

ScarceNet: Animal Pose Estimation with Scarce Annotations

Chen Li Gim Hee Lee

Department of Computer Science, National University of Singapore

lichen@u.nus.edu

gimhee.lee@comp.nus.edu.sg

Abstract

Animal pose estimation is an important but under-explored task due to the lack of labeled data. In this paper, we tackle the task of animal pose estimation with scarce annotations, where only a small set of labeled data and unlabeled images are available. At the core of the solution to this problem setting is the use of the unlabeled data to compensate for the lack of well-labeled animal pose data. To this end, we propose the ScarceNet, a pseudo label-based approach to generate artificial labels for the unlabeled images. The pseudo labels, which are generated with a model trained with the small set of labeled images, are generally noisy and can hurt the performance when directly used for training. To solve this problem, we first use a small-loss trick to select reliable pseudo labels. Although effective, the selection process is improvident since numerous high-loss samples are left unused. We further propose to identify reusable samples from the high-loss samples based on an agreement check. Pseudo labels are re-generated to provide supervision for those reusable samples. Lastly, we introduce a student-teacher framework to enforce a consistency constraint since there are still samples that are neither reliable nor reusable. By combining the reliable pseudo label selection with the reusable sample re-labeling and the consistency constraint, we can make full use of the unlabeled data. We evaluate our approach on the challenging AP-10K dataset, where our approach outperforms existing semi-supervised approaches by a large margin. We also test on the TigDog dataset, where our approach can achieve better performance than domain adaptation based approaches when only very few annotations are available. Our code is available at the project website ¹.

1. Introduction

The ability to understand animal behavior is fundamental to many applications, such as farming, ecology and surveillance. Animal pose estimation is a key step to un-

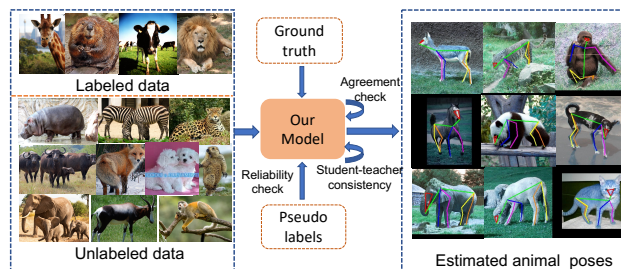


Figure 1. An illustration of our task. We aim to achieve accurate animal pose estimation with a small set of labeled images as well as unlabeled images.

derstand animal behavior and has attracted increasing attention in recent years [5, 16, 20, 23, 33]. Although great success is achieved for human pose estimation with the development of sophisticated deep learning models, these techniques cannot be directly used for animals due to the lack of labeled animal pose data. Existing works overcome this problem by learning from human pose data [5] or synthetic animal images [16, 20, 23]. However, there is a large domain gap between the real and synthetic (human) data. For example, the synthetic animal images in [20], which are generated from CAD models, only exhibit limited pose, appearance and background variations. As a result, the model trained with the synthetic data may not adapt well to real images, especially for images with crowded scene or self-occlusion. Moreover, the generation of synthetic images is a tedious process that also requires expert knowledge. The above-mentioned problems lead us to ask whether we can achieve accurate animal pose estimation with minimal effort for annotating? To answer this question, we focus on how to learn from scarce labeled data for animal pose estimation. As shown in Fig. 1, we aim to achieve accurate animal pose estimation with only a small set of labeled images and unlabeled images.

The data scarcity problem is solved by semi-supervised learning (SSL) for the classification task [3, 25, 28, 34]. One powerful class of SSL is pseudo labeling (PL), where artificial labels generated from a pretrained model are used together with labeled data to train a model. Impressive

¹<https://github.com/chaneyddtt/ScarceNet>

results have been achieved by applying PL [5, 16, 20] to domain adaptation based animal pose estimation. However, the PL-based methods suffer from confirmation bias when the network starts to output incorrect predictions with high confidence, which we refer as noisy pseudo labels. To solve this noisy pseudo label problem, we use the small-loss trick [9, 16] to select reliable pseudo labels, *i.e.* samples with small loss to update the network. This criteria is based on the memorization effect [2] of deep networks where the network learn from clean and easy instances first before eventually overfitting to noisy ones. Consequently, clean samples exhibit smaller loss compared to noisy ones at the beginning of training. Although the small-loss trick is effective in selecting reliable pseudo labels, it is wasteful because numerous high-loss samples are left unused.

In this paper, we propose the ScarceNet, a PL-based approach to learn an animal pose network from scarce annotations. We first train an animal pose estimation network on the small set of labeled data, and then generate pseudo labels for the unlabeled images with the pretrained model. The pseudo labels are often corrupted with noise due to the small size of the training data. We thus also select reliable pseudo labels using the small-loss trick. However, instead of ignoring or adding consistency regularization to the high-loss samples as done in previous works [16, 20], we propose to identify reusable samples from those high-loss samples based on an agreement check. The agreement check is to guarantee that the model predictions for the reusable samples are smooth and concurrently far from the decision boundary, and hence allow us to directly re-label reusable samples with the model predictions. Finally, we introduce a student-teacher network [28] and enforce a consistency constraint since there are still samples that are neither reliable nor reusable. By combining the reliable pseudo label selection with the reusable sample re-labeling and the consistency constraint, we can fully exploit the unlabeled data.

