self-supervision. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 5714–5724, 2019. 5
- M. Chen, K. Q. Weinberger, and J. Blitzer. Co-training for domain adaptation. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), pages 2456-2464, 2011. 2, 4, 9

- 16 T. Ohkawa et al.
- X. Chen, G. Wang, C. Zhang, T.-K. Kim, and X. Ji. Shpr-net: Deep semantic hand pose regression from point clouds. *IEEE Access*, 6:43425–43439, 2018. 3, 4, 6
- D. Damen, H. Doughty, G. M. Farinella, A. Furnari, J. Ma, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. Rescaling egocentric vision. *International Journal of Computer Vision (IJCV)*, early access, 2021.
- J. Deng, W. Li, Y. Chen, and L. Duan. Unbiased mean teacher for cross-domain object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 4091–4101, 2021.
- G. French, M. Mackiewicz, and M. H. Fisher. Self-ensembling for visual domain adaptation. In Proceedings of the International Conference on Learning Representations (ICLR), 2018.
- H. Fu, M. Gong, C. Wang, K. Batmanghelich, K. Zhang, and D. Tao. Geometryconsistent generative adversarial networks for one-sided unsupervised domain mapping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2427–2436, 2019. 5
- Q. Fu, X. Liu, and K. M. Kitani. Sequential decision-making for active object detection from hand. *CoRR*, abs/2110.11524, 2021.
- 22. Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the International Conference* on Machine Learning (ICML), pages 1050—-1059, 2016. 7, 10
- Y. Ganin and V. Lempitsky. Unsupervised domain adaptation by backpropagation. In Proceedings of the International Conference on Machine Learning (ICML), pages 1180–1189, 2015. 10, 11, 12
- 24. G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim. First-person hand action benchmark with RGB-D videos and 3D hand pose annotations. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 409–419, 2018. 2, 3, 4, 10, 12, 13
- 25. Y. Ge, D. Chen, and H. Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. 9
- O. Glauser, S. Wu, D. Panozzo, O. Hilliges, and O. Sorkine-Hornung. Interactive hand pose estimation using a stretch-sensing soft glove. ACM Transactions on Graphics, 38(4):41:1–41:15, 2019. 2, 4
- D. Goudie and A. Galata. 3D hand-object pose estimation from depth with convolutional neural networks. In *Proceedings of the IEEE International Conference* on Automatic Face & Gesture Recognition (FG), pages 406–413, 2017. 3, 4, 6
- K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Ham-28.burger, H. Jiang, M. Liu, X. Liu, M. Martin, T. Nagarajan, I. Radosavovic, S. K. Ramakrishnan, F. Ryan, J. Sharma, M. Wray, M.g Xu, E. Zhongcong Xu, C. Zhao, S. Bansal, D. Batra, V. Cartillier, S. Crane, T. Do, M. Doulaty, A. Erapalli, C. Feichtenhofer, A. Fragomeni, Q. Fu, C. Fuegen, A. Gebreselasie, C. Gonzalez, J. Hillis, X. Huang, Y. Huang, W. Jia, W. Khoo, J. Kolar, S. Kottur, A. Kumar, F. Landini, C. Li, Y. Li, Z. Li, K. Mangalam, R. Modhugu, J. Munro, T. Murrell, T. Nishiyasu, W. Price, P. R. Puentes, M. Ramazanova, L. Sari, K. Somasundaram, A. Southerland, Y. Sugano, R. Tao, M. Vo, Y. Wang, X. Wu, T. Yagi, Y. Zhu, P. Arbelaez, D. Crandall, D. Damen, G. M. Farinella, B. Ghanem, V. K. Ithapu, C. V. Jawahar, H. Joo, K. Kitani, H. Li, R. Newcombe, A. Oliva, H. Soo Park, J. M. Rehg, Y. Sato, J. Shi, M. Z. Shou, A. Torralba, Lo Torresani, M.i Yan, and J. Malik. Ego4D: Around the world in 3, 000 hours of egocentric video. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 18995–19012, 2022. 1, 2, 3, 10, 13

