

# Do It Yourself: Learning Semantic Correspondence from Pseudo-Labels

Olaf Dünkel<sup>1</sup>, Thomas Wimmer<sup>1,2</sup>, Christian Theobalt<sup>1,4</sup>, Christian Rupprecht<sup>3</sup>, Adam Kortylewski<sup>1,5</sup>

<sup>1</sup>Max Planck Institute for Informatics, Saarland Informatics Campus <sup>2</sup>ETH Zurich <sup>3</sup>University of Oxford

<sup>4</sup>Saarbrücken Research Center for Visual Computing, Interaction and AI <sup>5</sup>University of Freiburg

<https://genintel.github.io/DIY-SC>

## Abstract

Finding correspondences between semantically similar points across images and object instances is one of the everlasting challenges in computer vision. While large pre-trained vision models have recently been demonstrated as effective priors for semantic matching, they still suffer from ambiguities for symmetric objects or repeated object parts. We propose improving semantic correspondence estimation through 3D-aware pseudo-labeling. Specifically, we train an adapter to refine off-the-shelf features using pseudo-labels obtained via 3D-aware chaining, filtering wrong labels through relaxed cyclic consistency, and 3D spherical prototype mapping constraints. While reducing the need for dataset-specific annotations compared to prior work, we establish a new state-of-the-art on SPair-71k, achieving an absolute gain of over 4% and of over 7% compared to methods with similar supervision requirements. The generality of our proposed approach simplifies the extension of training to other data sources, which we demonstrate in our experiments.

## 1. Introduction

Finding correspondences between images remains a fundamental task in computer vision, having various applications in tracking [8, 10, 20], mapping and localization [24, 37], affordance understanding [25], pose estimation [58], analysis of characteristic object parts [16, 50], image and video generation or editing [12, 39, 53], 3D representation learning [31, 46, 56], or style transfer [23]. One remaining challenge in the correspondence estimation field is the task of finding matches across different instances of similar objects, *i.e.*, finding semantic correspondences. This is a highly semantic task with a certain amount of ambiguity, especially for man-made objects.

Recently, foundation model features have demonstrated surprisingly high zero-shot performance for this task [49, 62, 63]. However, these features still have various weaknesses for finding correspondences, such as ambiguities for



Figure 1. Our method, DIY-SC, is able to find semantic correspondences even for extreme appearance and shape changes. All matches are computed from the keypoints in the top-left image.

similar object parts [35] or symmetric objects [63]. In addition, supervised methods [63] indicate that there are more effective strategies to combine foundational features for solving semantic correspondence than simple concatenation [62] or weighted averaging [35]. Recent works [33, 59, 63] train models in a supervised manner using manual keypoint annotations and achieve strong performance when evaluated on the same dataset. However, as the required manual annotations are scarce and difficult to obtain, this strategy is not scalable to larger, more diverse datasets. It thus remains an open challenge to determine the best way to extract and enhance knowledge encoded in foundation models for finding semantic correspondences *without* relying on labor-intensive keypoint supervision.

Recent works reduce the necessary level of supervision for learning semantic correspondences by regularizing with

3D information [35] or injecting information about label definitions [9, 63]. However, they either require careful tuning of weighting factors [35] or rely on access to dataset-specific labeling conventions [63].

We address both limitations in our work by learning features that significantly outperform the previous SOTA without requiring manual keypoint label definitions: Our work demonstrates that training with self-generated labels, *i.e.*, pseudo-labels, where label quality is improved through weak 3D-aware supervision, is surprisingly effective for improving the semantic correspondence performance, also generalizing to heavy appearance changes (Fig. 1). To achieve this, we train a lightweight adapter supervised by the pseudo-labels, which effectively refines the foundation model’s features, as shown in Fig. 2. Specifically, we pursue a zero-shot approach [62] for generating pseudo-labels for image pairs exhibiting moderate viewpoint variation, a strategy that has been demonstrated to be effective in such contexts but fails for larger viewpoint variations [63]. We compose these labels over multiple image pairs to acquire higher label quality for harder correspondence pairs, *i.e.*, with larger viewpoint variance. We then employ a spherical prototype as a weak geometric prior [35] to reject wrong matches, which addresses the inherent challenge of *matchability* and reduces feature ambiguity.

In contrast to previous weakly supervised methods [9, 63], our approach does not rely on dataset-specific keypoint definitions, reducing the barriers to applying it to other datasets. To demonstrate this, we also train a model on the recent ImageNet-3D dataset [34]. Pre-training on this larger dataset notably also improves the performance on SPair-71k [36], demonstrating the effectiveness of our more generalizable approach.

To summarize our contributions:

- We show that pseudo-labeling is effective for learning better semantic correspondence features, and we demonstrate that the quality of pseudo-labels can be improved through 3D-aware chaining, relaxed cyclic-consistency constraints, and the integration of a weak geometric prior.
- We show that our strategy is scalable to a larger dataset, further improving the model’s performance.
- Finally, we set a new SOTA on SPair-71k with weak supervision, outperforming the previous best model by 4.5 absolute points.

## 2. Related Work

**Semantic correspondence.** Semantic matching, *i.e.*, finding correspondences between *different* instances of the same object class, is more challenging than geometric matching due to potential variations in appearance and shape. The scarcity and ambiguity of manually annotated data pose an extra challenge to learning-based methods for this task [51, 65].

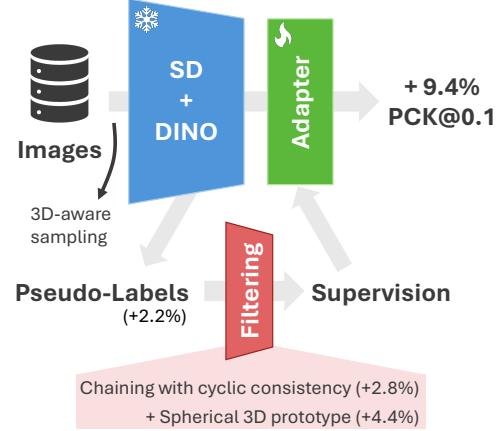


Figure 2. **Self-training using pseudo-labels.** Using foundational features, we generate pseudo-labels, which are subsequently filtered and used as supervision signals to train a light-weight adapter that refines the features for semantic matching.

While early works relied on hand-crafted descriptors [30, 32], deep learning enabled learning better-suited feature extractors [22, 38, 44, 60] and direct semantic correspondence detection networks [14, 23, 43]. Due to the limited ground truth data, techniques emerged that rely on weak supervision [6, 26, 52, 64], leverage warp supervision and cycle consistency losses [51, 52, 66], or rely on increasing label supervision with pseudo-labels [18, 21, 28].

With the advent of so-called foundation models, the semantic matching literature has come full circle: Recent studies demonstrated that features obtained from such models can be utilized for detecting semantic correspondences in a zero-shot manner [1, 2, 7, 15, 49, 62]. More specifically, DINO features [5, 40] have been shown to exhibit strong coarse semantic awareness useful for handling larger, cross-instance variations [2, 9, 48, 59, 62, 63], while features from diffusion models [45, 47] can complement them, resulting in better performance [9, 15, 29, 35, 49, 59, 62, 63].

While a simple nearest-neighbor search in feature space proves to be a powerful zero-shot method for correspondence estimation, prior work also uncovered systematic problems of this approach, *e.g.*, in the disambiguation of symmetric object parts [29, 33, 35, 46, 55, 63]. Therefore, several works have proposed appending adapter modules that are fine-tuned using supervision with ground-truth correspondences [59, 63]. While reducing the labeling requirements, methods that build a joint atlas for objects in multiple images usually require larger quantities of input images and perform test-time optimization [11, 39]. Zhang et al. [63] use keypoint-specific information to disambiguate left-right symmetries at test time. While their strategy is simple, not all symmetries can be solved by image transformations, *e.g.*, flipping, and the required keypoint-specific information is generally not available. Fundel et al. [9]

proposed fine-tuning distilled foundational features for geometric matching using 3D ground truth data. This approach is, however, inherently limited to the same object instance, resulting in suboptimal results for cross-instance matching. Mariotti et al. [35] proposed learning a map from features to the surface of a sphere using 3D pose information as a weak supervision signal. While this method successfully resolves most symmetries for simple objects, such as cars, a spherical prior does not accurately capture objects with more complex topologies.