In addition to the training scheme, we also modify the state-of-the-art human pose estimation network HRNet [27] and adapt it to the pseudo label-based learning framework. The original HRNet regresses the heatmaps of different joints from a shared feature representation. However, the negative effect of noisy pseudo label for any joint is propagated to all other joints through this shared feature representation. To mitigate this problem, we introduce a multi-branch HRNet (MBHRNet) where a joint-specific feature representation is learned for regressing the corresponding joint location in each branch.

We validate our approach on the challenging AP-10K dataset [33], which includes animal images from more than 50 species. Our approach outperforms existing SSL methods by a large margin, especially when the labeled data is very limited (for example, 5 labeled images per species). We also test our approach under the inter-family transfer

setting, where animal species from only one family are labeled and the rest are unlabeled. Our approach improves the HRNet baseline significantly, and even outperform the supervised counterpart for some species. We further test on the TigDog dataset [20], where we can outperform existing synthetic data based approaches when only 0.5% of the dataset is labeled. Our contribution can be summarized as follows:

- We propose a pseudo label based training scheme for the task of animal pose estimation from scarce labels.
- We introduce the MBHRNet to mitigate negative effect of noisy joints to other joints in the same image.
- We achieve state-of-the-art performance on several animal pose benchmarks under different settings.

2. Related Work

Animal Pose Estimation. Animal pose estimation, including both 2D [5, 16, 20] and 3D [4, 15, 37], has become an active research field in the last few years. The development of animal pose estimation techniques is relatively slower than the human counterpart due to the lack of large scale animal pose dataset like COCO [18] or MPII [1]. Existing works circumvent the requirement for animal pose data by transferring knowledge from other domains, such as human pose data [5] and synthetic animal data [16, 20, 23]. Although impressive results have been achieved in [5], substantial amounts of labeled animal images are still needed to facilitate the knowledge transfer. This is due to the large domain gap between human and animal images, including pose, appearance and body structure differences. To reduce the domain gap between the source and target domains, synthetic animal images are generated from CAD models and used as the source images in [20]. Three consistency criteria are proposed to generate pseudo labels for real images. Improved upon [20], Li *et al.* [16] propose a pseudo label updating strategy to prevent the network from overfitting to the noisy pseudo labels. Despite the domain gap is reduced by using synthetic animal images, the limited pose and appearance variations of the CAD models prohibit the synthetic data based approaches from adapting to real scenarios with crowded scene or self-occlusion. Moreover, synthesize realistic images is a tedious process, which also requires expert knowledge. In view of this, we focus on learning from scarce annotations to minimize human labor and achieve accurate animal pose estimation at the same time.

Semi-supervised Learning. Semi-supervised learning is powerful in leveraging unlabeled data to improve a model's performance when labeled data is limited. One of the most widely used SSL techniques is pseudo labeling [14, 31], which generates artificial labels for unlabeled images from

model predictions. Another technique is consistency regularization [13, 22, 28], which enforces that the model output should be consistent when the input is randomly perturbed. MixMatch [3] combines the consistency regularization with the entropy minimization to encourage the network to output confident predictions for unlabeled data. More recent Unsupervised Data Augmentation (UDA) [30] and FixMatch [25] combine pseudo labeling, consistency regularization with strong augmentations [6] and achieve superior performance. FlexMatch [34] further improves over UDA and Fixmatch by introducing a curriculum learning scheme. We also adopt the PL-based techniques in this work given its effectiveness. To mitigate the negative effect of noisy pseudo labels, we combine PL with reliable pseudo label selection and reusable sample re-labeling.

Learning from Noisy Labels. Learning from noisy labels has become an important task since deep neural networks are known to be susceptible to noisy labels. To reduce the negative effect of noisy labels, different approaches which can be divided into three categories have been proposed. The first category is to treat the true labels as hidden variables and use an extra network to estimate a noise transition matrix [8, 21]. These methods assume there is correlation between the corrupted and clean labels, hence cannot handle complicate cases with random noise. The second category is to design noise-robust loss functions [7], such as the generalized cross entropy [35] and the symmetric cross entropy [29]. The robust-loss based approaches are theoretically sound, but may only be capable of handling a certain noise rate. The last category [9, 11, 19] focuses on selecting potential clean samples to update the network. One promising selection criteria is the small-loss trick [9], which is designed based on the memorization effect of deep networks and has been widely applied to different tasks [16, 26, 32]. We also apply the small-loss trick to select reliable pseudo labels. However, instead of ignoring the high-loss samples as done in previous works, we combine the sample selection with the reusable sample re-labeling to make full use of the unlabeled data.

