

This ACCV 2024 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2024 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

Match me if you can: Semi-Supervised Semantic Correspondence Learning with Unpaired Images

Jiwon ${\rm Kim}^{1,2*}$ Byeongho ${\rm Heo}^1$ Sangdoo Yun^1 Seungryong ${\rm Kim}^3$ Dongyoon ${\rm Han}^1$

 1 NAVER AI Lab 2 LG AI Research $^{-3}$ KAIST

Abstract. Semantic correspondence methods have advanced to obtaining high-quality correspondences employing complicated networks, aiming to maximize the model capacity. However, despite the performance improvements, they may remain constrained by the scarcity of training keypoint pairs, a consequence of the limited training images and the sparsity of keypoints. This paper builds on the hypothesis that there is an inherent data-hungry matter in learning semantic correspondences and uncovers the models can be more trained by employing densified training pairs. We demonstrate a simple machine annotator reliably enriches paired key points via machine supervision, requiring neither extra labeled key points nor trainable modules from unlabeled images. Consequently, our models surpass current state-of-the-art models on semantic correspondence learning benchmarks like SPair-71k, PF-PASCAL, and PF-WILLOW and enjoy further robustness on corruption benchmarks. Our code is available at https://github.com/naver-ai/matchme.

1 Introduction

Learning dense correspondence between image pairs is a fundamental problem that facilitates many computer vision tasks [15, 18, 23, 27, 33, 51, 55]. In contrast to classical tasks, where images are captured in geometrically constrained settings such as different views of the same scene [8, 31] or neighboring frames in a video [11, 13, 45], the semantic correspondence task [1, 4, 12, 28] finds pixel-wise visual correspondences between images containing the same object or semantic meaning. Due to these unconstrained settings, it should handle the additional challenges of large intra-class variations in appearance and background clutter. Recent methods [2, 3, 16, 17, 25, 26, 29, 32, 34, 36, 59] were generally trained to fit on full-labeled datasets [4, 35] providing limited training pairs with manually annotated keypoint pairs. The rigorous requirements of pixel-level semantic correspondences lead to considerable time and expense in manual annotation by experts. This results in a limited quantity of available training data; we call it a data-hungry problem in semantic correspondence learning.

Various methods have focused on unsupervised strategies [10, 14, 19, 40, 50] to increase the amount of correspondence supervision on unlabeled data in a

^{*} Work done while at NAVER AI Lab, currently at LG AI Research.

Correspondence to Dongyoon Han.

self-supervised or weakly-supervised way. In particular, the weakly-supervised methods attempted to solve the problem by using a cycle consistency [19, 50] or pseudo-labels [10, 14] on real image pairs for unsupervised loss signal, but they still only rely on image pairs in the training set. The capability of heavy matching networks hinges on data quantity at first, but the training data remains significantly smaller than other computer vision tasks (*e.g.*, 1.2M images in ImageNet-1K [42]). Therefore, we argue that previous approaches [10, 14, 19, 50], attempting to densify points for training, may not be an underlying solution for the data-hungry problem.

In this paper, we present a fundamental approach dubbed MatchMe focusing on overcoming the insufficiency of both image and point pairs. We utilize unlabeled image pairs, having potentially rich semantic information that has remained unannotated. The unlabeled images newly supplied can be utilized to generate a bunch of novel pairs with originally labeled (Fig. 1(a)) or other unlabeled images (Fig. 1(b)); the newly created keypoint pairs densify labels for training (Fig. 1(c)). We adopt machine annotators [24,37,47,53] to acquire densified labels for simplicity as well. We conjecture that a machine annotator could offer reliable labels based on the findings from [53,56]. Additionally, our framework allows for improved label quality by iteratively updating the annotator with the current trained model in successive training cycles.

Our proposed method is demonstrated by applying it to recent matching architectures [2, 3] to show applicability. Experimental results prove that our method is effective and achieves state-of-the-art performance on every benchmark, including PF-PASCAL [4], PF-WILLOW [4], and SPair-71k [35]. MatchMe achieves state-of-the-art performance on all semantic correspondence benchmarks, showing accuracy gain of 2.0% and 2.4% on PF-WILLOW and SPair-71k (PCK@ $\alpha = 0.1$).

2 Background

2.1 Task Definition

The semantic correspondence task aims to predict the matching probability P between a semantically similar image pair. Given a training image pair S with source image $I_s \in \mathbb{R}^{H_s \times W_s}$ and target image $I_t \in \mathbb{R}^{H_t \times W_t}$, a matching function f with the network parameters θ predicts $P_{s,t} = f(I_s, I_t; \theta) \in \mathbb{R}^{H_s W_s \times H_t W_t}$ by considering the feature similarities across all the points in I_s and I_t . It minimizes the following problem with image pairs S and supervision $\hat{P}_{s,t} \in \mathbb{R}^{H_s W_s \times H_t W_t}$ between two images (I_s, I_t) :

$$\mathcal{L}_{\mathcal{S}} = \frac{1}{|\mathcal{S}|} \sum_{(I_s, I_t) \in \mathcal{S}} \sum_{i=1}^{H_t W_t} \hat{M}_{s,t}(i) \mathcal{D}(P_{s,t}(\cdot, i), \hat{P}_{s,t}(\cdot, i)),$$
(1)

where (\cdot, i) indicates the *i*-th column of a matrix and $\mathcal{D}(\cdot, \cdot)$ is a distance function. $\hat{M}_{s,t} \in \mathbb{B}^{H_t W_t}$ denotes a binary mask vector, in which $\hat{M}_{s,t}(i)$ corresponds to the

MatchMe 3



(b) New images with no keypoints

(c) Mined abundant keypoint pairs from previously untapped images

Fig. 1: Untapped annotation gems. Semantic correspondence learning usually suffers from data hunger, so few sparsely paired keypoints drawn by yellow lines in labeled data inherently limit the performance. (a) Labeled images in the SPair-71k benchmark [35] contain sparse manually annotated keypoint pairs. (b) Unlabeled images would become hidden supplementary sources for potentially increasing the density of pairs. (c) Newly expanded image pairs can provide abundant densified points to alleviate the underlying data-hungry matter. (c) illustrates that a wealth of novel machine-annotated keypoint pairs (indicated by blue-type lines) are generated by simply incorporating new unlabeled images.

existence of $\hat{P}_{s,t}(\cdot, i)$; we have

$$\hat{M}_{s,t}(i) = \begin{cases} 1, & \text{if } ||\hat{P}_{s,t}(\cdot,i)||_{\infty} > 0, \\ 0, & \text{otherwise.} \end{cases}$$
(2)

In a supervised learning framework, (I_s, I_t) is tied together, so minimization in Eq.(1) gives a matching function f under the fixed and limited image pairs S.