Our work also aims to improve the performance of semantic matching by refining foundational features. In contrast to Zhang et al. [63], we do not require dataset-specific keypoint definitions since our work only relies on weak supervision that is available at scale. Similar to Mariotti et al. [35], we leverage 3D information to reduce ambiguities. However, our work does not suffer from performance drops for objects with more complex topology and thus does not require tuning of weighting factors

**Pseudo-labeling.** Lee et al. [27] proposed pseudo-labeling, also referred to as self-training, as a technique for semi-supervised training of neural networks, where a small initial set of labels is propagated to a larger set of unlabeled instances. Various methods have been developed based on this concept, demonstrating the generality and scalability of this approach [3–5, 13, 20, 57], sometimes in conjunction with knowledge distillation from other pre-trained models. Previous works have expanded manual keypoint annotations and demonstrated improved semantic correspondence performance [18, 21, 28]. In our work, we propose to extract pseudo-labels using off-the-shelf foundational features (at training time), which we use to fine-tune a light-weight adapter for the task of semantic correspondence estimation.

### 3. Method

Foundational features show strong semantic awareness for zero-shot matching across instances. However, while the high feature similarity of similar object parts is reasonable from a semantic perspective, distinguishing them is crucial for accurate semantic correspondences. The performance gap between supervised and unsupervised semantic matching methods [63] shows that a learned combination and refinement of features outperforms a simple concatenation or averaging of multiple foundational features. As we aim to supervise a refinement stage without relying on human annotations, we draw on the effectiveness of pseudo-labeling for other tasks [4, 19, 20] and hypothesize its potential for improving features for semantic correspondence.

In this section, we first formalize the semantic correspondence task (Sec. 3.1). Then, we introduce our two-stage method: First, pseudo-labels are generated via a zero-shot matching strategy and subsequently filtered (Sec. 3.2). These pseudo-labels are then used to supervise the training

of a light-weight adapter  $f_p$  (Sec. 3.3). Finally, we detail how our method can be scaled to a larger dataset in Sec. 3.4.

#### 3.1. Problem Definition

The problem of finding semantic correspondences is defined as follows: Given a source image  $I^s$ , a target image  $I^t$ , and a query point  $p_i^s \in \mathbb{R}^2$  on the object in the source image, the task is to find the target point  $p_i^t \in \mathbb{R}^2$  on the target image that localizes the object part that semantically corresponds to the queried object part. One strategy to solve this task is to find the nearest neighbor in feature space via the cosine similarity  $\text{sim}(\cdot)$  of the feature maps  $\mathcal{F}^s$  and  $\mathcal{F}^t$ :

$$p_i^t = \text{NN}^{s \rightarrow t}(p_i^s) = \arg \max_{q_i^t \in I^t} \text{sim}(\mathcal{F}^s(p_i^s), \mathcal{F}^t(q_i^t)). \quad (1)$$

However, in contrast to zero-shot strategies that solely apply pre-trained feature extractors [62], we aim to improve the features such that the refined features are better suited for finding semantic correspondences. Since the features  $\mathcal{F}^{\text{DINO}}$  [40] and  $\mathcal{F}^{\text{SD}}$  [45] generalize well due to their large training corpus, we learn a light-weight adapter  $\mathcal{F} = f_p(\tilde{\mathcal{F}})$  on top of the concatenated features  $\tilde{\mathcal{F}} = [\mathcal{F}^{\text{DINO}}, \mathcal{F}^{\text{SD}}]$ , as proposed by Zhang et al. [63].

#### 3.2. Generation of High-Quality Pseudo-Labels

Given a pair of images  $(I^s, I^t)$  with their feature maps  $(\mathcal{F}^s, \mathcal{F}^t)$  and the instance masks  $(\mathcal{M}^s, \mathcal{M}^t)$ , the goal is to find matches  $(p_i^s, p_i^t)$  that can be used as pseudo-labels for supervising the training of the adapter  $f_p(\cdot)$ . A naïve strategy would apply the nearest-neighbor (NN) computation (Eq. (1)) of all points on the mask:

$$p_i^t = \text{NN}^{s \rightarrow t}(p_i^s) \quad (2)$$

with  $p_i^s \in \mathcal{M}^s$  and  $p_i^t \in \mathcal{M}^t$ . However, due to viewpoint variations, topological variations, or occlusions, not all points on the mask  $\mathcal{M}^s$  can be mapped to the target image. This challenge is commonly referred to as the *matchability* or *visibility* problem [51]. We show that training an adapter with these naïvely obtained pseudo-labels already improves results (see 5). However, as self-training heavily relies on the quality of the pseudo-labels, we aim to improve their quality by rejecting wrong matches. Pseudo-labels should be removed if no matching is possible (*matchability*) or if points are matched wrongly because of erroneous zero-shot correspondences, *e.g.*, because of *feature ambiguity* of object parts with a similar appearance.

Our work addresses the challenge of removing wrong labels in the following ways: (1) We perform 3D-aware sampling of image pairs along a chain  $(I^s, I^t)$  where image pairs have similar viewpoints and more accurate pseudo-labels are expected. (2) We enforce cyclic consistency for the pseudo-labels along the chain to reduce the number of spurious matches. (3) We reject wrong pseudo-labels by integrating a canonicalized spherical 3D prior.

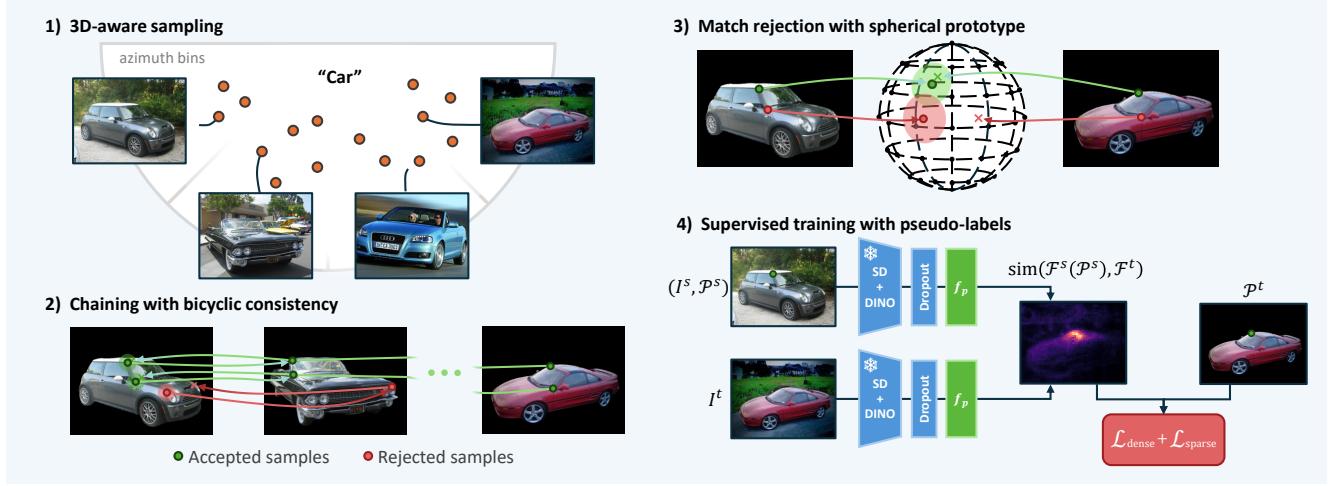


Figure 3. **Method overview.** We use azimuth information to sample image pairs for which higher zero-shot performance can be expected (1). We then chain the pairwise predictions to get correspondences for larger viewpoint changes, where we reject matches that do not fulfill a relaxed cyclic consistency constraint (2). We further filter pseudo-labels by rejecting pairs that can not be mapped to a similar location on a 3D spherical prototype (3). Finally, we use the resulting pseudo-labels to train an adapter  $f_p$  in a supervised manner (4).

$\Delta v$	$[0^\circ, 45^\circ)$	$(0^\circ, 90^\circ)$	$(45^\circ, 135^\circ)$	$(90^\circ, 180^\circ)$	$(135^\circ, 225^\circ)$
PCK@0.1	75.9	68.1	57.2	54.0	52.6

Table 1. **Zero-shot performance (SD+DINO) degrades for larger viewpoint changes.** The SPair-71k dataset has 8 azimuth bins with an equidistant size of  $45^\circ$ . We compute the ratio of correctly corresponded keypoints (PCK@0.1 *per-img*) for increasing viewpoint changes averaged over all categories.