3. Our Method: ScarceNet

We propose a PL-based approach to learn animal pose estimation from scarce annotations. The overall framework of our ScarceNet is illustrated in the left part of Fig. 2, where we propose a multi-branch HRNet (MBHRNet) as our animal pose estimation network. Let us denote an input image as \mathbf{x} and its ground truth animal pose (if available) as \mathbf{y} . Given a small set of labeled images $\mathcal{D}_l = \{\mathbf{x}_i^l, \mathbf{y}_i^l\}_{i=1}^{N_l}$ and unlabeled images $\mathcal{D}_u = \{\mathbf{x}_i^u\}_{i=1}^{N_u}$, the objective is to train the animal pose estimation network by leveraging both the labeled and unlabeled data. N_l and N_u represent the

number of labeled and unlabeled images, respectively, and $N_l \ll N_u$. The labeled images are supervised with the ground truth pose data. We then generate pseudo labels with the model trained on the labeled data for the unlabeled images. However, these generated pseudo labels are usually heavily corrupted with noise, especially when the labeled data is very limited. We thus apply the small-loss trick [9] to select a set of reliable pseudo labels. Despite its effectiveness, pseudo label selection by the small-loss trick tends to discard numerous high-loss samples. This results in high wastage since those discarded samples can still provide extra information for better discrimination. In view of this, we propose a reusable sample re-labeling step to further identify reusable samples from the high-loss samples via an agreement check and re-generate the corresponding pseudo labels for supervision. Lastly, we design a student-teacher framework to enforce consistency between the outputs of the student and teacher network.

3.1. Reliable Pseudo Label Selection

We first train the animal pose estimation network on labeled images \mathcal{D}_l with a supervised loss:

$$\mathcal{L}_s = \frac{1}{J} \sum_{i=1}^J \|H_i - \hat{H}_i\|^2, \quad (1)$$

where H_i is the ground truth heatmap, J denotes the total number of joints, and \hat{H}_i represents the output heatmap of the network. The pretrained network is then used to generate an initial set of pseudo labels for the unlabeled images \mathcal{D}_u following the procedure in [20]. The generated pseudo labels are generally noisy and can hurt the performance if directly used for training. To prevent this, we select reliable pseudo labels, denoted as \mathbb{S} , from the initial pseudo labels using the small-loss trick. Specifically, we only make use of samples with loss smaller than a threshold to update the network. Formally, this can be expressed as:

$$\mathcal{L}_p = \frac{\sum_{i=1}^J \{\mathcal{L}_p^i < l_r\}_1 \mathcal{L}_p^i}{\sum_{i=1}^J \{\mathcal{L}_p^i < l_r\}_1}, \text{ where } \mathcal{L}_p^i = \|\tilde{H}_i - \hat{H}_i\|^2. \quad (2)$$

\tilde{H}_i represents the pseudo label for the i^{th} joint and \mathcal{L}_p^i is the corresponding pseudo label based loss. $\{\text{condition}\}_1$ is a condition function, which outputs 1 when the condition is true and 0 if false. l_r denotes the threshold loss value, and more details on how to decide l_r are provided in the supplementary material. This small-loss trick is designed based on the memorization effect of deep network, and has been shown effective in selecting reliable labels in previous works [16, 26].

3.2. Reusable Sample Re-labeling

Only small-loss samples are selected in the reliable pseudo label selection process, and the remaining pseudo

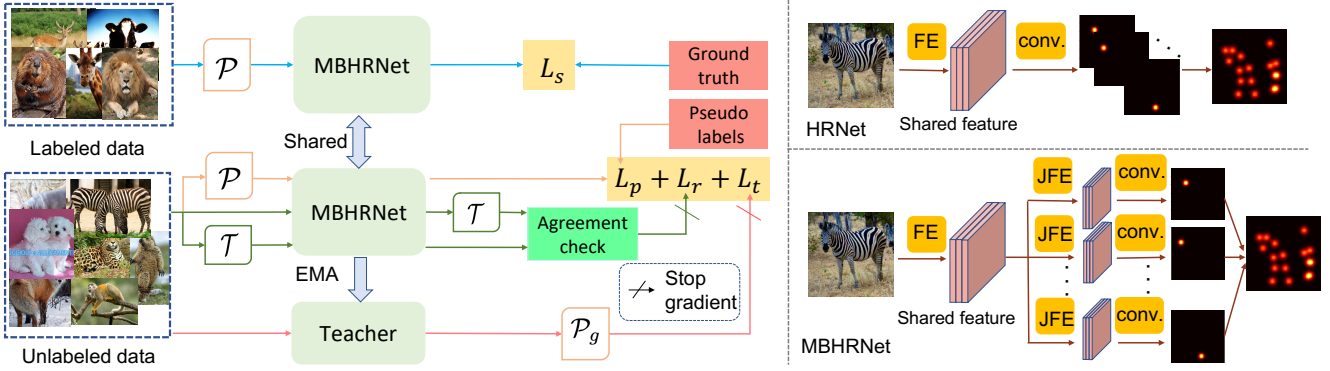


Figure 2. The left part shows the overall framework of our ScarceNet, where \mathcal{P} and \mathcal{T} represent strong and weak augmentation, respectively. The right part shows a comparison between the HRNet (top) and the MBHRNet (bottom), where ‘FE’ represents feature extractor, ‘conv.’ represents convolutional layer and ‘JFE’ represents joint specific feature extractor.