2.2 Motivation

Previous methods traditionally aimed to design a novel matching network to gain a high-quality correlation map based on high-dimensional convolutions or Transformers. However, they leverage complicated learning frameworks with large models and distinct data augmentations due to relying on limited annotated keypoint pairs. Sophisticated learning frameworks [10, 14, 50] or heavy models [2, 3, 32, 41] with a matching function f, can fit a model to insufficient data space more; involving data augmentation methods [14] diversify the images I_s and I_t . However, they do not take into account the cardinality of the image pairs |S|; still, the recent unsupervised methods that try to densify keypoint pairs [10, 14] in each image pair do not consider the additional image pairs.

We argue using fixed annotated pairs in training inherently restricts performance; this is more likely because the annotated pairs are very sparse [4,35] (see Fig. 1(a)). Furthermore, the available image count is inadequate to offset the limited number of pairs. Given the constraints, we reframe the issue as a sample optimization problem instead of Eq.(1) as:

$$\min_{\theta, S'} \frac{1}{|\mathcal{S}'|} \sum_{(I_s, I_t) \in \mathcal{S}'} \sum_{i=1}^{H_t W_t} \hat{M}_{s,t}(i) \mathcal{D}(P_{s,t}(\cdot, i), \hat{P}_{s,t}(\cdot, i)),$$
(3)

where the objective has a newly added variable S' and corresponding supervision. However, a direct optimization of this problem seems like an NP-hard problem. Therefore, we relax the problem by managing the image pair variable to be expanded, having a lower bound of Eq.(1) by simply untying the link between source and target image pairs:

$$\min_{\theta} \frac{1}{|\mathcal{S}'|} \sum_{(I_s, I_t) \in \mathcal{S}'} \sum_{i=1}^{H_t W_t} \hat{M}_{s,t}(i) \mathcal{D}(P_{s,t}(\cdot, i), \hat{P}_{s,t}(\cdot, i)), \qquad (4)$$

where $S' \supseteq S$. Minimizing Eq.(4) will give the trained weight with a lower value than minimizing Eq.(1), where a larger cardinality has a lower objective value. Our concern now moves on to how to acquire additional pair sets in S' over the original paired images S.

3 Method

This section introduces how we enlarge the training pairs by effectively using unlabeled data that lack annotations, as illustrated in Fig. 2. It is worth noting that our method serves as a demonstration of the intended purpose, suggesting that more complex approaches could outperform ours.

3.1 Mining Untapped Annotation Gems

Suppose we have a superset that contains images with or without labels having C object classes; there are n_c samples in each class c. Ideally, $n_c(n_c - 1)$ pairs for each class could be utilized for supervised training (with labels). Namely, the possible image pairs for class c is

$$\mathcal{U}_c = \{ (I_s, I_t) \mid s \in c, t \in c, s \neq t \},\tag{5}$$

and the set of entire pairs in a training set is $\mathcal{U} = \bigcup_{c=1}^{C} \mathcal{U}_c$. Due to the impracticality of labeling full-image-pairs in \mathcal{U} , sparsely and partially labeled keypoint-level supervisions are typically available. A set of labeled data \mathcal{S}_c for supervised training consists of image pairs for each class label c, which can be defined as a subset of \mathcal{U}_c :

$$S_c = \{ (I_s, I_t) \in \mathcal{U}_c \mid ||\hat{M}_{s,t}||_{\infty} > 0 \},$$
(6)



Fig. 2: Schematic illustration of our method. Unlabeled images are iteratively labeled by a progressively evolving machine annotator, where incrementally increasing noise are injected to challenge the training. Therefore, by learning increasingly challenging images, the model's generalization ability continues to improve.

and the total set of image pairs is $S = \{S_c\}_{c=1}^C$. Prior methods [2,3,16,17,32,36,59]only trained with training image pairs S under the shared protocol based upon semantic correspondence benchmarks [4,35]. This bottlenecks the model's performance due to insufficient training image pairs to instantiate full dense correspondences under large intra-class variations [3,21,36]. Therefore, expanding the training pairs towards the entire pairs \mathcal{U} by mining new annotation gems is a straightforward yet underlying solution. There can be approaches to picking which image pairs to use in training, but we alternatively involve the image set closer to the entire image pairs \mathcal{U} to demonstrate that such a simple choice works. We further conjecture that this is more likely to cover different appearances and difficulty levels between image pairs, enabling the learned model to have improved generalizability.

To mine correspondence supervisions for novel training pairs having no human annotations, we utilize a self-training technique [53, 56] using a machine annotator $f_{\rm T}$, which is trained on labeled training pairs \mathcal{S} , generates rough correspondences $\bar{P}_{s,t} = f_{\rm T}(I_s, I_t; \theta_{\rm T}) \in \mathbb{R}^{H_s W_s \times H_t W_t}$ on newly involved training pairs where a large fraction of pairs do not belong to the existing training pairs.