**(1) 3D-aware image pair sampling.** The zero-shot approach, which utilizes features from Stable Diffusion [45] and DINOv2 [40], performs well for objects that do not undergo significant changes in their topology and appearance, and are presented from a similar viewpoint. However, the performance clearly degrades for larger viewpoint changes, as studied by Zhang et al. [63]. We study the sensitivity of the semantic correspondence performance for varying viewpoint differences using the azimuth angle annotations of the object-centric semantic correspondence dataset SPair-71k [36] in Tab. 1. We find that the zero-shot performance decreases for larger viewpoint changes, *i.e.*, if the viewpoint’s azimuth angle  $\phi$  varies by more than  $45^\circ$ .

This observation motivates the generation of pseudo-labels for image pairs with a similar viewpoint. However, more challenging image pairs where the viewpoint heavily deviates are also required to learn a strong point matcher.

To account for this, we propose propagating matches through a  $K$ -tuple  $(I_1, \dots, I_K)$ , where each pair  $(I_k, I_{k+1})$  represents an image pair of the same category with a small but non-zero viewpoint variation (Fig. 3, (1)), such that  $d_{\text{circ}}(\phi_k, \phi_{k+1}) < 90^\circ$  with the spherical distance  $d_{\text{circ}}(\cdot)$ , which is equivalent to one azimuth bin in the SPair-71k

dataset. The matches  $\mathcal{K}^k = \{(p_i^k, p_i^{k+1}) : p_i^k \in \mathcal{P}^k \subseteq \mathcal{M}^k, p_i^{k+1} \in \mathcal{P}^{k+1} \subseteq \mathcal{M}^{k+1}\}$  are propagated via recursive application of the NN operator for the subsequent pairs:

$$\mathcal{P}^{k+1} = \text{NN}^{k \rightarrow k+1}(\mathcal{P}^k) \quad (3)$$

$$= \text{NN}^{k \rightarrow k+1}(\text{NN}^{k-1 \rightarrow k}(\mathcal{P}^{k-1})) = \dots$$

$$\mathcal{P}^1 = \mathcal{M}^1. \quad (4)$$

Through this strategy, each considered image pair is expected to have a higher pseudo-label quality than the naïve strategy, where pseudo-labels are generated for tuples with potentially very different poses  $(I_1, I_k)$ .

**(2) Cyclic consistency of propagated pseudo-labels.** While the aforementioned strategy can improve the quality of matches, it does not necessarily prevent spurious matches that can occur due to feature ambiguities, *e.g.*, because of similar appearance. In order to reduce such spurious matches or matching of incompatible regions, we apply a cyclic consistency constraint for the matched points, as proposed by Aberman et al. [1]. A match  $(p_i^s, p_i^t)$  is cyclic consistent if the following holds:

$$\text{NN}^{s \rightarrow t}(p_i^s) = p_i^t \quad \text{and} \quad \text{NN}^{t \rightarrow s}(p_i^t) = p_i^s, \quad (5)$$

with the nearest neighbor operator NN in feature space, as introduced in Eq. (1).

We observe that this cyclic consistency constraint rejects many keypoint tuples  $(p_i^s, p_i^t)$ , as the zero-shot matching approach has difficulties mapping back to the exact location of  $p_i^s$ . For this reason, we relax the cyclic consistency constraint, where we do not enforce exact consistency but allow

a small deviation from the source point:

$$p_i^t = \text{NN}^{s \rightarrow t}(p_i^s), \quad \hat{p}_i^s = \text{NN}^{t \rightarrow s}(p_i^t), \quad (6)$$

and  $\|\hat{p}_i^s - p_i^s\|_2 < r_{\max}$

with the rejection radius  $r_{\max}$ , see Fig. 3 (2). We iteratively apply this relaxed cyclic consistency constraint to all matches  $\mathcal{K}^k$ , which filters out incompatible pairs and keeps matches capturing concepts that could be propagated through the set of  $K$  images. We choose  $K = 4$  in our experiments since this considers the full range of viewpoint variations of  $180^\circ = 4 \cdot 45^\circ$ , given that the azimuth annotations for SPair-71k are only coarse bins of  $45^\circ$ .

**(3) Rejection of wrong pseudo-labels with a canonicalized spherical 3D prior.** Although our chaining strategy removes spurious matches, wrong matches might still occur, in particular for cases of left-/right-ambiguity and repeated object parts, as, *e.g.*, illustrated in Fig. 3, where simple nearest-neighbor search on foundational features fails.

Therefore, we aim to specifically remove such wrong matches by leveraging a canonicalized object-centric spherical 3D prior that captures objects in an aligned coordinate system. Mariotti et al. [35] propose a *spherical mapper* that maps DINOv2 feature patches  $x_i^{\text{DINO}} = \mathcal{F}^{\text{DINO}}(p_i)$  to points  $\psi_i \in \mathcal{S}^2$  on a canonicalized sphere  $\psi_i = f_s(x_i^{\text{DINO}})$  by using only coarse view point information and instance masks during training. The key advantage of this strategy is that it naturally assigns object parts to regions on a spherical prototype across different instances and categories, which allows rejecting wrong samples in case they are in different areas of the sphere, *e.g.*, when considering visually similar wheels on the left or the right side of a car.

Mariotti et al. [35] demonstrate that this approach significantly improves the performance on SPair-71k when combining their spherical features with the concatenated features of SD and DINO via a weighted average. However, as this approach modulates the original feature similarity, the performance deteriorates for categories that are not well represented by a sphere, *e.g.*, non-rigid categories. In contrast, our approach uses the spherical mapper only for removing potentially wrong pseudo-labels. That filter should have a high true positive rate while still keeping a sufficiently high true negative rate: Supervising with few correct matches is more desirable than supervising with many matches that are partially systematically wrong, as we show in our experiments. For this purpose, we compute the spherical points  $\mathcal{P}^s$  and  $\mathcal{P}^t$  for all matches:

$$\Psi^s = f_s(\mathcal{F}^{\text{DINO}}(\mathcal{P}^s)) \quad \text{and} \quad \Psi^t = f_s(\mathcal{F}^{\text{DINO}}(\mathcal{P}^t)). \quad (7)$$

Then, we reject all matches  $(p_i^s, p_i^t)$  where

$$\text{sim}(\psi_i^s, \psi_i^t) < \theta_{\text{th}}$$

with the threshold  $\theta_{\text{th}} < 0.15 \cdot \pi$  that is selected to allow disambiguating left and right and potentially repeated object parts, such as the wheels of a car. While this thresholding may also remove correct matches for classes that are not well-represented by a rigid object prototype, we observe that the number of matches is still sufficiently large to serve as a dense supervision signal. Most importantly, only removing pseudo-labels does not modulate the original zero-shot matches. This is in contrast to regularizing with spherical mapper features, which deteriorates the localization accuracy of the original foundational features [35].

### 3.3. Supervised Training with Pseudo-Labels

After having generated the pseudo-labels, we train the adapter  $f_p(\cdot)$  using supervised training with the following two losses: First, we train a CLIP [41]-inspired sparse contrastive loss as proposed by Luo et al. [33]:

$$\mathcal{L}_{\text{sparse}} = CL(\mathcal{F}^s(\mathcal{P}^s), \mathcal{F}^t(\mathcal{P}^t)). \quad (8)$$

This loss maximizes feature similarity for corresponding points while minimizing similarity to non-matching points. Second, we also supervise with the dense loss

$$\mathcal{L}_{\text{dense}} = \sum \|\hat{p}_i^t - (p_i^t + \epsilon)\|_2 \quad \text{with} \quad \hat{p}_i^t = \text{WindowSoftArgmax}(\mathcal{F}^s(p_i^s)^\top \mathcal{F}^t), \quad (9)$$

with a window soft-argmax [23, 63] and Gaussian noise samples  $\epsilon$ . The dense loss is particularly important for us since it propagates the gradient also to areas in the feature map without labels, *i.e.*, unmatched areas. This learned combination and refinement of foundational features is a decisive improvement over previously proposed simple concatenation [62, 63] or weighted averaging of feature similarities of a weak geometric regularizer [35].

### 3.4. Scaling to a Larger Dataset

The comparatively small size of the SPair-71k dataset, with only 18 categories and 1,800 images in total, motivates scaling to a larger dataset with more diverse classes and a greater number of images. The spherical mapper [35] is potentially applicable to larger object-centric datasets with 3D annotations, but the simple geometric prior limits the improvement when applied via weighted average, as proposed by Mariotti et al. [35]. The geometric-aware flipping used by the current SOTA methods [9, 63] is not applicable to other datasets since it requires access to the keypoint label definitions, only presented in correspondence datasets.