labels with large loss are discarded. This is improvident since the high-loss samples can also contribute to improving the network if used properly. To this end, we propose to further identify reusable samples, denoted as \mathbb{U} , from the unused samples to fully exploit the unlabeled data. We define reusable samples as the ones that may also provide reliable knowledge to the network despite their large losses. To identify the reusable samples, we introduce an agreement check. Specifically, given an input \mathbf{x}^u , we generate another view of this sample by adding a random transformation \mathcal{T} , which consists of rotation and scaling. Both samples are fed into the pose estimation network to produce the corresponding heatmaps, denoted as:

$$H^{v_1} = f(\mathbf{x}^u), H^{v_2} = f(\mathcal{T}\mathbf{x}^u). \quad (3)$$

The location for the i^{th} joint is obtained by applying argmax operation to the corresponding output heatmap:

$$\mathbf{h}_i^{v_1} = \operatorname{argmax}(H_i^{v_1}), \mathbf{h}_i^{v_2} = \operatorname{argmax}(H_i^{v_2}). \quad (4)$$

Unlike the classification task where the output is transformation invariant [12, 32], the output joint location is transformation equivariant. Consequently, we have to transform $\mathbf{h}_i^{v_1}$ before the agreement check, *i.e.*:

$$d_i = \|\mathcal{T}\mathbf{h}_i^{v_1} - \mathbf{h}_i^{v_2}\|^2. \quad (5)$$

Subsequently, we can determine whether a sample is reusable based on the distance between the outputs of the two views. Specifically, a sample \mathbf{x}_i^u is reusable when the distance fulfills $d_i < d_r$, where d_r represents the distance threshold. This agreement check is based on the agreement maximization principle [24], which has also been used for out-of-distribution detection [32] and stable sample selection [12]. Intuitively, the reusable samples must satisfy smoothness in the neighborhood of the model prediction

and must also be far from the decision boundary. Otherwise, the output becomes inconsistent when the input is randomly transformed.

The next step is to re-label the identified reusable samples since they exhibit large loss based on the initial pseudo labels. We directly use the prediction from the current model as the new pseudo labels to train the network. The motivation is twofold: 1) The model has been trained with unlabeled images, hence it is able to output more accurate prediction comparing to the model trained only on the labeled images. 2) The selection criteria for reusable samples indicates that the current model prediction is already reliable. Finally, the loss function for the reusable samples can be defined as:

$$\mathcal{L}_r = \frac{\sum_{i=1}^J \{d_i < d_r\}_1 \mathcal{L}_r^i}{\sum_{i=1}^J \{d_i < d_r\}_1}, \text{ where } \mathcal{L}_r^i = \|\bar{H}_i - \hat{H}_i\|^2. \quad (6)$$

\bar{H}_i represents the heatmap generated from the joint location of current model prediction. By selecting the reusable samples, we are able to make use of more unlabeled images and prevent the network from overfitting to the initial noisy pseudo labels at the same time.

3.3. Student-teacher Consistency

Both the reliable pseudo label selection and reusable sample selection are based on specific constraints. Samples that do not fulfill either constraint are still left unused. In view of this, we further introduce a teacher network and enforce the consistency between the student and teacher network. This is inspired by the mean teacher framework [28], which enforces the consistency constraint for unlabeled data for semi-supervised classification. The teacher network shares the same structure with the student network, and the weights are updated with an exponential moving average

(EMA) of the student model:

$$\theta_{\text{st}}^t = \alpha \theta_{\text{st}}^{t-1} + (1 - \alpha) \theta^t. \quad (7)$$

θ_{st} and θ denote the weights of the teacher and student network, respectively. The superscript t represents the training step, and α is a smoothing coefficient of the EMA update. We can see that the teacher network is a temporal ensemble of networks, and thus it can provide more accurate pseudo labels. As shown in Fig. 2, we add random perturbations \mathcal{P} to the input of the student network, which consists of both geometry-based perturbations \mathcal{P}_g such as rotation, scaling and flipping, and image-based perturbation \mathcal{P}_m such as noise, occlusion and blurring. Consistency is then enforced between the teacher and student network as:

$$\mathcal{L}_t = \frac{1}{J} \sum_{i=1}^J \|\mathcal{P}_g H_i^t - \hat{H}_i\|^2, \quad (8)$$

where H_i^t denotes the output of the teacher network. Note that only the geometry-based perturbation \mathcal{P}_g is added to the output of the teacher network since the output joint location is invariant to the image-based perturbation.

3.4. Multi-head Animal Pose Estimation Network

We adopt the human pose estimation network HRNet as our backbone. The HRNet has achieved state-of-the-art performance for human pose estimation by combining high-resolution feature representations with multi-resolution feature fusion. However, as shown in the right top part of Fig. 2, the final heatmaps of HRNet are regressed from a shared feature representation with one linear layer. This shared mechanism can have negative effect in the presence of noisy labels. For example, the feature representation is affected by backpropagation when the pseudo label for any joint is noisy and thus further affect the prediction of other joints in the same image. To mitigate this problem, we introduce a multi-branch head to the HRNet to learn a specific feature representation for each joint. As shown in the right bottom part of Fig. 2, the network has J output branches, where each branch learns a specific feature representation via a joint specific feature extractor (JFE) to regress the corresponding joint location. In this way, the adverse effect from the noisy pseudo label of a joint gets absorbed in its respective branch.