The model f is then trained to learn the novel correspondences with the loss:

$$\mathcal{L}_{\mathcal{U}} = \sum_{(I_s, I_t) \in \mathcal{U}} \sum_{i=1}^{H_t W_t} \bar{M}_{s,t}(i) \mathcal{D}(P_{s,t}(\cdot, i), \bar{P}_{s,t}(\cdot, i)),$$
(7)

where $P_{s,t} = f(I_s, I_t; \theta)$ is the predicted correspondence from the model. We use $\overline{M}_{s,t}$ as the binary mask for gating the correspondence based on a confidencebased strategy, widely used in [39,44,52,54,57]. By using a matching confidence $\overline{C}_{s,t}(\cdot, i)$, which defined from pair-wise scores between all locations in I_s and i in I_t with soft-argmax operation, it is thus formulated as:

$$\bar{M}_{s,t}(i) = \begin{cases} 1, & \text{if } ||\bar{C}_{s,t}(\cdot,i)||_{\infty} > \tau, \\ 0, & \text{otherwise,} \end{cases}$$
(8)

where τ is the score threshold. The correspondences generated with high confidence encourage the model to be trained without erroneous supervision from ambiguous or textureless areas. In this work, different from the previous methods [2,3,9,21,25,32,34,36,59], trained on sparse keypoint pairs per limited labeled image pairs, we can significantly increase the number of image pairs available for training by including unlabeled data. Additionally, we use a large corpus of keypoint pairs, densely filling most object parts by using dense predictions from a machine annotator without the expense of manual annotation made by a human expert.

Our training objective is $\mathcal{L} = \mathcal{L}_{\mathcal{U}} + \lambda \mathcal{L}_{\mathcal{S}}$, combining the loss function in Eq.(1) and Eq.(7) to leverage both the existing labeled pairs and the ones annotated by machine annotator. The weighting parameter λ adjusts the learning dynamics between the losses; we use $\lambda=1$ for simplicity. Note that recent approaches [6, 30, 58] employed features from pre-trained text-to-image diffusion models for semantic correspondence, which somewhat resembles our machine annotator. However, our model is much lighter and is trained on much smaller data under a simpler setup.

3.2 Iterative Labeling and Training

Our framework enjoys further improvements via iterative training. Specifically, each iteration repeats the training process using the model trained in the previous iteration as a new machine annotator and trains a new model. We define *l*-th training iteration as $(f_{\rm T}^l, f^l)$, consisting of a pair of the machine annotator and in-training-model. The first generation of annotator model $f_{\rm T}^0$ is trained from scratch on labeled training pairs S. The subsequent annotator models use the model trained in the preceding generations, *i.e.*, $f_{\rm T}^l = f^{l-1}$.

During training, we augment the input images of the model to diversify them further and let the model learn with more challenging ones. We use photometric augmentation \mathcal{N}_p [2, 3, 14] and geometric augmentation \mathcal{N}_g [14, 40, 50], to the source and target image for training. As a result, the model is trained on more Table 1: Comparison with state-of-the-art methods on SPair-71k. Per-class and overall PCK ($\alpha_{bbox} = 0.1$) results are shown in the table. Numbers in bold indicate the best performance, and underlined ones are the second best. The averaged PCK of each MatchMe significantly improves the baseline by a large margin, surpassing the state-of-the-art methods. This superiority is mostly consistent across various regimes, including supervised regimes such as a recent work HCCNet, semi-supervised regimes such as SemiMatch and SCORRSAN, and unsupervised regimes such as DIFT.

Methods	All	Aero	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Dog	Horse	MBike	Person	Plant	Sheep	Train	TV
HPF [<mark>34</mark>]	28.2	25.2	18.9	52.1	15.7	38.0	22.8	19.1	52.9	17.9	33.0	32.8	20.6	24.4	27.9	21.1	15.9	31.5	35.6
SCOT [29]	35.6	34.9	20.7	63.8	21.1	43.5	27.3	21.3	63.1	20.0	42.9	42.5	31.1	29.8	35.0	27.7	24.4	48.4	40.8
DHPF [36]	37.3	38.4	23.8	68.3	18.9	42.6	27.9	20.1	61.6	22.0	46.9	46.1	33.5	27.6	40.1	27.6	28.1	49.5	46.5
PMD [26]	37.4	38.5	23.7	60.3	18.1	42.7	39.3	27.6	60.6	14.0	54.0	41.8	34.6	27.0	25.2	22.1	29.9	70.1	42.8
MMNet [59]	40.9	43.5	27.0	62.4	27.3	40.1	50.1	37.5	60.0	21.0	56.3	50.3	41.3	30.9	19.2	30.1	33.2	64.2	43.6
CHM [32]	46.3	49.6	29.3	68.7	29.7	45.3	48.4	39.5	64.9	20.3	60.5	56.1	46.0	33.8	44.3	38.9	31.4	72.2	55.5
PMNC [21]	50.4	54.1	35.9	74.9	36.5	42.1	48.8	40.0	72.6	21.1	67.6	58.1	50.5	40.1	54.1	43.3	35.7	74.5	59.9
CATs [2]	49.9	52.0	34.7	72.2	34.3	49.9	57.5	43.6	66.5	24.4	63.2	56.5	52.0	42.6	41.7	43.0	33.6	72.6	58.0
TransforMatcher [16]	53.7	59.2	39.3	73.0	41.2	52.5	66.3	55.4	67.1	26.1	67.1	56.6	53.2	45.0	39.9	42.1	35.3	75.2	68.6
CATs++ [3]	59.8	60.6	46.9	82.5	41.6	56.7	65.1	50.4	72.8	29.2	75.9	65.3	62.6	50.9	56.1	54.6	48.0	80.8	75.0
SemiMatch [14]	50.7	53.6	37.0	74.6	32.3	47.5	57.7	42.4	67.4	23.7	64.2	57.3	51.7	43.8	40.4	45.3	33.1	74.1	65.9
SCORRSAN [10]	55.3	57.1	40.3	78.3	38.1	51.8	57.8	47.1	67.9	25.2	71.3	63.9	49.3	45.3	49.8	48.8	40.3	77.7	69.7
$DIFT_{sd}$ [46]	52.9	61.2	53.2	79.5	31.2	45.3	39.8	33.3	77.8	34.7	70.1	51.5	57.2	50.6	41.4	51.9	46.0	67.6	59.5
HCCNet [17]	54.8	59.9	40.6	70.5	39.8	55.9	<u>65.1</u>	56.8	66.6	25.6	69.2	59.6	58.7	46.7	40.3	43.6	39.6	$\underline{82.2}$	65.4
MatchMe (ours)	62.0	63.4	51.1	83.2	44.8	53.1	66.9	53.4	<u>74.8</u>	<u>30.4</u>	76.8	66.6	68.1	55.2	60.7	59.1	48.1	84.9	75.9

challenging image pairs than those in previous steps, which enhances the model's generalization ability, resulting in a superior model that surpasses the performance of the previous step. It can be defined as $P_{s,t} = f(\mathcal{N}_p(I_s), \mathcal{N}_p(\mathcal{N}_g(I_t))))$. Photometric augmentation is applied to both the source and target images, while geometric augmentation is applied only to the target image, considering computational efficiency. The machine-generated labels $\bar{P}_{s,t}$ are warped to align spatial position changes by applying the same geometric transformation.