In contrast, our work relies solely on categories, masks, and 3D object poses as a weak supervision signal, which are available on a larger scale. Therefore, our method can be scaled to larger datasets, leading to generalizable semantic correspondence matching by training on objects that are

		✈	⌚	🐦	🚂	🔧	🚌	🚗	🐺	🏛	🐄	🐎	🏎	🏃	🏆	🐏	🚍	💻	avg	
ASIC	[11]	57.9	25.2	68.1	24.7	35.4	28.4	30.9	54.8	21.6	45.0	47.2	39.9	26.2	48.8	14.5	24.5	49.0	24.6	36.9
DINOv2	[40]	72.7	62.0	85.2	41.3	40.4	52.3	51.5	71.1	36.2	67.1	64.6	67.6	61.0	68.2	30.7	62.0	54.3	24.2	55.6
DIIFT	[49]	63.5	54.5	80.8	34.5	46.2	52.7	48.3	77.7	39.0	76.0	54.9	61.3	53.3	46.0	57.8	57.1	71.1	63.4	57.7
SD + DINOv2	[62]	73.0	64.1	86.4	40.7	52.9	55.0	53.8	78.6	45.5	77.3	64.7	69.7	63.3	69.2	58.4	67.6	66.2	53.5	64.0
DistillDIIFT* (U.S.)	[9]	74.6	60.4	88.7	42.5	53.5	55.0	54.6	80.8	42.7	78.6	72.0	71.4	62.2	70.7	53.1	68.6	65.2	61.6	65.1
SphMap <sup>†</sup>	[35]	75.3	63.8	87.7	48.2	50.9	74.9	71.1	81.7	47.3	81.6	66.9	73.1	65.4	61.8	55.5	70.2	75.0	58.5	67.8
TLR	[63]	78.0	<u>66.4</u>	<u>90.2</u>	44.5	<b>60.1</b>	66.6	60.8	82.7	<b>53.2</b>	82.3	69.5	<u>75.1</u>	<u>66.1</u>	71.7	<u>58.9</u>	<u>71.6</u>	<u>83.8</u>	55.5	69.6
DistillDIIFT* (W.S.)	[9]	<b>78.2</b>	63.8	90.1	45.0	54.6	68.0	63.7	<u>83.2</u>	49.3	<u>82.6</u>	<b>74.5</b>	73.8	63.5	<u>72.0</u>	56.2	<u>71.0</u>	<b>86.2</b>	<u>66.5</u>	70.6
Ours		77.2	<b>69.1</b>	<b>90.8</b>	<b>54.2</b>	<u>57.9</u>	<b>83.7</b>	<b>77.5</b>	<b>86.5</b>	<u>53.1</u>	<b>86.7</b>	<u>73.1</u>	<b>78.5</b>	<b>72.5</b>	<b>74.0</b>	<b>73.5</b>	<b>76.0</b>	77.2	<b>69.5</b>	<b>74.4</b>
SphMap- $\mathcal{S}^2$ (IN3D)		74.3	60.9	82.8	49.4	50.4	76.2	73.3	71.7	47.9	73.0	53.4	68.9	69.4	50.9	34.8	54.4	62.5	57.1	61.9
Ours (IN3D)		75.9	68.7	90.1	55.2	56.1	82.9	76.5	82.7	55.4	83.6	71.2	75.2	71.3	64.3	56.2	69.8	78.6	61.2	71.2
Ours (IN3D $\rightarrow$ SPair)		77.6	<u>70.3</u>	<u>91.0</u>	53.6	<u>58.7</u>	<u>84.8</u>	<u>80.3</u>	86.1	<u>54.1</u>	<u>87.1</u>	<u>74.0</u>	<u>79.2</u>	72.0	<u>75.4</u>	71.9	<u>76.4</u>	<u>77.4</u>	<u>71.9</u>	<u>75.1</u>

Table 2. **Per-category PCK@0.1 scores (per-keypoint) on SPair-71k.** Best and second best are highlighted. When pre-training on ImageNet3D [34], results are improved. <sup>†</sup>SphMap avg is slightly higher than in the original paper where they used macro averaging. We trained using their code. \*DistillDIIFT was evaluated using their trained checkpoints. SphMap and DIY-SC require 3D pose annotations.

not present in the SPair-71k dataset. We demonstrate this by training our whole pipeline on the recently proposed ImageNet-3D [34] dataset, which has around 86k images. We train the spherical mapper similarly to Mariotti et al. [35], where we use masks acquired from SAM [42]. As more accurate 3D pose labels are present in the ImageNet-3D dataset, we do not supervise the viewpoint direction only using the bin of the azimuth angle, as proposed by Mariotti et al. [35] for the SPair-71k dataset. Instead, we extend their formulation by computing the correlation coefficient on the two-sphere  $\mathcal{S}^2$  with the pose annotation:

$$\mathcal{L}_{\mathcal{S}^2} = \sum_{\mathbf{I}, \mathbf{I}'} \|\psi_{\mathbf{I}} \cdot \psi_{\mathbf{I}'} - \mu(f_{\mathbf{S}}(\mathbf{I})) \cdot \mu(f_{\mathbf{S}}(\mathbf{I}'))\|^2, \quad (10)$$

with  $\psi_{\mathbf{I}} \in \mathcal{S}^2$  and  $\psi_{\mathbf{I}'} \in \mathcal{S}^2$  computed from the 3D pose annotations of the objects in the two images. This not only supervises the azimuth viewpoint information (yaw) but also the pitch of an object and, therefore, improves canonicalization. We provide more details in the supplementary.

This strategy enables us to scale our proposed method to more classes and images from a significantly larger dataset, ultimately leading to improved generalization.

## 4. Experiments

**Details on pseudo-label generation.** Using our sampling strategy, we generate 30k pseudo-labeled image pairs per category from SPair-71k. We use a maximum of 50 randomly sampled keypoints for supervised training for each image pair. Following Zhang et al. [62], we resize the input images to  $960^2$  and  $840^2$  to extract the SD and DINOv2 features, resulting in feature maps of resolution  $60 \times 60$ . We choose the rejection radius  $r_{\max}$  for the relaxed bicyclic consistency constraint such that a deviation of one feature patch is allowed.

**Details on adapter training.** Similar to Zhang et al. [63], we use four bottleneck layers with 5M parameters for the adapter. The model is trained using the AdamW optimizer

Models	SPair-71k			AP-10k (PCK@0.1)			
	0.1	0.05	0.01	I.S.	C.S.	C.F.	
SD + DINOv2	[62]	59.9	44.7	7.9	62.9	59.3	48.3
DistillDIIFT* (U.S.)	[9]	60.8	45.4	8.0	-	-	-
SphMap <sup>†</sup>	[35]	64.4	48.2	8.4	65.4	63.1	51.0
TLR	[63]	65.4	49.1	9.9	68.7	64.6	52.7
DistillDIIFT* (W.S.)	[9]	65.3	49.8	8.9	-	-	-
Ours		<b>71.6</b>	<b>53.8</b>	<b>10.1</b>	<u>70.6</u>	<u>69.1</u>	<u>57.8</u>
Ours (DINOv2)		<u>70.6</u>	<u>51.1</u>	9.0	<b>71.2</b>	<b>69.8</b>	<b>58.3</b>
TLR (sup)	[63]	82.9	72.6	21.6	70.1	68.3	58.4

Table 3. **Results for different PCK levels (per-image) on the SPair-71k and the AP-10k dataset.** Reproduced supervised results for the model [63] solely trained on SPair-71k. Results for AP-10K are intra-species (I.S.), cross-species (C.S.), and cross-family (C.F.), as introduced in [63]. <sup>†</sup> and \* as in Tab. 2. DINOv2: Only application and refinement of DINOv2 features.

with a weight decay rate 0.001, a learning rate of  $5 \cdot 10^{-3}$ , and the one-cycle scheduler for 200k steps.

**Datasets.** Recent works [35, 49, 62, 63] find SPair-71k [36] to be the most challenging and informative semantic matching benchmark due to its larger size and more challenging matches, containing images from 18 different categories with 600 to 900 image pairs per category. We also evaluate the methods on the recently proposed AP-10k [61] semantic correspondence dataset [63] without training on it.