3.5. Strong-weak Augmentation

We also introduce a strong-weak augmentation training scheme, where two sets of augmentation are applied to the inputs, denoted as \mathcal{T} and \mathcal{P} in previous sections. Particularly, we apply strong augmentations \mathcal{P} for the learning task with backpropagation and weak augmentation \mathcal{T} for the re-labeling task. As shown in Fig. 2, strong augmentation \mathcal{P} is applied to the labeled images in the blue path

and unlabeled images in the orange path. These two paths are responsible for updating the network, and hence strong augmentation can improve the generalization capacity. In comparison, the weak augmentation \mathcal{T} is applied to the green path because this path is responsible for identifying and re-labeling reusable samples. Overly-strong augmentation can have adversarial affects on the agreement check process since strong augmentation tends to result in inconsistent predictions. More details on the augmentation are provided in the supplementary material.

3.6. Overall Framework

We train our network in two stages. In the first stage, we train the MBHRNet on \mathcal{D}_l with the supervised loss in Eqn. (1). The trained network is applied to generate pseudo labels for the unlabeled data. In the second stage, the student network is initialized with the pretrained model and trained with both \mathcal{D}_l and \mathcal{D}_u . The objective function for the second stage consists of both labeled and unlabeled losses as follows:

$$\mathcal{L} = \lambda_1 \mathcal{L}_s + \lambda_2 \mathcal{L}_p + \lambda_3 \mathcal{L}_r + \lambda_4 \mathcal{L}_t, \quad (9)$$

where $\lambda_1, \lambda_2, \lambda_3$, and λ_4 are the weights for different losses. The overall training procedure is presented in Algorithm 1.

Algorithm 1 Overall training procedure of the second stage

Input: labeled data \mathcal{D}_l and unlabeled data \mathcal{D}_u , threshold loss value l_r , distance threshold d_r , batch size B

- 1: **for** $\{x_1, x_2, \dots, x_B\} \in \mathcal{D}_l \cup \mathcal{D}_u$ **do**
 - 2: **if** $x_i \in \mathcal{D}_l$ **then**
 - 3: Compute supervised loss \mathcal{L}_s according to Eqn. (1)
 - 4: **else**
 - 5: **if** $x_i \in \mathcal{S}$ **then**
 - 6: Compute the pseudo labeled based loss \mathcal{L}_p according to Eqn. (2)
 - 7: **else if** $x_i \in \mathcal{U}$ **then**
 - 8: Compute loss \mathcal{L}_r for reusable samples by Eqn. (6)
 - 9: **end if**
 - 10: Compute the consistency loss \mathcal{L}_t according to Eqn. (8)
 - 11: **end if**
 - 12: Compute the total loss \mathcal{L} according to Eqn. (9)
 - 13: Update student network $\theta \leftarrow \theta - \eta \Delta \mathcal{L}$
 - 14: Update teacher network by Eqn. (7)
 - 15: **end for**
-

4. Experimental Results

Training details. We adopt the HRNet-w32 as the backbone and train our network in two stages. We first train the MBHRNet on the labeled data with the supervised loss. Following [27], we train the network with the Adam optimizer for 210 epochs. The initial learning rate is set to 10^{-3} , and is dropped to 10^{-4} and 10^{-5} at the 170th and 200th epochs, respectively. The trained model is then used to generate pseudo labels following [20]. In the second stage, we initialized the student network from the first stage and train the whole framework with both labeled and unlabeled data for another 150 epochs. More training details are provided in the supplementary material.

Datasets. We evaluate our method on the AP-10K dataset [33] and the TigDog dataset [20]. The AP-10K dataset is a recently proposed animal pose dataset, which contains more diverse animal species than previous datasets [5, 20]. The dataset contains 10,015 labeled images in total, which are split into train, validation, and test sets with a ratio of 7:1:2 per animal species. The TigDog dataset includes 8380 images for the horse category and 6523 images for the tiger category. This dataset is used in several domain adaptation works [16, 20] to show that the knowledge learned from synthetic animals can be transferred to real animal images. We also evaluate on this dataset to show that better performance can be achieved by labeling just a few images, hence can reduce human labor compared to generating synthetic data.

Evaluation metrics. We adopt the Mean Average Precision (mAP) as the evaluation metric for the AP-10K dataset. The Percentage of Correct Keypoints (PCK), which computes the percentage of joints that are within a normalized distance to the ground truth locations, is used for the TigDog dataset following [20].

	5	10	15	20	25
HRNet [27]	0.360	0.463	0.511	0.547	0.588
UDA [30]	0.429	0.519	0.566	0.580	0.628
FixMatch [25]	0.478	0.544	0.589	0.601	0.631
FlexMatch [34]	0.466	0.555	0.596	0.618	0.646
Ours	0.533	0.597	0.632	0.654	0.681

Table 1. Results (mAP) on the AP-10K when 5, 10, 15, 20 and 25 images per species are labeled. Best results in bold.

4.1. Results on the AP-10K Dataset.

We evaluate our approach under two scenarios on the AP-10K dataset. In the first scenario, a set of images are labeled for each animal species. The second scenario is that only the images in one family are labeled and we test on

other animal families. This setting is more challenge since the appearance, size and background environment are quite different for different animal species.