4 Experiments

4.1 Experimental Setups

Benchmarks. Experiments are conducted on three standard benchmarks for semantic correspondence learning: PF-PASCAL [4], PF-WILLOW [4], and SPair-71k [35] consisting of image pairs with human-annotated keypoints from 20, 4, and 18 categories, respectively. As in previous works [5], we split the PF-PASCAL dataset [4] into about 700, 300, and 300 images for training, validation, and testing, respectively. For the SPair-71k dataset [35], we use 53,340 for training, 5,384 for validation, and 12,234 for testing. To verify generalization capacity, the PF-WILLOW dataset [4] is used for testing only.

Evaluation metric. Following [34], the percentage of correct keypoint (PCK@ α_k) is used for the evaluation metric by setting α_k , a tolerance margin, having a value $\in \{0, 1\}$. PCK can be computed as the ratio of correctly estimated keypoint pairs to the total number of keypoint pairs using the Euclidean distance between them within the pixel margin $\alpha_k \cdot \max(H_k, W_k)$. By setting $k \in \{\text{img, bbox}\}, H_k$ and W_k are the width and height of either image or the object's bounding box.



Fig. 3: Qualitative results on SPair-71k in comparison with the competing SOTA methods. The point-to-point matches are drawn by linking key point pairs with line segments. Green and red lines denote correct and incorrect predictions with respect to the ground-truth pairs, respectively. We observe that ours outperforms the counterparts significantly across all the sample image pairs.

Implementation details. We demonstrate our proposed method with two simple baselines, CATs [2] and CATs++ [3]. We use the best-performing model's weight trained in a supervised setting on labeled data for the initial annotator to generate labels. Only weak photometric augmentations, such as color-jitter and gray-scale, are used with a given probability of 0.2 to prevent early over-fitting. We employ stronger data augmentations to benefit generated labels more. Following the literature [2, 3], a combination of strong photometric augmentation at a frequency of 0.4 is used along with geometric augmentation [22, 40, 48]. The confidence threshold τ for the generated correspondences is commonly set as 0.7 for all image pairs in training datasets. We pick the best trained model as a new machine annotator for successive iterative training.

	PF-	PASC	CAL	PF-WILLOW				
Methods		$\alpha_{\rm img}$		$\alpha_{ m bbox}$				
	0.05	0.1	0.15	0.05	0.1	0.15		
HPF	60.1	84.8	92.7	45.9	74.4	85.6		
DHPF	75.7	90.7	95.0	49.5	77.6	89.1		
MMNet	77.6	89.1	94.3	-	-	-		
CHM	80.1	91.6	-	52.7	79.4	-		
CATs	75.4	92.6	96.4	50.3	79.2	90.3		
TransforMatcher	-	80.8	91.8	-	76.0	-		
CATs++	84.9	93.8	96.8	56.7	81.2	-		
PMNC	82.4	90.6	-	-	-	-		
SemiMatch	80.1	93.5	96.6	54.0	82.1	92.1		
SCORRSAN	81.4	92.9	96.1	54.1	80.0	89.8		
DIFT_{sd}	-	-	-	58.1	81.2	-		
HCCNet	-	92.4	-	-	74.5	-		
MatchMe (ours)	84.9	94.3	96.7	59.6	83.6	92.9		

Table 2: Comparison with state-of-the-art methods on PF-PASCAL and PF-WILLOW. Numbers in bold denote the best, and underlined ones are the second best. MatchMe outperforms the competing methods again, like in Spair-71k.

4.2 Comparison on Benchmarks

We evaluate our method in comparison with the SOTA methods [2,3,16,17,32,59] trained based on the supervised protocol with existing keypoint annotations. We also compare with similar methods [10,14], using both generated labels and GT labels and a method without supervision [46]

On SPair-71k. Tab. 1 shows PCKs ($\alpha_{bbox}=0.1$) on all 18 object classes, including the overall mean PCK. Our overall averaged PCK=62.0% significantly outperforms the current state-of-the-art methods. We achieve +2.2% of PCK improvements over the baseline [3]. It demonstrates that the matching networks, especially with correlation enhancement architecture, have been under-trained with sparse and limited keypoint supervision. Furthermore, our consistently superior performance in both per-class and average PCK compared to state-of-the-art methods across various regimes [3, 10, 14, 17, 46] indicate improved generalizability, which facilitates handling large intra-class variation and deformation between instances within the same object class.

Additionally, as shown in Fig. 3, we visualize the sampled example pairs with the predicted matches for MatchMe and the competing methods showing the best performance in both the supervised regime, such as CATs [2], CATs++ [3] and the regime similar to ours, such as SemiMatch [14], SCORRSAN [10], using both supervision from generated labels and GT keypoint labels. This demonstrates that our method estimates correspondences between image pairs more accurately than others, even under a large discrepancy in viewpoint and scale.

On PF-PASCAL and PF-WILLOW. Tab. 2 summarizes our results on the PF-PASCAL and PF-WILLOW datasets compared with the other competing

Table 3: PCK comparison among training methods. For a fair comparison, we use the fixed baseline CATs++ [3] for all semi-supervised training methods. While all lead to performance improvements, ours enjoy the most significant enhancement, which highlights the need to prioritize the data-hungry matter.