- S. Hampali, M. Rad, M. Oberweger, and V. Lepetit. Honnotate: A method for 3D annotation of hand and object poses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3196–3206, 2020. 3, 4, 7, 10, 11, 13, 14
- 30. Y. Hasson, G. Varol, D. Tzionas, I. Kalevatykh, M. J. Black, I. Laptev, and C. Schmid. Learning joint reconstruction of hands and manipulated objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11807–11816, 2019. 4
- 31. G. Hidalgo, Z. Cao, T. Simon, S.-E. Wei, Y. Raaj, H. Joo, and Y. Sheikh. Open-Pose. https://github.com/CMU-Perceptual-Computing-Lab/openpose. 12
- 32. W. Huang, P. Ren, J. Wang, Q. Qi, and H. Sun. AWR: Adaptive weighting regression for 3D hand pose estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 11061–11068, 2020. 5
- J. Jiang, Y. Ji, X. Wang, Y. Liu, J. Wang, and M. Long. Regressive domain adaptation for unsupervised keypoint detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6780–6789, 2021. 3, 4, 10, 11, 12
- 34. H. Joo, H. Liu, L. Tan, L. Gui, B. C. Nabbe, I. A. Matthews, T. Kanade, S. Nobuhara, and Y. Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3334–3342, 2015. 2, 4, 12
- 35. S. Kim, H. -G. Chi, X. Hu, A. Vegesana, and K. Ramani. First-person view hand segmentation of multi-modal hand activity video dataset. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020. 4
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Proceedings of the International Conference on Learning Representations (ICLR), 2014.
- K. Lee, A. Shrivastava, and H. Kacorri. Hand-priming in object localization for assistive egocentric vision. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 3422–3432, 2020.
- Y.-J. Li, X. Dai, C.-Y. Ma, Y.-C. Liu, K. Chen, B. Wu, Z. He, K. Kitani, and P. Vajda. Cross-domain object detection via adaptive self-training. *CoRR*, abs/2111.13216, 2021.
- 39. H. Liang, J.G. Yuan, D. Thalmann, and N. Magnenat-Thalmann. AR in hand: Egocentric palm pose tracking and gesture recognition for augmented reality applications. In *Proceedings of the ACM International Conference on Multimedia* (MM), pages 743–744, 2015. 1
- 40. J. Likitlersuang, E. R. Sumitro, T. Cao, R. J. Visée, S. Kalsi-Ryan, and J. Zariffa. Egocentric video: A new tool for capturing hand use of individuals with spinal cord injury at home. *Journal of Neuroengineering and Rehabilitation (JNER)*, 16(1):83, 2019. 1
- Y.-C. Liu, C.-Y. Ma, Z. He, C.-W. Kuo, K. Chen, P. Zhang, B. Wu, Z. Kira, and P. Vajda. Unbiased teacher for semi-supervised object detection. In *Proceedings* of the International Conference on Learning Representations (ICLR), 2021. 5
- 42. M. Long, H. Zhu, J. Wang, and M. I. Jorda. Unsupervised domain adaptation with residual transfer networks. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 136–144, 2016. 5
- 43. Y. Lu and W. W. Mayol-Cuevas. Understanding egocentric hand-object interactions from hand pose estimation. *CoRR*, abs/2109.14657, 2021. 4
- 44. R. McKee, D. McKee, D. Alexander, and E. Paillat. NZ sign language exercises. Deaf Studies Department of Victoria University of Wellington, http://www.victoria.ac.nz/llc/llc_resources/nzsl. 12

- 18 T. Ohkawa et al.
- 45. L. Melas-Kyriazi and A. K. Manrai. Pixmatch: Unsupervised domain adaptation via pixelwise consistency training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12435–12445, 2021. 5
- 46. G. Moon, S.-I. Yu, H. Wen, T. Shiratori, and K. M. Lee. InterHand2.6M: A dataset and baseline for 3D interacting hand pose estimation from a single RGB image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 548–564, 2020. 2, 4
- 47. F. Mueller, F. Bernard, O. Sotnychenko, D. Mehta, S. Sridhar, D. Casas, and C. Theobalt. GANerated Hands for real-time 3D hand tracking from monocular RGB. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 49–59, 2018. 4
- F. Mueller, D. Mehta, O. Sotnychenko, S. Sridhar, D. Casas, and C. Theobalt. Real-time hand tracking under occlusion from an egocentric RGB-D sensor. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1163–1172, 2017. 4
- N. Neverova, C. Wolf, F. Nebout, and G. W. Taylor. Hand pose estimation through semi-supervised and weakly-supervised learning. *Computer Vision and Image Un*derstanding, 164:56–67, 2017. 3, 4, 6
- A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 9912, pages 483–499, 2016.
- 51. T. Ohkawa, R. Furuta, and Y. Sato. Efficient annotation and learning for 3d hand pose estimation: A survey. *CoRR*, abs/2206.02257, 2022. 2, 4
- 52. T. Ohkawa, N. Inoue, H. Kataoka, and N. Inoue. Augmented cyclic consistency regularization for unpaired image-to-image translation. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, pages 362–369, 2020. 5
- 53. T. Ohkawa, T. Yagi, A. Hashimoto, Y. Ushiku, and Y. Sato. Foreground-aware stylization and consensus pseudo-labeling for domain adaptation of first-person hand segmentation. *IEEE Access*, 9:94644–94655, 2021. 2, 3, 4, 10, 11, 12
- 54. H. Pham, Z. Dai, Q. Xie, and Q. V. Le. Meta pseudo labels. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 11557–11568, 2021. 3, 8
- 55. V. Prabhu, S. Khare, D. Kartik, and J. Hoffman. SENTRY: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 8558–8567, 2021. 5
- C. Qian, X. Sun, Y. Wei, X. Tang, and J. Sun. Realtime and robust hand tracking from depth. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1106–1113, 2014.
- 57. S. Qiao, W. Shen, Z. Zhang, B. Wang, and A. L. Yuille. Deep co-training for semi-supervised image recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 11219, pages 142–159, 2018. 9
- P. Ren, H. Sun, Q. Qi, J. Wang, and W. Huang. SRN: stacked regression network for real-time 3D hand pose estimation. In *Proceedings of the British Machine* Vision Conference (BMVC), 2019. 5
- K. Saito, Y. Ushiku, and T. Harada. Asymmetric tri-training for unsupervised domain adaptation. In Proceedings of the International Conference on Machine Learning (ICML), pages 2988–2997, 2017. 4, 9
- N. Santavas, I. Kansizoglou, L. Bampis, E. G. Karakasis, and A. Gasteratos. Attention! A lightweight 2d hand pose estimation approach. *CoRR*, abs/2001.08047, 2020. 4