**Metrics.** We follow the standard settings [11, 17, 35, 36, 59, 63] and evaluate the semantic correspondence performance via the Percentage of Correct Keypoints (PCK). It is defined by computing the ratio of correctly predicted matched keypoints that lie within a radius of  $R = \alpha \cdot \max(h, w)$  around the correct ground truth match, where we report the *bbox* variant with  $h$  and  $w$  referring to the height and width of the bounding box of the considered object, respectively. Since various prior works report it differently, we report both the PCK when averaged over keypoints (*per kpt*) and when averaged over images (*per img*).

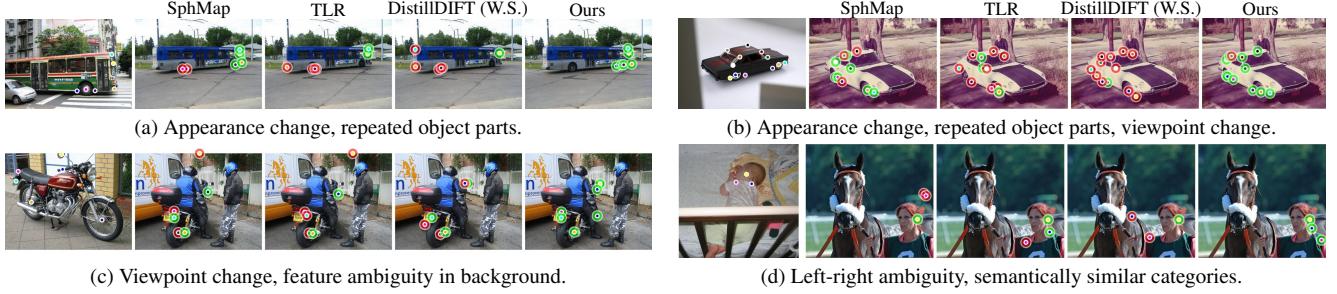


Figure 4. **Qualitative results for challenging examples.** We present four challenging examples of the SPair-71K test dataset where current SOTA models mostly fail. Correct matches that lie within the  $\text{PCK}_{\text{bbox}} @ 0.1$  radius of the ground-truth label are color-coded in green, while incorrectly matched keypoints are depicted in red.

**Prior work.** We compare our method with recent zero-shot and weakly-supervised strategies. *DIIFT* [49], *DINOv2* [40], and *SD + DINOv2* [62] extract features from foundation models and perform nearest-neighbor matching in feature space. *DistillDIIFT (U.S.)* [9] is further fine-tuned for correspondence detection on 3D data. *ASIC* [11] uses DINO features to establish coarse correspondences and refines them at test-time. Spherical mapper (*SphMap*) [35] uses class labels, object masks, and 3D pose information as weak supervision signals during training, similar to us. We note, however, that we do not use any of this information at test time. Telling Left from Right (*TLR*) [63] and *DistillDIIFT (W.S.)* [9] make use of the keypoint label definitions to flip images at test time. We compare to zero-shot approaches and the spherical mapper [35] trained on ImageNet-3D with our extended 2-sphere configuration.

#### 4.1. Experimental Results

**Quantitative results on SPair-71k.** The category-specific results in Tab. 2 show that our model achieves strong gains across nearly all categories, setting a new SOTA on SPair-71k and demonstrating the effectiveness of our pseudo-labeling strategy. Improvements are especially large for categories that are symmetric or contain repeated object parts, *e.g.*, +15.7% for bus or +14.0% for car. While Mariotti et al. [35] already improved performance for such cases, our approach outperforms theirs by effectively combining SD and DINOv2 features through pseudo-label supervision. Unlike their method, we retain original features and simply reject matches that strongly deviate from a spherical prototype.

All in all, we clearly outperform all previous methods on all PCK-levels (Tab. 2 and Tab. 3), including methods that use the same weak supervision as we do [35] and methods that include dataset-specific information about the definition of the keypoint labels [9, 63]. Surprisingly, training an adapter only on top of DINOv2 features with the same pseudo-labels results in competitive semantic correspondence performance compared to the adapter on top

Models	SPair-71k			AP-10k (PCK@0.1)		
	0.1	0.05	0.01	I.S.	C.S.	C.F.
Ours (SPair)	71.6	53.8	10.1	70.6	69.1	57.8
SphMap- $\mathcal{S}^2$ (IN3D)	58.5	44.0	7.7	56.6	53.2	38.9
Ours (IN3D)	68.0	51.1	9.8	67.8	65.8	53.3
Ours (IN3D $\rightarrow$ SPair)	<b>72.2</b>	<b>54.6</b>	<b>10.6</b>	<b>71.1</b>	<b>69.4</b>	<b>58.1</b>

Table 4. **Effect of pre-training on ImageNet-3D (PCK per img).** SphMap- $\mathcal{S}^2$  represents the weighted application of our 2-sphere-based spherical mapper, evaluated following Mariotti et al. [35].

of SD+DINOv2 features, even outperforming this strategy for the animal AP-10K dataset. This demonstrates that DINOv2 features achieve competitive semantic correspondence performance if refined effectively.

**Evaluation on AP-10k.** Additionally, we also evaluate on the AP-10k dataset and report the results in Tab. 3. We find that, despite not being trained on this dataset and only having a few animal categories in the SPair-71k training dataset, we outperform the current SOTA, the geometry-aware approach by Zhang et al. [63], showcasing the generalization capabilities of our method. We additionally compare to a fully-supervised model [63] trained on SPair-71k. We observe a severe performance drop of this model (−17.4% PCK@0.1 for I.S.) compared to fully-supervised training on AP-10k [63, Tab.3]. On the other hand, our approach does not suffer such a performance drop and, thus, achieves comparable results to a model trained with manual annotations. We hypothesize that this is due to the denser and more diverse pseudo-labels compared to the fixed set of manually annotated keypoints that are being overfitted in supervised training.

**Scaling to a larger dataset.** Our weakly-supervised training strategy allows scaling the training to larger object-centric datasets that do not include keypoint annotations. While pre-training improves the model performance on SPair-71k and AP-10k (Tab. 3 and Tab. 4), we surprisingly even outperform the previous SOTA on SPair-71k despite not having trained on the SPair-71k dataset, showcasing strong generalization capabilities. This is in contrast to the

naïve weighted application of the ImageNet-3D spherical mapper, where the results significantly deteriorate, potentially explained by the fact that this model does not generalize well to the unseen categories in SPair-71k. We also study the impact of automatically acquired 3D pose labels in the supplementary.

**Qualitative results.** We present qualitative results in Fig. 4 comparing our model to current SOTA models on four challenging examples of the SPair-71k test dataset. While the spherical mapper effectively addresses cases of symmetries and repeated object parts for rigid objects, it comes with the inherent challenge of finding the correct weighting factor for interpolating between foundational and processed spherical feature similarities, where the optimal weight depends on the considered category and sample. For example, while the weighting is too low in (a), the weighting is too large in (d) since the ill-defined background features of SphMap result in a match with an object in the background. Additionally, its effectiveness is limited to categories that are well represented by a spherical prototype, *e.g.*, a potential reason for the motorbike failure case in (c). TLR and DistillDIFT both suffer from the fact that simply flipping an object does not resolve all ambiguities, such as for categories with repeated object parts (a,b), rotated objects (b), or cases with left-/right ambiguities (d). DistillDIFT is less prone to incorrect matches caused by background clutter (c), whereas TLR and SphMap suffer from the fact that they rely on the potentially noisy nearest neighbor of SD and DINOV2 features. However, DistillDIFT still fails to disambiguate the human and the horse in (d), and it suffers from left-right ambiguity if the flipping is not effective (d). Furthermore, as visible in (b), DistillDIFT’s performance deteriorates for pairs with heavy appearance change, which may be attributed to the multi-view fine-tuning on 3D instances, which does not generalize well to cross-instance correspondences. In contrast, our method identifies high-quality semantic correspondences for these challenging examples, as it successfully resolves feature ambiguities and does not rely on the weighting of two similarity terms.