We show results for the first scenario in Tab. 1, where we evaluate when 5, 10, 15, 20 and 25 images per animal species are labeled. There are 50 species in total, which results in 250, 500, 720, 1000 and 1250 labeled images out of 7023 images. We compare with state-of-the-art semi-supervised approaches, including UDA [30], FixMatch [25] and FlexMatch [34]. ‘HRNet’ represents HRNet trained only with labeled images. We reimplement those semi-supervised approaches according to the open repository since they only show results for classification task. We can see that our approach outperform existing approaches by a large margin, especially when very few labeled images are available, such as the case of 5 labels per species. The reason is that UDA, FixMatch and FlexMatch only retain the samples with reliable pseudo labels based on the confidence scores. In comparison, we combine the reliable pseudo label selection with the reusable sample re-labeling and teacher-student consistency, which helps to fully utilize the unlabeled data. Moreover, our results are the closest to the fully supervised setting [33], *i.e.* 0.738 mAP, when only 25 images per animal species are labeled. We also show qualitative results for the 25 labels per animal species setting in Fig. 3.

We also show results for the case when images from only one animal family are labeled in Tab. 2. Following the inter-family transfer learning setting in [33], we train our model with labeled images of the Bovidae family and the remaining unlabeled images, and test on families including Cervidae, Equidae and Hominidae. There are seven animal species in the Bovidae family that includes a total of 755 images. The results under the ‘Generalization’, ‘Few-Shot’ and ‘Transfer’ settings are taken from [33]. ‘Generalization’ means that the model is trained only with images in the Bovidae, ‘Few-Shot’ means that 20 labels from the testing species are used, and ‘Transfer’ means that all labels of the testing species are used. We can see that we achieve better performance than the generalization setting in all cases except for ‘Moose’. Moreover, we even outperform the ‘Transfer’ setting for ‘Deer’ and ‘Horse’ although we do not use any labels from the target species. Our results for the ‘Chimpanzee’ and ‘Gorilla’ are not satisfactory although still comparable to the ‘Few-Shot’ setting. We show the failure cases of ‘Chimpanzee’ and ‘Gorilla’ in the last two examples of Fig. 3. We can see that ‘Chimpanzee’ and ‘Gorilla’ are very different from animal species in the Bovidae family (the first four examples of the last row) in terms of appearance and pose, which makes the knowledge transfer hard without using any labels from the target species.

Species	Setting	Performance	Species	Setting	Performance
Deer	Generalization	0.723	Moose	Generalization	0.587
	Few-Shot	0.742		Few-Shot	0.648
	Transfer	0.751		Transfer	0.726
	Ours	0.767		Ours	0.547
Horse	Generalization	0.592	Zebra	Generalization	0.324
	Few-Shot	0.635		Few-Shot	0.480
	Transfer	0.718		Transfer	0.708
	Ours	0.749		Ours	0.562
Chimpanzee	Generalization	0.009	Gorilla	Generalization	0.017
	Few-Shot	0.022		Few-Shot	0.144
	Transfer	0.550		Transfer	0.662
	Ours	0.057		Ours	0.033

Table 2. Results (mAP) for the Cervidae, Equidae and Hominidae when only images from the Bovidae are labeled.

	Real	Cycgan [36]	BDL [17]	Cycada [10]	CC-SSL [20]	MDAM-MT [16]	Ours
Horse	78.98	51.86	62.33	55.57	70.77	79.50	73.05
Tiger	81.99	46.47	52.26	51.48	64.14	67.76	74.88
Average	80.48	49.17	57.30	53.53	67.52	73.66	73.83

Table 3. PCK@0.05 accuracy of the TigDog dataset. ‘Real’ means full supervision. Cycgan, BDL, Cycada, CC-SSL, MDAM-MT are domain adaptation approaches using synthetic data. Our approach is trained with only 36 labeled images per category.

	Horse Accuracy								Tiger Accuracy								Average
	Eye	Chin	Shoulder	Hip	Elbow	Knee	Hooves	Mean	Eye	Chin	Shoulder	Hip	Elbow	Knee	Hooves	Mean	
CC-SSL	84.60	90.26	69.69	85.89	68.58	68.73	61.33	70.77	96.75	90.46	44.84	77.61	55.82	42.85	64.55	64.14	67.52
MDAM-MT	91.05	93.37	77.35	80.67	73.63	81.83	73.67	79.50	97.01	91.18	46.63	78.08	50.86	61.54	70.84	67.76	73.66
Ours	89.42	93.60	74.63	83.61	74.31	79.91	73.99	79.18	99.13	92.82	46.96	75.36	50.46	63.51	74.82	68.98	74.13

Table 4. PCK@0.05 accuracy of the TigDog dataset for domain adaptation setting. ‘Mean’ represents the accuracy for each animal category, and ‘Average’ represents the average accuracy for all animal categories. Best results in bold.

	5	10	15	20	25
Full	0.533	0.597	0.632	0.655	0.681
- RSR	0.521	0.586	0.621	0.633	0.668
- MT	0.487	0.568	0.614	0.628	0.659
- MB	0.520	0.585	0.624	0.645	0.670
- AUG	0.487	0.565	0.608	0.625	0.654

Table 5. Ablation study on the AP-10K dataset when 5, 10, 15, 20 and 25 images per animal species are labeled.

