

# Improving Semantic Correspondence with Viewpoint-Guided Spherical Maps

Octave Mariotti      Oisín Mac Aodha      Hakan Bilen

University of Edinburgh

[groups.inf.ed.ac.uk/vico/research/sphericalmaps](https://groups.inf.ed.ac.uk/vico/research/sphericalmaps)

## Abstract

Recent self-supervised models produce visual features that are not only effective at encoding image-level, but also pixel-level, semantics. They have been reported to obtain impressive results for dense visual semantic correspondence estimation, even outperforming fully-supervised methods. Nevertheless, these models still fail in the presence of challenging image characteristics such as symmetries and repeated parts. To address these limitations, we propose a new semantic correspondence estimation method that supplements state-of-the-art self-supervised features with 3D understanding via a weak geometric spherical prior. Compared to more involved 3D pipelines, our model provides a simple and effective way of injecting informative geometric priors into the learned representation while requiring only weak viewpoint information. We also propose a new evaluation metric that better accounts for repeated part and symmetry-induced mistakes. We show that our method succeeds in distinguishing between symmetric views and repeated parts across many object categories in the challenging SPair-71k dataset and also in generalizing to previously unseen classes in the AwA dataset.

## 1. Introduction

Semantic correspondence (SC) estimation aims to find local regions that correspond to the same semantic entities across a collection of images, where each image contains a different instance of the same object category [30]. SC has been studied in both supervised [9, 26, 59] and unsupervised [2, 50] settings. Thanks to the recent progress in self-supervised learning (SSL) [6, 40, 60], recent SSL-based approaches have been shown to obtain strong performance on multiple benchmarks [1, 2, 54]. However, they often act as simple part detectors, without having any 3D understanding of object extent, or relative location of object parts. As these SSL methods are typically trained using only 2D-based objectives, they are not able to learn 3D-aware representations, and often converge to similar features for object parts that share appearance but not fine-grained semantics.

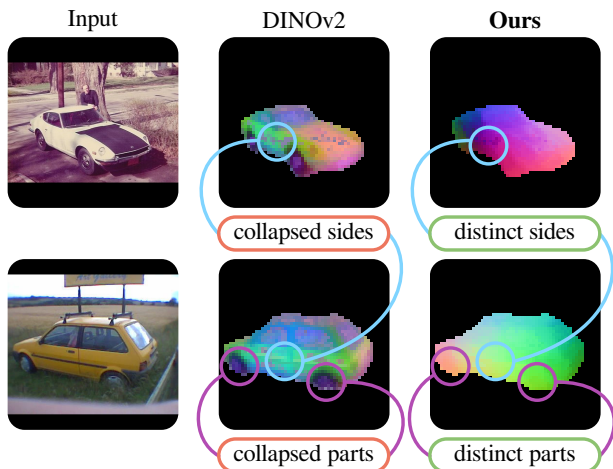


Figure 1. Features from self-supervised methods such as DINOv2 [40] have been used to discover parts and regions of objects. However, these features fail to correctly distinguish (i) object symmetries, e.g. the left and right side of the car have the same features and (ii) individual parts, e.g. the wheels are represented by the same features irrespective of their location on the car. Our approach through use of a weak geometric spherical prior addresses these issues. Note, we use the 3D PCA projection of the features for DINOv2 and our learned spherical mapping for our method.

The two main failure modes of current methods are depicted in Fig. 1. First, they are often confused by the symmetries that objects can exhibit (e.g. left vs. right side of a car), producing reflected versions of the feature maps when presented with the two different sides of a symmetric object. Second, they predict similar features for similar looking parts (e.g. the front left vs. back left wheel of a car), and fail to discriminate between part instances. Such symmetries and repeated structures are ubiquitous in both human-made and naturally occurring object categories.

We posit that these shortcomings can be addressed by enforcing an explicit 3D structure into correspondence learning, allowing the model to disambiguate multiple repeated parts, and accounting for the fact that some object parts might not be visible in all images. Indeed, scene-based correspondence pipelines use geometric verification in order to guarantee robustness in their matches [44], while other works have shown that detailed 3D meshes can be used to

discover category-based correspondences through a render-and-compare pipeline [29, 62]. However, these approaches typically require 3D mesh supervision which is difficult and costly to obtain at scale.