## 4.2. Ablations

We thoroughly ablate all components of our method and present the main results in Tab. 5. First, supervision with naïvely generated pseudo-labels already outperforms the zero-shot approach. This is due to the fact that learning the light-weight adapter naturally stabilizes feature maps for semantic correspondence estimation by merging the complementary information captured by SD and DINOV2 features. Constraining pseudo-labels to the mask of an object further enhances the performance. We also observe that the relaxed cyclic consistency constraint outperforms both unfiltered pseudo-labels and exact cyclic consistency filtering. The performance further increases when employing our chaining

pseudo	cyc. cons.	relaxed c.c.	chaining	sph. rej.	PCK@0.1
✓					65.0
✓	✓				67.2
✓		✓			66.9
✓		✓	✓		68.4
✓		✓	✓		70.0
✓				✓	72.9
✓		✓	✓	✓	<b>74.4</b>

Table 5. **Ablations on SPair-71k.** All introduced components bring a significant improvement. The baseline is evaluated using the SD+DINO zero-shot approach with window soft argmax.

ing strategy that improves the pseudo-label quality because (1) NN matching is only performed for easier image pairs, *i.e.*, a similar object pose, and (2) the bicyclic consistency constraint reduces spurious matches. Adding the sphere-based rejection strategy substantially improves all model instantiations by reducing wrong matches for objects with symmetries and repeated object parts that make up a significant ratio of the SPair-71k dataset. This rejection also clearly enhances the naive pseudo-labeling strategy, combining all modules to achieve the best performance.

**Failure Cases.** While our method outperforms previous methods, it still fails for some challenging cases of SPair-71K, such as images containing objects that are upside-down. We show visual examples and discuss failure cases of our method in the supplementary.

## 5. Conclusion

We presented DIY-SC, a method for finding semantic correspondences with weak supervision. We demonstrated the effectiveness of training with pseudo-labels for semantic correspondence: By generating labels using a zero-shot approach with foundational features and carefully designing methods for rejecting wrong matches, we provide a strong supervision signal that can be used to train an adapter that refines foundational features. Our experiments showed improvements of **+4.5%** and **+6.8%** over the previous SOTA for PCK@0.1 on SPair-71k, measured per-keypoint and per-image, respectively. Our ablation studies have shown that the proposed chaining and filtering steps, as a means of enhancing the quality of pseudo-labels, improve the corresponding performance significantly.

**Limitations and future work.** Although our weak supervision strategy scales to larger datasets, reducing the supervision requirements is of importance to further scale the approach. See the supplementary for a preliminary feasibility analysis. Furthermore, while the match rejection with the spherical mapper shows significant improvements, it is less effective for objects that are not represented well by a sphere, which motivates extending this strategy. We hope that the proposed improved correspondence features can enhance tasks that rely on semantic features.

## Acknowledgments

Adam Kortylewski acknowledges support via his Emmy Noether Research Group funded by the German Research Foundation (DFG) under Grant No. 468670075. Thomas Wimmer is supported through the Max Planck ETH Center for Learning Systems.

## References

- [1] Kfir Aberman, Jing Liao, Mingyi Shi, Dani Lischinski, Baoquan Chen, and Daniel Cohen-Or. Neural best-buddies: Sparse cross-domain correspondence. *TOG*, 2018. [2](#), [4](#)
- [2] Shir Amir, Yossi Gandalman, Shai Bagon, and Tali Dekel. Deep ViT features as dense visual descriptors. *ECCVW What is Motion For?*, 2022. [2](#)
- [3] YM Asano, C Rupprecht, and A Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020. [3](#)
- [4] Souhaib Attaiki and Maks Ovsjanikov. NCP: Neural correspondence prior for effective unsupervised shape matching. In *NeurIPS*, 2022. [3](#)
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Herve Jegou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. [2](#), [3](#)
- [6] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Show, match and segment: Joint weakly supervised learning of semantic matching and object co-segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [2](#)
- [7] Xinle Cheng, Congyue Deng, Adam W Harley, Yixin Zhu, and Leonidas Guibas. Zero-shot image feature consensus with deep functional maps. In *ECCV*, 2024. [2](#)
- [8] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. Tap-vid: A benchmark for tracking any point in a video. In *NeurIPS*, 2022. [1](#)
- [9] Frank Fundel, Johannes Schusterbauer, Vincent Tao Hu, and Björn Ommer. Distillation of diffusion features for semantic correspondence. In *WACV*, 2025. [2](#), [5](#), [6](#), [7](#), [1](#)
- [10] Shenyuan Gao, Chunluan Zhou, Chao Ma, Xinggang Wang, and Junsong Yuan. Aiatrack: Attention in attention for transformer visual tracking. In *ECCV*, 2022. [1](#)
- [11] Kamal Gupta, Varun Jampani, Carlos Esteves, Abhinav Shrivastava, Ameesh Makadia, Noah Snavely, and Abhishek Kar. ASIC: Aligning sparse in-the-wild image collections. In *ICCV*, 2023. [2](#), [6](#), [7](#)
- [12] Yoav HaCohen, Eli Shechtman, Dan B. Goldman, and Dani Lischinski. Non-rigid dense correspondence with applications for image enhancement. In *SIGGRAPH*, 2011. [1](#)
- [13] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T. Freeman. Unsupervised semantic segmentation by distilling feature correspondences. In *ICLR*, 2022. [3](#)
- [14] Kai Han, Rafael S. Rezende, Bumsub Ham, Kwan-Yee K. Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. SC-Net: Learning semantic correspondence. In *ICCV*, 2017. [2](#)
- [15] Eric Hedlin, Gopal Sharma, Shweta Mahajan, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Unsupervised semantic correspondence using stable diffusion. In *NeurIPS*, 2023. [2](#)
- [16] Lin Huang, Tomas Hodan, Lingni Ma, Linguang Zhang, Luan Tran, Christopher Twigg, Po-Chen Wu, Junsong Yuan, Cem Keskin, and Robert Wang. Neural correspondence field for object pose estimation. In *ECCV*, 2022. [1](#)
- [17] Shuaiyi Huang, Luyu Yang, Bo He, Songyang Zhang, Xuming He, and Abhinav Shrivastava. Learning semantic correspondence with sparse annotations. In *ECCV*, 2022. [6](#)
- [18] Yiwen Huang, Yixuan Sun, Chenghang Lai, Qing Xu, Xiaomei Wang, Xuli Shen, and Weifeng Ge. Weakly supervised learning of semantic correspondence through cascaded online correspondence refinement. In *ICCV*, 2023. [2](#), [3](#)
- [19] Patrick Kage, Jay C Rothenberger, Pavlos Andreadis, and Dimitrios I Diochnos. A review of pseudo-labeling for computer vision. *arXiv preprint arXiv:2408.07221*, 2024. [3](#)
- [20] Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-tracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*, 2024. [1](#), [3](#)
- [21] Jiwon Kim, Kwangrok Ryoo, Junyoung Seo, Gyuseong Lee, Daehwan Kim, Hansang Cho, and Seungryong Kim. Semi-supervised learning of semantic correspondence with pseudo-labels. In *CVPR*, 2022. [2](#), [3](#)
- [22] Seungryong Kim, Dongbo Min, Stephen Lin, and Kwanghoon Sohn. DCTM: Discrete-continuous transformation matching for semantic flow. In *ICCV*, 2017. [2](#)
- [23] Seungryong Kim, Dongbo Min, Somi Jeong, Sunok Kim, Sangryul Jeon, and Kwanghoon Sohn. Semantic attribute matching networks. In *CVPR*, 2019. [1](#), [2](#), [5](#)
- [24] Filippos Kokkinos and Iasonas Kokkinos. To the point: Correspondence-driven monocular 3d category reconstruction. In *NeurIPS*, 2021. [1](#)
- [25] Zihang Lai, Senthil Purushwalkam, and Abhinav Gupta. The functional correspondence problem. In *ICCV*, 2021. [1](#)
- [26] Shiyi Lan, Zhiding Yu, Christopher Choy, Subhashree Radhakrishnan, Guilin Liu, Yuke Zhu, Larry S. Davis, and Anima Anandkumar. DiscoBox: Weakly supervised instance segmentation and semantic correspondence from box supervision. In *ICCV*, 2021. [2](#)
- [27] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *learning*, 2013. [3](#)
- [28] Xin Li, Deng-Ping Fan, Fan Yang, Ao Luo, Hong Cheng, and Zicheng Liu. Probabilistic model distillation for semantic correspondence. In *CVPR*, 2021. [2](#), [3](#)
- [29] Xinghui Li, Jingyi Lu, Kai Han, and Victor Adrian Prisacariu. Sd4match: Learning to prompt stable diffusion model for semantic matching. In *CVPR*, 2024. [2](#)
- [30] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *PAMI*, 2010. [2](#)
- [31] Haolin Liu, Xiaohang Zhan, Zizheng Yan, Zhongjin Luo, Yuxin Wen, and Xiaoguang Han. Stable-SCore: A stable

registration-based framework for 3d shape correspondence. *arXiv preprint arXiv:2503.21766*, 2025. 1