	Methods	PCK
(a)	Baseline	59.8
(b) (c) (d) (e)	 (a) + CNNGeoU [19] (a) + PWarpC [50] (a) + SCORRSAN [10] (a) + MatchMe (ours) 	$\begin{array}{c} 60.1 \ (+0.3) \\ 60.5 \ (+0.7) \\ 61.0 \ (+1.2) \\ \textbf{62.0} \ (+2.2) \end{array}$

methods trained on PF-PASCAL from each initialized model (*i.e.*, usually pretrained on ImageNet [42]). We also fine-tune our model, pre-trained on SPair-71k with the unlabeled data from PASCAL VOC 2012, on the PF-PASCAL dataset to evaluate the generalization capability of our model on different datasets. MatchMe records the new state-of-the-art PCK value 94.3 that beats the previous state-of-the-art value of 93.8, which is almost saturated, on PF-PASCAL.

On PF-WILLOW, MatchMe outperform the baseline [3] by 2.9% / 2.4%(α =0.05/0.1), surpassing other competing methods across different regimes. Note that our method not only achieves higher PCKs than the competing methods on PF-PASCAL but also outperforms them by a more significant margin on PF-WILLOW. This signifies the generalization capability of our method and discloses that ours learns a general representation, which can be applied to various datasets, different from the baselines usually overfitted on a specific dataset. **Controlled experiments for learning methods.** We conduct controlled comparisons between our method and existing semi-supervised methods [10, 19, 49]. All the methods are trained with the fixed baseline CATs++ [2] for a fair comparison to evaluate the methods' uniqueness in improving each method's performance without any potential influence from the model difference. We use SPair-71k [35], which contains fixed, sparsely-annotated pairs, for a comprehensive comparison. We strive to report the best results for each method via parameter searches.

Tab. 3 first shows MatchMe outperforms all the competitors. Specifically, (a) and (b), using a cycle consistency, and (c), using generated labels as sources of unsupervised loss signal, show limited performance improvement compared to the baseline supervised learner (a). This is attributed to their narrow focus on augmenting labels within a limited amount of labeled data. Unlike them, our method (d) focuses on densifying generated labels even with *unlabeled* data as well as labeled data, thereby highly boosting performance.

4.3 Analyzing Our Method

Robustness evaluation. Here, we construct a new benchmark for semantic correspondence estimation robustness, named SPair-C, following the regime [7].

Table 4: Robust evaluation on SPair-C. We report the PCK numbers of our model and the baseline evaluated on clean (SPair-71k) and the newly introduced corrupted dataset (Spair-C), using the pre-trained weights of our model and the baseline model [2] provided by official code on SPair-71k.We report detailed PCKs for all the corruptions and further report the averaged PCKs to facilitate comparison. Numbers in bold indicate the best performance, and underlined ones are the second best.

Mathada Sar		Noise			ise Blur			Weather				Digital				Corrun	Clean	
Methods	Bev.	Gauss.	Shot	Impulse	Speckle	Defocus	Gaussian	Snow	Frost	Fog	Spatter	Bright	Contrast	Saturate	Pixel	JPEG	Corrup.	Clean
	1	45.0	44.8	42.0	45.9	41.0	45.5	39.3	42.3	43.7	48.3	48.8	46.3	48.9	47.7	46.2	45.0	
	2	40.7	41.1	38.1	44.3	35.9	38.1	28.8	34.6	41.9	42.7	48.1	44.4	48.1	46.8	44.5	41.2	
	3	35.3	35.5	35.2	37.9	27.7	29.8	28.8	29.1	38.3	37.5	47.3	40.9	48.4	43.3	43.8	37.3	40.0
Dasenne	4	27.9	27.8	27.9	34.1	21.6	23.6	23.6	27.9	35.7	35.1	46.1	31.3	46.5	36.5	40.4	32.4	49.9
	5	19.7	22.7	21.8	29.4	17.1	16.1	23.9	25.4	27.6	29.0	44.1	20.5	43.9	33.6	37.0	27.5	
	avg.	33.7	34.4	33.0	38.3	28.7	30.6	28.9	31.9	37.4	38.5	46.9	36.7	47.2	41.6	42.4	36.7	
	1	47.9	48.5	45.7	49.5	46.9	50.4	44.2	47.0	49.3	51.0	52.3	50.8	52.1	50.7	49.8	49.1	
	2	44.7	45.2	42.4	47.8	43.4	44.9	35.7	40.6	47.4	46.0	51.7	49.6	51.6	50.5	48.8	46.0	
MotohMo	3	39.5	40.2	39.5	41.9	36.5	39.2	35.3	34.6	44.9	42.1	51.5	46.8	51.7	45.5	47.0	42.4	52.0
Matchne	4	31.4	30.9	30.8	38.3	29.2	31.6	30.1	33.2	43.5	39.1	50.3	37.5	50.1	38.0	43.5	37.2	33.0
	5	20.4	24.3	22.5	32.6	22.6	19.3	31.1	29.5	36.2	32.6	48.2	17.4	39.2	34.9	39.2	30.0	
	avg.	36.8	37.8	36.2	42.0	35.7	37.1	35.3	37.0	44.3	42.2	50.8	40.4	48.9	43.9	46.1	40.9	1

We highlight that SPair-C is the first dataset with corruption and noise for dense correspondence learning and will be useful for future evaluations. We verify our method's robustness by evaluating whether the model can robustly predict correspondences on corrupted images. Additional details for the dataset are provided in the supplementary material.

Tab. 4 shows the overall PCK values throughout 15 corruptions in SPair-C for MatchMe with CATs versus CATs. We observe that MatchMe consistently outperforms the baseline [2] in terms of the average PCK values across different severities of a single corruption (located in the column of the table) and across different types of corruption with the same severity (shown in the row of the table). Moreover, the average PCK for all 75 corruptions of MatchMe is 40.9, which represents a 4.2 improvement over the baseline, surpassing the gap of 3.1 observed in the clean SPair dataset. This improvement is presumably attributed to novel labeled keypoints mined by our method, which allowed the model to extract more robust features, even in the presence of corrupted pixels. As a result, the model's robustness to typical corruptions, commonly found in natural images, further reinforces our generalization capability effectively.