Motivated by the limitations, and the promise of the recent SSL representations [40], we propose a novel method that incorporates the missing 3D geometry into object representations. We make use of a weak 3D prior by representing object features on the surface of a sphere which encodes both geometry and appearance information. Our method consistently maps features corresponding to the same semantic parts (*e.g.* left headlight of a car) in images of different object instances to the same spherical coordinates, hence enabling a 3D aware and accurate correspondence. However, learning such a mapping without the ground-truth 3D supervision is nontrivial because of repeated parts and object symmetries. To this end, we propose to use much weaker supervision via coarse camera viewpoint annotations to help separate these confusing features during training. We couple this with geometric constraints to prevent collapsed representations for repeated parts. As a result, our method learns to separate object parts based on their position and spatial configuration in an image.

Furthermore, we point out that the standard evaluation metric for SC, Percentage of Correct Keypoints (PCK), is not able to properly evaluate symmetry-related failures as it only considers keypoints that are *simultaneously visible* in both the source and target image. For instance, a system predicting high similarity between a front left wheel in the source and front right wheel in the target is unlikely to be penalized for it, as the target image is unlikely to contain both a front left *and* front right wheel due to self-occlusions. In such cases PCK will simply not evaluate what the nearest neighbor of the source wheel is. Rather than simple accuracy-based metric, we argue that average precision is a more robust scoring mechanism and propose a new metric, Keypoint Average Precision (KAP), to account for potential spurious matches.

Our contributions are: i) a novel SC approach that learns to separate visually similar parts by mapping them to different location on a sphere, ii) a set of simple geometric constraints that force separation of visually similar parts during training, and iii) a new evaluation protocol that better accounts for the failure cases of current methods.

## 2. Related work

**Image matching.** Visual correspondences were first explored as exact matches between different views of the same scene with applications to stereo vision [24, 38], object detection [14, 33], structure from motion [20, 44, 52], and SLAM [12, 19]. Correspondences can also be computed across pixels in the frames of a video, which is referred to as optical flow [5, 7]. Traditionally they are obtained by computing hand-crafted visual descriptors [11, 32, 43].

**Semantic correspondence (SC).** In SC estimation [30], the goal is not only to find matches across different views of the exact same object instance, but also across different instances of the same object category. While earlier work [30] match two images through densely sampled SIFT [32] features, later works extended pairwise correspondence learning to a global matching problem [61] and established region correspondences using object proposals and their geometric relations over dense descriptors [17]. Subsequently, CNNs trained with only image-level supervision have been shown to discover correspondences [31], paving the way for many deep learning correspondence systems [10, 18, 26, 59]. Typical correspondence-specific CNN architectures rely on aggregating features along the depth of a network [35, 53, 59] or use 4D convolutions spanning two images simultaneously [41]. Recently, transformers have been shown to be particularly suited to correspondence discovery thanks to their attention mechanism [9].

While the early works in unsupervised correspondence discovery [2, 27, 50] use on the convolutional representations, recent methods [1, 47, 54, 57] build on the success of transformer-based self-supervised representations [6, 40, 60] to discover salient and robust features, which subsequent works have applied to detect semantic correspondences successfully, often surpassing supervised methods. The features from generative image models have also been used in SC [22, 49, 57]. Two closely related works [16, 39] propose to create category-wide representations by aligning self-supervised feature maps. Unlike our approach, these works are limited to 2D, and find correspondences between images without knowing the underlying 3D structure of the world which renders them incapable of distinguishing different sides of objects (*e.g.* left *vs.* right side of a car). A concurrent work [58] also identifies geometry-related issues in SSL-based correspondence but targets a supervised training regime where our approach does not necessitate keypoint supervision.

**3D priors for correspondence.** In literature, 3D priors have been used for SC, mainly in the form of annotated meshes [8, 29, 37, 46, 62], where annotations can either be from image to mesh, or between meshes. [29, 37] in particular use a dense parametrization of the mesh surface similar to our sphere. When no mesh is available, precise camera pose can be used to recover sparse correspondences through constrained geometric optimization [48] or by training a neural field [56]. In comparison, our approach requires no meshes and only very coarse viewpoint supervision.