[32] David G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004. 2

[33] Grace Luo, Lisa Dunlap, Dong Huk Park, Aleksander Holynski, and Trevor Darrell. Diffusion hyperfeatures: Searching through time and space for semantic correspondence. In *NeurIPS*, 2023. 1, 2, 5

[34] Wufei Ma, Guofeng Zhang, Qihao Liu, Guanning Zeng, Adam Kortylewski, Yaoyao Liu, and Alan Yuille. Imagenet3d: Towards general-purpose object-level 3d understanding. In *NeurIPS*, 2024. 2, 6, 1

[35] Octave Mariotti, Oisin Mac Aodha, and Hakan Bilen. Improving semantic correspondence with viewpoint-guided spherical maps. In *CVPR*, 2024. 1, 2, 3, 5, 6, 7

[36] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Spair-71k: A large-scale benchmark for semantic correspondence. *arXiv preprint arXiv:1908.10543*, 2019. 2, 4, 6

[37] Raul Mur-Artal, J. M. M. Montiel, and Juan D. Tardos. ORB-SLAM: A versatile and accurate monocular SLAM system. *T-RO*, 2015. 1

[38] David Novotny, Diane Larlus, and Andrea Vedaldi. AnchorNet: A weakly supervised network to learn geometry-sensitive features for semantic matching. In *CVPR*, 2017. 2

[39] Dolev Ofri-Amar, Michal Geyer, Yoni Kasten, and Tali Dekel. Neural congealing: Aligning images to a joint semantic atlas. In *CVPR*, 2023. 1, 2

[40] Maxime Oquab, Timothée Darcret, Théo Moutakanni, Huy V Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *TMLR*, 2023. 2, 3, 4, 6, 7

[41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 5

[42] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädele, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 6

[43] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *CVPR*, 2017. 2

[44] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. End-to-end weakly-supervised semantic alignment. In *CVPR*, 2018. 2

[45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3, 4

[46] Leonhard Sommer, Olaf Dünkel, Christian Theobalt, and Adam Kortylewski. Common3d: Self-supervised learning of 3d morphable models for common objects in neural feature space. In *CVPR*, 2025. 1, 2

[47] Nick Stracke, Stefan Andreas Baumann, Kolja Bauer, Frank Fundel, and Björn Ommer. CleanDIFT: Diffusion features without noise. *arXiv preprint arXiv:2412.03439*, 2024. 2

[48] Saksham Suri, Matthew Walmer, Kamal Gupta, and Abhinav Shrivastava. Lift: A surprisingly simple lightweight feature transform for dense vit descriptors. In *ECCV*, 2024. 2

[49] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent correspondence from image diffusion. In *NeurIPS*, 2023. 1, 2, 6, 7

[50] Anh Thai, Weiyao Wang, Hao Tang, Stefan Stojanov, James M Rehg, and Matt Feiszli. 3×2: 3d object part segmentation by 2d semantic correspondences. In *ECCV*, 2024. 1

[51] Prune Truong, Martin Danelljan, Fisher Yu, and Luc Van Gool. Warp consistency for unsupervised learning of dense correspondences. In *ICCV*, 2021. 2, 3

[52] Prune Truong, Martin Danelljan, Fisher Yu, and Luc Van Gool. Probabilistic warp consistency for weakly-supervised semantic correspondences. In *CVPR*, 2022. 2

[53] Jiangshan Wang, Yue Ma, Jiayi Guo, Yicheng Xiao, Gao Huang, and Xiu Li. Cove: Unleashing the diffusion feature correspondence for consistent video editing. In *NeurIPS*, 2025. 1

[54] Zehan Wang, Ziang Zhang, Tianyu Pang, Chao Du, Hengshuang Zhao, and Zhou Zhao. Orient anything: Learning robust object orientation estimation from rendering 3d models. *arXiv:2412.18605*, 2024. 1

[55] Thomas Wimmer, Peter Wonka, and Maks Ovsjanikov. Back to 3d: Few-shot 3d keypoint detection with back-projected 2d features. In *CVPR*, 2024. 2

[56] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*, 2024. 1

[57] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. Self-training with noisy student improves ImageNet classification. In *CVPR*, 2020. 3

[58] Yan Xu, Kwan-Yee Lin, Guofeng Zhang, Xiaogang Wang, and Hongsheng Li. RNNPose: Recurrent 6-DoF object pose refinement with robust correspondence field estimation and pose optimization. In *CVPR*, 2022. 1

[59] Fei Xue, Sven Elflein, Laura Leal-Taixé, and Qunjie Zhou. MATCHA: Towards matching anything. *arXiv preprint arXiv:2501.14945*, 2025. 1, 2, 6

[60] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. LIFT: Learned invariant feature transform. In *ECCV*, 2016. 2

[61] Hang Yu, Yufei Xu, Jing Zhang, Wei Zhao, Ziyu Guan, and Dacheng Tao. AP-10k: A benchmark for animal pose estimation in the wild. In *NeurIPS*, 2021. 6

[62] Junyi Zhang, Charles Herrmann, Junhwa Hur, Luisa Polanía Cabrera, Varun Jampani, Deqing Sun, and Ming-Hsuan Yang. A tale of two features: Stable diffusion complements dino for zero-shot semantic correspondence. In *NeurIPS*, 2023. 1, 2, 3, 5, 6, 7

[63] Junyi Zhang, Charles Herrmann, Junhwa Hur, Eric Chen, Varun Jampani, Deqing Sun, and Ming-Hsuan Yang. Telling

left from right: Identifying geometry-aware semantic correspondence. In *CVPR*, 2024. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)

- [64] Kaifeng Zhang, Yang Fu, Shubhankar Borse, Hong Cai, Fatih Porikli, and Xiaolong Wang. Self-supervised geometric correspondence for category-level 6d object pose estimation in the wild. In *ICLR*, 2023. [2](#)
- [65] Kaiyan Zhang, Xinghui Li, Jingyi Lu, and Kai Han. Semantic correspondence: Unified benchmarking and a strong baseline. *arXiv preprint arXiv:2505.18060*, 2025. [2](#)
- [66] Tinghui Zhou, Philipp Krahenbuhl, Mathieu Aubry, Qixing Huang, and Alexei A. Efros. Learning dense correspondence via 3d-guided cycle consistency. In *CVPR*, 2016. [2](#)

# Do It Yourself: Learning Semantic Correspondence from Pseudo-Labels

## Supplementary Material

### A. Details on ImageNet-3D Trained Model

In this subsection, we present more details about the model that was trained on ImageNet-3D, including more details about the improved spherical mapper.

**Pose Conversion in ImageNet-3D.** We reformulate the loss objective for taking into account viewpoint information as presented in Eq. (10). We acquire the needed labels in the following way: Given the rotation matrix  $R$  presenting the 3D pose in the ImageNet-3D dataset [34], we compute the corresponding coordinate on the 2-sphere  $\psi = [\theta, \phi] \in S^2$  as follows:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = R \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix},$$

$$\theta = \arccos(z), \quad \phi = \text{atan2}(y, x)$$

**Experimental Details** For training the spherical mapper, we use the same hyperparameter as Mariotti et al. [35] and we train for 200 epochs on the ImageNet-3D dataset. The generation of the pseudo-labels and training of the adapter follow the same hyperparameters as our presented model for SPair-71k. We train on this larger dataset for 400k steps.

### B. Discussion of Failure Cases

While we show significant improvements for most object categories, our method does not improve results for all classes compared to the SOTA. Our approach performs worse when objects are vertically flipped, *e.g.*, for the airplane category or bicycles, as presented in the challenging example in Tab. 6a. In such cases, a limiting factor is that the polar angle is not available for 3D-aware sampling and training of the spherical mapper. Our method also fails for heavy perspective changes Tab. 6b.



(a) Our model can fail for objects that are up-side down.  
(b) Matching objects under heavy scale changes is challenging.

Table 6. Examples of failure cases.