4.2. Results on the TigDog Dataset

The TigDog Dataset is used in previous domain adaptation works [16, 20] to show that models trained from synthetic images can be used for real ones. We also evaluate on this dataset to show that our approach can achieve comparable or even better performance compared with state-of-the-art domain adaptation works with very few labeled images. Note that we use the same ResNet backbone as [20] for all experiments on the TigDog dataset for fair comparison. The results are shown in Tab. 3, where the numbers are

directly taken from [16]. ‘Real’ denotes fully-supervised setting, Cycgan [36], BDL [17], Cycada [10], CC-SSL [20], MDAM-MT [16] are domain adaptation approaches trained with synthetic animal images and ‘Ours’ means our approach trained with 36 images per animal category. Note that 36 is the minimum labeled data needed for each category to achieve better performance than existing unsupervised domain adaptation approaches, *i.e.* we only need to label 72 out of 14,903 images for the Tigdog dataset. The results demonstrate the advantage of our setting since annotating such few images does not take much human labor.

We also test our approach on the domain adaptation setting to further verify the effectiveness. The results are shown in Tab. 4, which includes accuracy for different joints, the mean accuracy for each animal category and the average accuracy over all categories. We can see that our approach achieves the best performance in most of the cases. Note that our approach and CC-SSL have simpler networks compared to MDAM-MT, which introduce an extra refinement module for self-knowledge distillation. More training details are provided in the supplementary material.

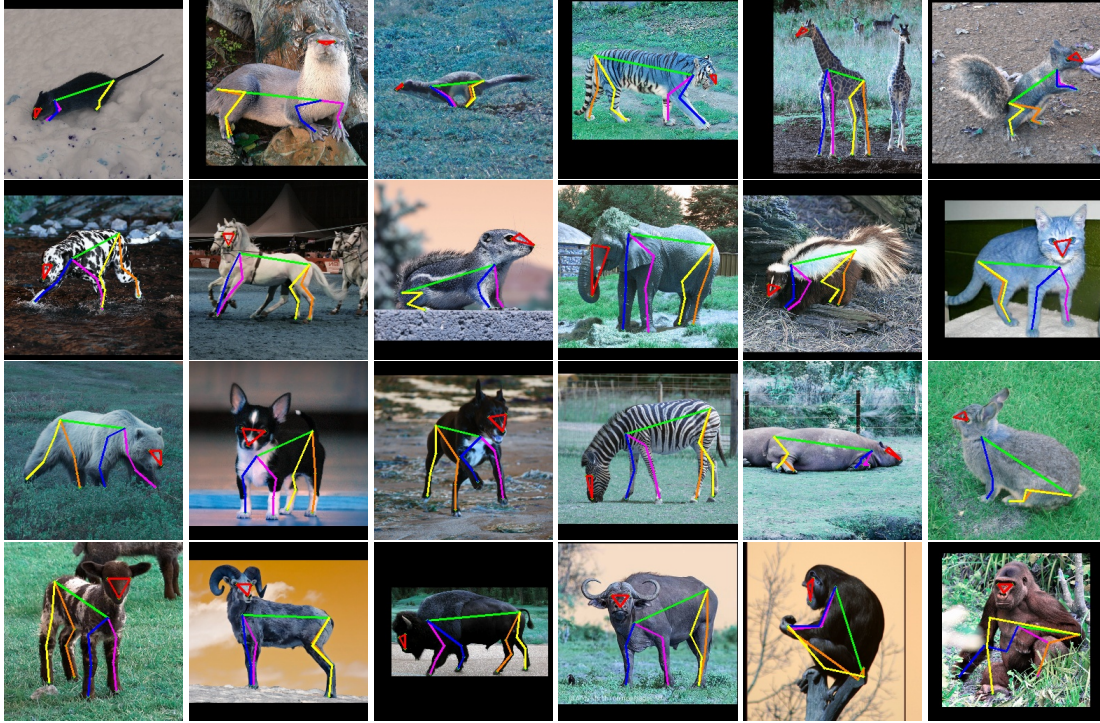


Figure 3. Qualitative results for the AP-10K dataset when 25 images per species is labeled (except for the last two examples). The last two examples show the **failure cases** when only images from the Bovidae family are labeled.

4.3. Ablation Study

We conduct ablation study on the AP-10K dataset. Each component is removed from the full model to verify the corresponding contribution. The results are shown in Tab. 5, where ‘RSR’ denotes reusable sample re-labeling, ‘MT’ denotes student teacher consistency, ‘MB’ denotes multi-branch architecture and ‘AUG’ denotes strong-weak augmentation. Note that we remove the weak augmentation \mathcal{T} and replace the strong augmentation \mathcal{P} with commonly used rotation and scale augmentation for the ablation of ‘AUG’. We can see that the performance drops when each component is removed. Specially, the accuracy drops significantly for the 5 labels per animal species setting when the student teacher consistency or strong-weak augmentation is removed. This can be attributed to two reasons. Firstly, the strong-weak augmentation enforces the output of the strong augmented images to be consistent with the pseudo labels that are generated using weak augmented images. In this way, the network is enforced to learn a smooth representation around each image, which helps to improve the performance especially in the low-data regime. Secondly, The initial set of pseudo labels are extremely noisy when the model is trained with very few labeled data. The teacher network, which is a temporal ensemble of network, is able to provide more stable supervision for the student network.