Impact of iterative training. Fig. 4 shows the improved PCK values as iterative training progresses, which empirically proves the effectiveness of our iterative training framework. We verify the effectiveness using two baseline models, including CATs and CATs++. As observed in Fig. 4, performance consistently improves with each iteration, and the difference between the first and the last iterations is nearly 3% under the same training hyperparameters. This suggests that improved annotations from stronger models on challenging images by more intense data augmentations can effectively boost performance. This exposes the data-hungry issue in earlier baselines, suggesting they were likely undertrained.

Training time efficiency. A pertinent question regarding our method is the extent of overall training time required by our iterative learning protocol. For-



Fig. 4: PCK at each iteration in iterative training. We report PCK values at each iteration to show the effectiveness of our training framework. We use identical architecture for the teacher and student and set the iterative training interval to 50 epochs for simplicity. The left and right figures are the results of MatchMe trained upon CATs and CATs++ backbones, respectively. This indicates that the baseline models were, in fact, undertrained and possess the capacity for further training, highlighting the data-hungry problem.

tunately, due to the efficiency of our approach to handling the expanded data, this is not a significant concern. Our method's multi-iteration training with more data may increase the training time, but our rapid convergence (8 for ours vs. 42 epochs for [10]) offsets this, maintaining time efficiency despite repeated iterations.

PCK analysis by variation factors. The averaged PCK is insufficient to demonstrate the performance of the matching models precisely because it is evaluated without considering variation factors. Therefore, the desirable model should show an even PCK performance among the diverse difficulty levels under various factors. To confirm the accurate performance comparison, we conduct the PCK analysis based on the variation factors and difficulty levels on SPair-71k in Tab. 5. In experiments, our models show robustness by maintaining higher PCK values consistently across different difficulty levels for various factors compared to the other models. This proves that a large amount of unlabeled data used by our method can cover a wide range of data distribution, including diverse difficulty levels and scene variations.

Ablation study. We comprehensively analyze each component in our method in Tab. 6. For a fair comparison, we train all the variants on SPair-71k under the same experimental setting. The ablation results show the impacts of each component consisting of MatchMe. Compared to (a) MatchMe, (b) demonstrates the impact of iterative training by running the same epochs to show the impact (*i.e.*,

Table 5: PCK analysis of state-of-the-art methods on SPair-71k. All methods commonly show lower PCK than average PCK as difficulty levels of labeled data become more difficult, but MatchMe shows the best at each level in all but two.

Mathada	A 11	View-point				Scale	,	Truncation				Occlusion				
Methods		Easy	Medi	Hard	Easy	Medi	Hard	None	Src	Tgt	Both	None	Src	Tgt	Both	
NC-NET [41]	26.4	34.0	18.6	12.8	31.7	23.8	14.2	29.1	22.9	23.4	21.0	29.0	21.1	21.8	19.6	
HPF [34]	28.2	35.6	20.3	15.5	33.0	26.1	15.8	31.0	24.6	24.0	23.7	30.8	23.5	22.8	21.8	
CATs [2]	49.9	54.6	44.5	43.6	54.5	49.2	36.3	53.7	42.0	49.4	40.6	52.7	44.5	44.9	42.0	
CATs++[3]	59.8	63.5	55.9	53.0	62.8	59.6	50.2	61.8	55.2	56.3	54.8	63.6	52.4	57.9	50.8	
SemiMatch [14]	50.8	54.8	44.1	46.2	55.3	50.2	36.6	54.2	43.2	50.0	42.7	53.6	45.1	44.9	43.3	
SCORRSAN [10]	55.3	59.2	51.2	48.9	58.7	55.0	45.0	59.2	46.1	55.0	46.9	57.8	50.2	50.7	48.7	
MatchMe (ours)	62.0	66.0	57.9	55.0	64.7	61.4	54.4	65.5	54.5	60.9	55.1	64.3	57.3	57.6	56.2	

Table 6: Ablation study with the components. We perform a factor analysis of the elements used for training MatchMe-CATs. We compute PCKs on SPair-71k for each component. We observe all the components contribute to the PCK improvements.

	Componenta	$\alpha_{ m b}$	box
	Components	0.05	0.10
(a)	$\texttt{MatchMe} \ (ours)$	29.6	53.0
(b)	(a) - Iterative training	28.8 (-0.8)	52.0 (-1.0)
(c)	(b) - Data noise	28.5 (-0.3)	51.6 (-0.4)
(d)	(c) - Unpaired data	26.9 (-1.6)	49.9 (-1.7)

(b) trained at once, and (a) trained for a total of three iterative training divided by 50 epochs). (c) shows that data noise also contributes to the performance of MatchMe. (d) shows the benefits that novel generated labels are densified at pixel-level and image-level, respectively, by showing a large margin of 1.6, compared to (c). The result demonstrates that the data-hungry issue in the semantic correspondence task is the most crucial in performance degradation since their performance gains are more significant than others.

5 Related Work

Semantic correspondence learning. Recent methods for semantic correspondence [2,3,16,17,25,26,29,32,36,59] inevitably train complicated matching networks to maximize performance in a supervised manner with limited qualified dataset [4,35], which leads to high computational demands and poor generalization capability across datasets.

Some unsupervised strategies [10, 14, 19, 50] extend their unsupervised loss to the supervised regime and significantly improve the performance of the previous supervised approaches. This shows that the performance of the existing supervised model was not fully learned due to a lack of data. Specifically, the methods [19, 50] use a cycle consistency for unsupervised loss signal, and the

others [10, 14] utilize pseudo-labels, generated by the model's prediction between real images, combined with confidence measures to guarantee the quality of pseudo-labels. Recently, methods [6, 30, 58] have emerged that implement unsupervised semantic correspondence by tapping into the inherent knowledge embedded in pre-trained diffusion models trained on large text-image datasets to facilitate semantic correspondence. Leveraging the knowledge in pre-trained models for unseen data may share a similar spirit, but our framework does not require billion-scale data for generative modeling like those. Our method employs a lighter pre-trained model as a guidance labeler to address the scarcity of both image and point pairs by utilizing massive unlabeled data through the pre-trained model.