### C. Training with Less Supervision

To explore the scalability of our proposed strategy to larger datasets without 3D pose annotations or masks, we study

how the performance deteriorates for SPair-71k when not having access to the viewpoint annotation. For this, we extract pose information using Orient-Anything [54] and extract object masks using `rembg` for unsupervised foreground extraction. Using SAM masks, which are of slightly higher quality but not completely unsupervised (although feasible through, *e.g.*, using GroundingSAM), and the automatically extracted poses, the PCK@0.1 *per-img* of our method drops to 69.6% compared to when using GT pose annotations (71.6%). Using `rembg` masks, it drops further to 68.0%, which is expected but still around 2.7p better than the previous best weakly supervised method [9]. While [47] uses dataset-specific information about the keypoint label convention, our approach only requires class labels. We report the results in Tab. 7.

3D pose label	mask label	PCK@0.1
Ground Truth	SAM	71.6
Orient-Anything	SAM	69.6
Orient-Anything	<code>rembg</code>	68.0

Table 7. PCK@0.1 *per-image* on SPair-71k without ground-truth viewpoint annotations.

### D. Pre-Training with Pseudo-Labels

We explore whether pre-training with our pseudo-labels also improves the supervised performance. For this purpose, we fine-tune the adapter with ground truth labels of the SPair-71k dataset, which improves the supervised performance from 82.9% [63] to 83.5% (PCK@0.1 *per-img*).

### E. Pseudo-Label Generation without SD

When training a refiner of DINOv2 features with pseudo-labels that are acquired only from DINOv2 features, *i.e.*, not from SD+DINOv2 as in the main paper, the performance drops to 67.17% (*per-img*) and to 70.29% (*per-kpt*), which is still on par with recent SOTA models.

Channels	PCK@0.1	PCK@0.05	PCK@0.01
128	74.08	56.28	11.26
384	74.39	56.87	11.53
768	74.43	56.76	11.22
1536	74.63	57.13	11.60

Table 8. PCK metrics for varying numbers of feature channels.

## F. Ablation of Number of Feature Channels

Learning a refining module allows reducing the number of channels of the features used for nearest neighbor computation. The performance does not heavily drop with a lower number of channels (Tab. 8), which might be a valuable option for memory-constrained applications.

## G. More Qualitative Results

We show detected correspondences for uncurated image pairs from the SPair-71k test dataset for DistillDIFT, TLR, SphMap, and our method in Fig. 5 and Fig. 6.

DistillDIFT

TLR

SphMap

Ours

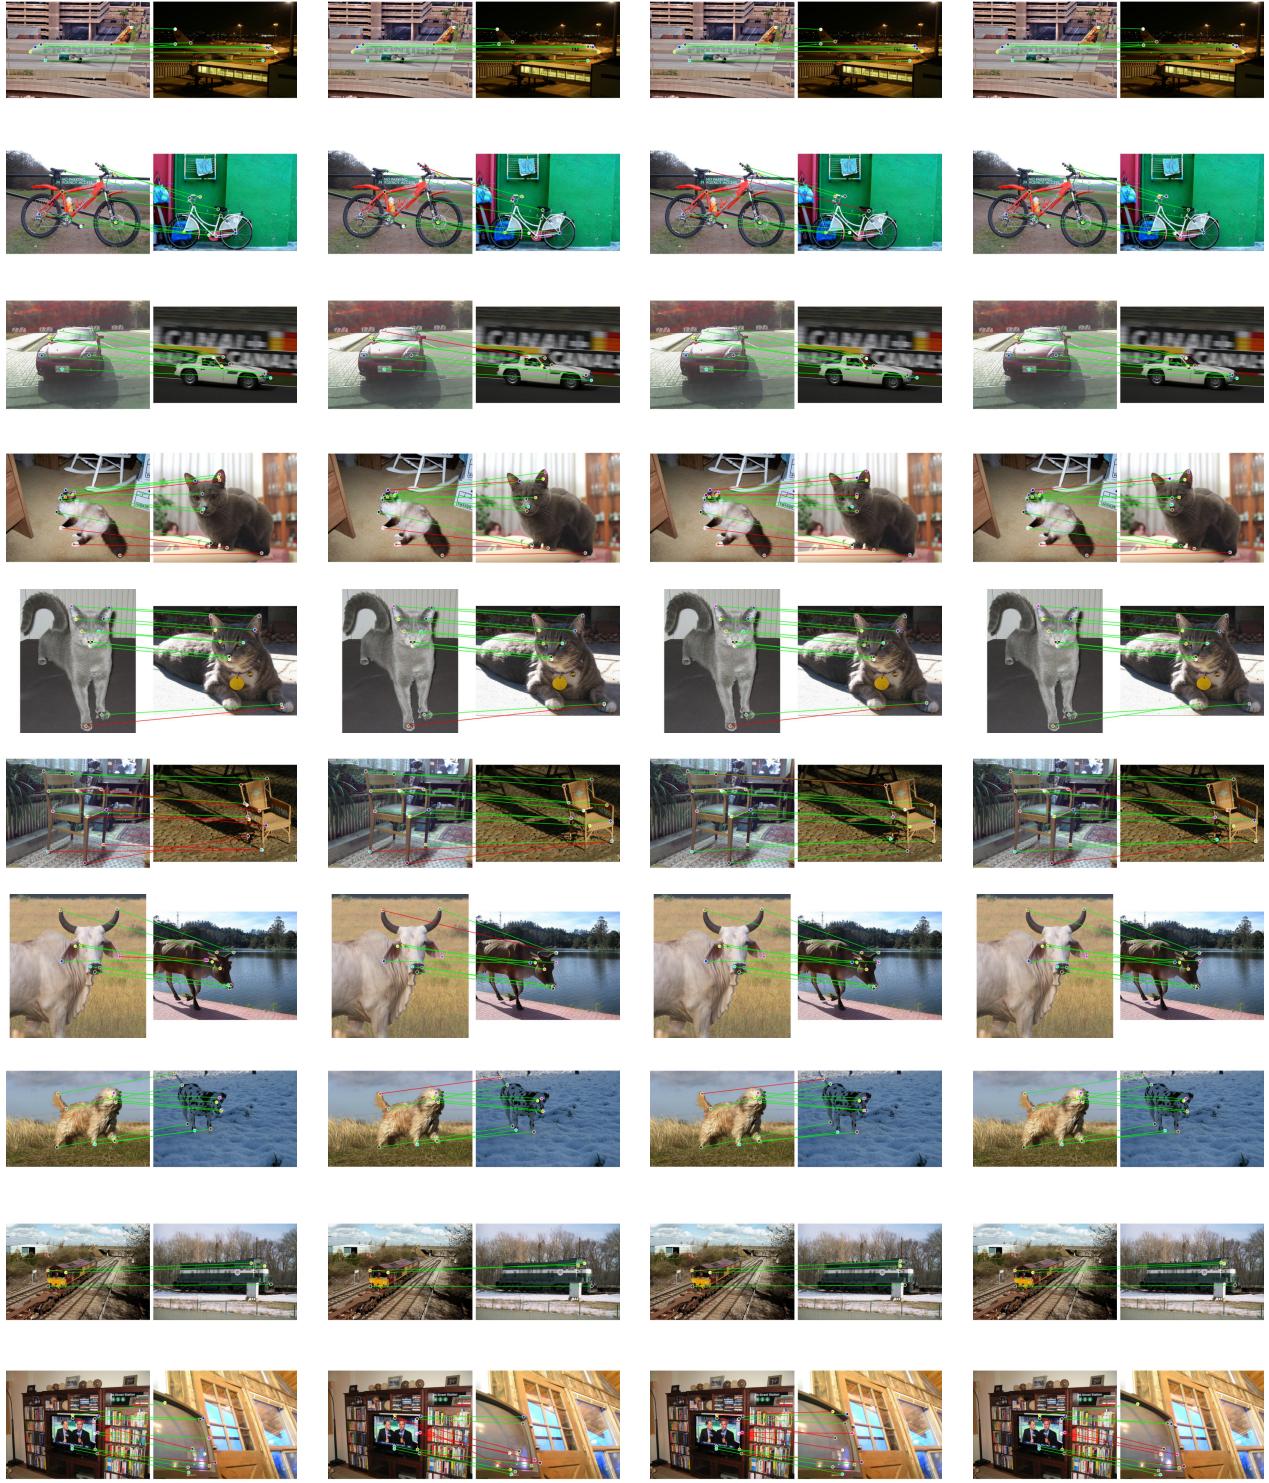


Figure 5. Uncurated image pairs of SPair-71K dataset of three SOTA models and ours.



Figure 6. Uncurated image pairs of SPair-71K dataset of three SOTA models and ours.