- T. Simon, H. Joo, I. Matthews, and Y. Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4645–4653, 2017.
- 62. S. Sridhar, F. Mueller, M. Zollhoefer, D. Casas, A. Oulasvirta, and C. Theobalt. Real-time joint tracking of a hand manipulating an object from RGB-D input. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 294–310, 2016. 4
- O. Taheri, N. Ghorbani, M. J. Black, and D. Tzionas. GRAB: A dataset of wholebody human grasping of objects. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 581–600, 2020.
- 64. A. Tarvainen and H. Valpola. Mean teachers are better role models: Weightaveraged consistency targets improve semi-supervised deep learning results. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. 3, 5, 6, 8, 9, 10, 11, 12
- A. Urooj and A. Borji. Analysis of hand segmentation in the wild. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 4710–4719, 2018.
- 66. L. O. Vasconcelos, M. Mancini, D. Boscaini, S. R. Bulò, B. Caputo, and E. Ricci. Shape consistent 2d keypoint estimation under domain shift. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, pages 8037–8044, 2020. 2, 4
- 67. T. H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pere. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of* the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2512–2521, 2019. 2, 5
- Y. Wang, C. Peng, and Y. Liu. Mask-pose cascaded CNN for 2d hand pose estimation from single color image. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 29(11):3258–3268, 2019. 3, 4, 6
- S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 4724–4732, 2016.
- M.-Y. Wu, P.-W. Ting, Y.-H. Tang, E. T. Chou, and L.-C. Fu. Hand pose estimation in object-interaction based on deep learning for virtual reality applications. *Journal of Visual Communication and Image Representation*, 70:102802, 04 2020.
- Q. Xie, Z. Dai, E. Hovy, T. Luong, and Q. Le. Unsupervised data augmentation for consistency training. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 5
- 72. L. Yan, B. Fan, S. Xiang, and C. Pan. CMT: cross mean teacher unsupervised domain adaptation for VHR image semantic segmentation. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. 9
- L. Yang, S. Chen, and A. Yao. Semihand: Semi-supervised hand pose estimation with consistency. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 11364–11373, 2021. 2, 5, 6
- 74. L. Yang, J. Li, W. Xu, Y. Diao, and C. Lu. Bihand: Recovering hand mesh with multi-stage bisected hourglass networks. In *Proceedings of the British Machine* Vision Conference (BMVC), 2020. 4
- 75. S. Yuan, Q. Ye, B. Stenger, S. Jain, and T.-K. Kim. BigHand2.2M benchmark: Hand pose dataset and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2605–2613, 2017. 2, 4

- 20 T. Ohkawa et al.
- 76. C. Zhang, G. Wang, X. Chen, P. Xie, and T. Yamasaki. Weakly supervised segmentation guided hand pose estimation during interaction with unknown objects. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, (ICASSP)*, pages 2673–2677, 2020. 3, 4, 6
- 77. X. Zhou, A. Karpur, C. Gan, L. Luo, and Q. Huang. Unsupervised domain adaptation for 3d keypoint estimation via view consistency. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 11216, pages 141–157, 2018. 2, 4
- C. Zimmermann, M. Argus, and T. Brox. Contrastive representation learning for hand shape estimation. CoRR, abs/2106.04324, 2021. 3, 10, 12, 13
- 79. C. Zimmermann and T. Brox. Learning to estimate 3D hand pose from single RGB images. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), pages 4913–4921, 2017. 4
- C. Zimmermann, D. Ceylan, J. Yang, B. Russell, M. J. Argus, and T. Brox. Frei-HAND: A dataset for markerless capture of hand pose and shape from single RGB images. In *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), pages 813–822, 2019. 2, 4
- Y. Zou, Z. Yu, B. V. Kumar, and J. Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 289–305, 2018.