Semi-supervised learning methods for semantic correspondence learning. Previous literature [24, 37, 47, 53] were proposed to utilize a teacher-student structure mainly for semi-supervised learning. Those methods are distinguished from the earlier methods pseudo-labeling [20, 43] and consistency regularization methods [38] due to their use of data and the resulting performance achieved. A teacher model, generally trained on a small set of labeled data, generates pseudolabels on a larger unlabeled data to guide the student model, and then the student is jointly trained on a combination of labeled and pseudo-labeled images. Recent studies have applied the teacher-student framework for pixel-level semisupervised learning, specifically for the semantic correspondence task [10, 26]. They employ a teacher model to generate additional pseudo-labels using knowledge from keypoint periphery [10] or hypotheses [26] across labeled image pairs. On the other hand, our method labels overlooked unlabeled data using a machine annotator, continually repeating the process by assigning the learned student back to the teacher.

6 Conclusion

In this paper, we have proposed a simple baseline that leverages unpaired images for semantic correspondence learning. Instead of using a sizeable complicated model with strong data augmentations to augment paired images, we have aimed to break the stereotype of using given labeled image pairs by expanding the training pairs with machine-annotated unpaired images. Only with a machineannotation-based framework for labeling the unpaired images, our method could beat the state-of-the-art models on SPair-71k, PF-PASCAL, and PF-WILLOW by large margins. Additionally, our approach could continuously improve performance by repeating the training process with increasingly challenging image pairs after each step. It also turns out that a resultant model has become more robust to corrupted images.

Limitations. Going beyond the scale of Spair-71k and tackling more challenging datasets unrelated to the semantic correspondence task would reveal a more generalized impact of our work. Furthermore, an exciting direction can be utilizing recently proposed Transformer-based architectures to deal with unpaired data with expanded data.

References

- Bristow, H., Valmadre, J., Lucey, S.: Dense semantic correspondence where every pixel is a classifier. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 4024–4031 (2015) 1
- Cho, S., Hong, S., Jeon, S., Lee, Y., Sohn, K., Kim, S.: Semantic correspondence with transformers. arXiv preprint arXiv:2106.02520 (2021) 1, 2, 3, 5, 6, 7, 8, 9, 10, 11, 13
- Cho, S., Hong, S., Kim, S.: Cats++: Boosting cost aggregation with convolutions and transformers. arXiv preprint arXiv:2202.06817 (2022) 1, 2, 3, 5, 6, 7, 8, 9, 10, 13
- Ham, B., Cho, M., Schmid, C., Ponce, J.: Proposal flow: Semantic correspondences from object proposals. IEEE Transactions on Pattern Analysis and Machine Intelligence 40(7), 1711–1725 (2017) 1, 2, 3, 5, 7, 13
- Han, K., Rezende, R.S., Ham, B., Wong, K.Y.K., Cho, M., Schmid, C., Ponce, J.: Scnet: Learning semantic correspondence. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1831–1840 (2017) 7
- Hedlin, E., Sharma, G., Mahajan, S., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: Unsupervised semantic correspondence using stable diffusion. Advances in Neural Information Processing Systems 36 (2024) 6, 14
- Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. arXiv preprint arXiv:1903.12261 (2019) 10
- Hosni, A., Rhemann, C., Bleyer, M., Rother, C., Gelautz, M.: Fast cost-volume filtering for visual correspondence and beyond. IEEE Transactions on Pattern Analysis and Machine Intelligence 35(2), 504–511 (2012) 1
- Huang, S., Wang, Q., Zhang, S., Yan, S., He, X.: Dynamic context correspondence network for semantic alignment. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2010–2019 (2019) 6
- Huang, S., Yang, L., He, B., Zhang, S., He, X., Shrivastava, A.: Learning semantic correspondence with sparse annotations. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIV. pp. 267–284. Springer (2022) 2, 3, 7, 9, 10, 12, 13, 14
- Hui, T.W., Tang, X., Loy, C.C.: Liteflownet: A lightweight convolutional neural network for optical flow estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8981–8989 (2018) 1
- Hur, J., Lim, H., Park, C., Chul Ahn, S.: Generalized deformable spatial pyramid: Geometry-preserving dense correspondence estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1392–1400 (2015) 1
- Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: Flownet 2.0: Evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2462–2470 (2017) 1
- Kim, J., Ryoo, K., Seo, J., Lee, G., Kim, D., Cho, H., Kim, S.: Semi-supervised learning of semantic correspondence with pseudo-labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19699– 19709 (2022) 2, 3, 6, 7, 9, 13
- Kim, S., Min, D., Jeong, S., Kim, S., Jeon, S., Sohn, K.: Semantic attribute matching networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 12339–12348 (2019) 1