5. Limitations

Our approach achieves impressive performance even when very few labeled data is available. However, there are still many remaining challenges for the task of animal pose estimation for limited labeled data. The large variations of appearance and pose across different animal species makes it hard for the network to generalize well to different animal species, *e.g.* train with labels from Bovidae and test on Hominidae. We leave this for future research.

6. Conclusion

We propose a pseudo label based training scheme for animal pose estimation from scarce annotations. We combine the reliable pseudo label selection with the reusable sample re-labeling and the student-teacher consistency to make full use of the unlabeled data. We also introduce the MBHRNet to mitigate the negative effect of noisy joints. Extensive experiments have been conducted to verify the effectiveness of our proposed approach.

Acknowledgement. This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2021-024), and the Tier 2 grant MOE-T2EP20120-0011 from the Singapore Ministry of Education.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt 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*, pages 3686–3693, 2014. 2
- [2] Devansh Arpit, Stanisław Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International conference on machine learning*, pages 233–242. PMLR, 2017. 2
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32, 2019. 1, 3
- [4] Benjamin Biggs, Oliver Boyne, James Charles, Andrew Fitzgibbon, and Roberto Cipolla. Who left the dogs out? 3d animal reconstruction with expectation maximization in the loop. In *European Conference on Computer Vision*, pages 195–211. Springer, 2020. 2
- [5] Jinkun Cao, Hongyang Tang, Hao-Shu Fang, Xiaoyong Shen, Cewu Lu, and Yu-Wing Tai. Cross-domain adaptation for animal pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9498–9507, 2019. 1, 2, 6
- [6] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 3
- [7] Aritra Ghosh, Himanshu Kumar, and P Shanti Sastry. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017. 3
- [8] Jacob Goldberger and Ehud Ben-Reuven. Training deep neural-networks using a noise adaptation layer. In *ICLR*, 2017. 3
- [9] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in neural information processing systems*, pages 8527–8537, 2018. 2, 3
- [10] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998, 2018. 7
- [11] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pages 2304–2313, 2018. 3
- [12] Zhanghan Ke, Daoye Wang, Qiong Yan, Jimmy Ren, and Rynson WH Lau. Dual student: Breaking the limits of the teacher in semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6728–6736, 2019. 4
- [13] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. 3
- [14] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013. 2
- [15] Chen Li and Gim Hee Lee. Coarse-to-fine animal pose and shape estimation. *Advances in Neural Information Processing Systems*, 34:11757–11768, 2021. 2
- [16] Chen Li and Gim Hee Lee. From synthetic to real: Unsupervised domain adaptation for animal pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1482–1491, 2021. 1, 2, 3, 6, 7
- [17] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6936–6945, 2019. 7
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755, 2014. 2
- [19] Eran Malach and Shai Shalev-Shwartz. Decoupling” when to update” from” how to update”. In *Advances in Neural Information Processing Systems*, pages 960–970, 2017. 3
- [20] Jiteng Mu, Weichao Qiu, Gregory D Hager, and Alan L Yuille. Learning from synthetic animals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12386–12395, 2020. 1, 2, 3, 6, 7
- [21] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1944–1952, 2017. 3
- [22] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. *Advances in neural information processing systems*, 28, 2015. 3
- [23] Moira Shooter, Charles Malleon, and Adrian Hilton. Sydog: A synthetic dog dataset for improved 2d pose estimation. *arXiv preprint arXiv:2108.00249*, 2021. 1, 2
- [24] Vikas Sindhwani, Partha Niyogi, and Mikhail Belkin. A co-regularization approach to semi-supervised learning with multiple views. In *Proceedings of ICML workshop on learning with multiple views*, volume 2005, pages 74–79. Citeseer, 2005. 4
- [25] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 1, 3, 6

- [26] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *International Conference on Machine Learning*, pages 5907–5915, 2019. [3](#)
- [27] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. [2](#), [6](#)
- [28] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#), [3](#), [4](#)
- [29] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 322–330, 2019. [3](#)
- [30] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33:6256–6268, 2020. [3](#), [6](#)
- [31] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10687–10698, 2020. [2](#)
- [32] Yazhou Yao, Zeren Sun, Chuanyi Zhang, Fumin Shen, Qi Wu, Jian Zhang, and Zhenmin Tang. Jo-src: A contrastive approach for combating noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5192–5201, 2021. [3](#), [4](#)
- [33] Hang Yu, Yufei Xu, Jing Zhang, Wei Zhao, Ziyu Guan, and Dacheng Tao. Ap-10k: A benchmark for animal pose estimation in the wild. *arXiv preprint arXiv:2108.12617*, 2021. [1](#), [2](#), [6](#)
- [34] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34:18408–18419, 2021. [1](#), [3](#), [6](#)
- [35] Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31, 2018. [3](#)
- [36] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [7](#)
- [37] Silvia Zuffi, Angjoo Kanazawa, Tanya Berger-Wolf, and Michael J Black. Three-d safari: Learning to estimate zebra pose, shape, and texture from images” in the wild”. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5359–5368, 2019. [2](#)