- 16 Kim et al.
- Kim, S., Min, J., Cho, M.: Transformatcher: Match-to-match attention for semantic correspondence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8697–8707 (2022) 1, 5, 7, 9, 13
- Kim, S., Min, J., Cho, M.: Efficient semantic matching with hypercolumn correlation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 139–148 (2024) 1, 5, 7, 9, 13
- Kokkinos, F., Kokkinos, I.: Learning monocular 3d reconstruction of articulated categories from motion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1737–1746 (2021) 1
- Laskar, Z., Kannala, J.: Semi-supervised semantic matching. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018) 2, 10, 13
- Lee, D.H., et al.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on challenges in representation learning, ICML. p. 896 (2013) 14
- Lee, J.Y., DeGol, J., Fragoso, V., Sinha, S.N.: Patchmatch-based neighborhood consensus for semantic correspondence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 13153–13163 (2021) 5, 6, 7
- Lee, J., Kim, D., Ponce, J., Ham, B.: Sfnet: Learning object-aware semantic correspondence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2278–2287 (2019) 8
- Lee, J., Kim, E., Lee, Y., Kim, D., Chang, J., Choo, J.: Reference-based sketch image colorization using augmented-self reference and dense semantic correspondence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5801–5810 (2020) 1
- Li, H., Wu, Z., Shrivastava, A., Davis, L.S.: Rethinking pseudo labels for semisupervised object detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 1314–1322 (2022) 2, 14
- Li, S., Han, K., Costain, T.W., Howard-Jenkins, H., Prisacariu, V.: Correspondence networks with adaptive neighbourhood consensus. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10196–10205 (2020) 1, 6, 13
- Li, X., Fan, D.P., Yang, F., Luo, A., Cheng, H., Liu, Z.: Probabilistic model distillation for semantic correspondence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7505–7514 (2021) 1, 7, 13, 14
- Li, X., Liu, S., De Mello, S., Kim, K., Wang, X., Yang, M.H., Kautz, J.: Online adaptation for consistent mesh reconstruction in the wild. Advances in Neural Information Processing Systems 33, 15009–15019 (2020) 1
- Liu, C., Yuen, J., Torralba, A.: Sift flow: Dense correspondence across scenes and its applications. IEEE Transactions on Pattern Analysis and Machine Intelligence 33(5), 978–994 (2010) 1
- Liu, Y., Zhu, L., Yamada, M., Yang, Y.: Semantic correspondence as an optimal transport problem. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4463–4472 (2020) 1, 7, 13
- Luo, G., Dunlap, L., Park, D.H., Holynski, A., Darrell, T.: Diffusion hyperfeatures: Searching through time and space for semantic correspondence. Advances in Neural Information Processing Systems 36 (2024) 6, 14
- Melekhov, I., Tiulpin, A., Sattler, T., Pollefeys, M., Rahtu, E., Kannala, J.: Dgcnet: Dense geometric correspondence network. In: IEEE Winter Conference on Applications of Computer Vision. pp. 1034–1042. IEEE (2019) 1

- Min, J., Cho, M.: Convolutional hough matching networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2940–2950 (2021) 1, 3, 5, 6, 7, 9, 13
- Min, J., Kang, D., Cho, M.: Hypercorrelation squeeze for few-shot segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6941–6952 (2021) 1
- Min, J., Lee, J., Ponce, J., Cho, M.: Hyperpixel flow: Semantic correspondence with multi-layer neural features. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3395–3404 (2019) 1, 6, 7, 13
- Min, J., Lee, J., Ponce, J., Cho, M.: Spair-71k: A large-scale benchmark for semantic correspondence. arXiv preprint arXiv:1908.10543 (2019) 1, 2, 3, 4, 5, 7, 10, 13
- Min, J., Lee, J., Ponce, J., Cho, M.: Learning to compose hypercolumns for visual correspondence. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16. pp. 346–363. Springer (2020) 1, 5, 6, 7, 13
- Pham, H., Dai, Z., Xie, Q., Le, Q.V.: Meta pseudo labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2021) 2, 14
- Rasmus, A., Berglund, M., Honkala, M., Valpola, H., Raiko, T.: Semi-supervised learning with ladder networks. Advances in Neural Information Processing Systems 28 (2015) 14
- Rizve, M.N., Duarte, K., Rawat, Y.S., Shah, M.: In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. arXiv:2101.06329 (2021) 6
- Rocco, I., Arandjelovic, R., Sivic, J.: Convolutional neural network architecture for geometric matching. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6148–6157 (2017) 2, 6, 8
- Rocco, I., Cimpoi, M., Arandjelovic, R., Torii, A., Pajdla, T., Sivic, J.: Ncnet: Neighbourhood consensus networks for estimating image correspondences. IEEE Transactions on Pattern Analysis and Machine Intelligence (2020) 3, 13
- 42. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV (2015) 2, 10
- Shi, W., Gong, Y., Ding, C., Tao, Z.M., Zheng, N.: Transductive semi-supervised deep learning using min-max features. In: Proceedings of the European Conference on Computer Vision. pp. 299–315 (2018) 14
- 44. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In: Advances in Neural Information Processing Systems (2020) 6
- 45. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8934–8943 (2018) 1
- Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. arXiv preprint arXiv:2306.03881 (2023) 7, 9
- 47. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in Neural Information Processing Systems (2017) 2, 14
- Truong, P., Danelljan, M., Timofte, R.: Glu-net: Global-local universal network for dense flow and correspondences. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6258–6268 (2020) 8

- 18 Kim et al.
- Truong, P., Danelljan, M., Yu, F., Van Gool, L.: Warp consistency for unsupervised learning of dense correspondences. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 10346–10356 (2021) 10
- Truong, P., Danelljan, M., Yu, F., Van Gool, L.: Probabilistic warp consistency for weakly-supervised semantic correspondences. arXiv preprint arXiv:2203.04279 (2022) 2, 3, 6, 10, 13
- Xie, G.S., Xiong, H., Liu, J., Yao, Y., Shao, L.: Few-shot semantic segmentation with cyclic memory network. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7293–7302 (2021) 1
- Xie, Q., Dai, Z., Hovy, E., Luong, T., Le, Q.: Unsupervised data augmentation for consistency training. In: Advances in Neural Information Processing Systems (2020) 6
- Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020) 2, 5, 14
- Xu, Y., Shang, L., Ye, J., Qian, Q., Li, Y.F., Sun, B., Li, H., Jin, R.: Dash: Semisupervised learning with dynamic thresholding. In: International Conference on Machine Learning (2021) 6
- Yang, G., Ramanan, D.: Volumetric correspondence networks for optical flow. Advances in Neural Information Processing Systems 32 (2019) 1
- 56. Yun, S., Oh, S.J., Heo, B., Han, D., Choe, J., Chun, S.: Re-labeling imagenet: from single to multi-labels, from global to localized labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2340– 2350 (2021) 2, 5
- 57. Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., Shinozaki, T.: Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In: Advances in Neural Information Processing Systems (2021) 6
- Zhang, J., Herrmann, C., Hur, J., Polania Cabrera, L., Jampani, V., Sun, D., Yang, M.H.: A tale of two features: Stable diffusion complements dino for zero-shot semantic correspondence. Advances in Neural Information Processing Systems 36 (2024) 6, 14
- Zhao, D., Song, Z., Ji, Z., Zhao, G., Ge, W., Yu, Y.: Multi-scale matching networks for semantic correspondence. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3354–3364 (2021) 1, 5, 6, 7, 9, 13