Alternatively, it is possible to estimate 3D models of object categories from image supervision alone [3, 15, 23, 25, 36, 55], from which correspondences are extracted. However, this is very challenging by itself, requiring complex pipelines capable of modeling accurate mesh deformations, pose, and rendering, which are all failure-prone.

### 3. Method

#### 3.1. Problem statement

**General formulation.** Given an unordered monocular image collection, we wish to learn a mapping  $f(\mathbf{I}, x) = \mathbf{z}$  that takes in an RGB image  $\mathbf{I} : \Lambda \rightarrow \mathbb{R}^3$ ,  $\Lambda \subset \mathbb{R}^2$  and a pixel location  $x \in \Lambda$  to an  $n$ -dimensional embedding  $\mathbf{z} \in \mathcal{Z} \subset \mathbb{R}^n$ . In particular, we would like  $f$  to map *semantic* concepts from different instances of the same category (e.g. center of front-right wheel of a car) to the same point in the embedding space  $\mathcal{Z}$ . Formally,  $f(\mathbf{I}, x) = f(\mathbf{I}', x')$ , when the pixels  $x$  in  $\mathbf{I}$ , and  $x'$  in  $\mathbf{I}'$  correspond to the same semantic part and instance. In practice, the corresponding location for  $(\mathbf{I}, x)$  in  $\mathbf{I}'$  can be obtained by finding the nearest neighbor of  $f(\mathbf{I}, x)$  in the embedding space, i.e.  $\arg \min_{x'} d_{\mathcal{Z}}(f(\mathbf{I}, x), f(\mathbf{I}', x'))$ , where  $d_{\mathcal{Z}}$  is a distance function over  $\mathcal{Z}$ . For simplicity, we will use  $f(\mathbf{I})$  to refer to the feature map obtained by evaluating  $f$  on all pixels.

An ideal embedding function should be general enough as to not only match related semantic parts across different object instances (e.g. match the front-left wheel across two different car models) but also specific enough to distinguish between different part instances within the same object instance (e.g. front-left wheel vs. front-right wheel). Learning this from image-based self-supervision alone is a challenging task due to the visual similarity of repeated parts and self-occlusion. As a result, it is quite challenging for correspondence estimators to learn that cars have four distinct wheels from single images without explicit supervision.

**Spherical mapping.** Instead, it is beneficial to endow a correspondence estimator with priors about the underlying 3D structure of objects. To this end, previous methods [29, 62] use detailed 3D meshes via a render-and-compare approach. However, they rely on the presence of groundtruth 3D meshes and the precise viewpoint information. Instead, following [50, 51], we use a simplified 3D structure, specifically a sphere  $S^2 \subset \mathbb{R}^3$  (denoted as  $S$  for simplicity) based on the assumption that the object categories of interest are approximately homeomorphic to spheres (i.e. objects without holes). While this may at first appear particularly ill-suited for some categories, e.g. bicycles, the geometric simplicity of  $S$  provides a convenient parameterization of an object’s surface by allowing easy computation of geometric constraints to separate visual and semantic similarity. In particular,  $S$  provides a simple geodesic distance through the dot product, making it straightforward to push different semantic parts apart, which we use to distinguish repeated object features. Formally, we are interested in building a *spherical mapping*  $f_S(\mathbf{I}, x)$  that maps pixels in  $\mathbf{I}$  to the surface of  $S$ . Unlike an arbitrary  $n$  dimensional space,  $S$  forms an object-centric specific coordinate system where each region ideally repre-

sents a distinct semantic part *and* instance on the surface of the object.

#### 3.2. Learning formulation

In the absence of ground truth correspondences, ensuring the semantic consistency of our spherical mapping function is very challenging. [50] enforces correspondence between image pairs through synthetically augmenting training images via thinplate splines. Unfortunately, such augmentations do not produce realistic 3D transformations such as 3D rotations which can result in the self-occlusion of parts. Instead, similar to [16, 39], we formulate the learning objective as one of aligning features from a pretrained SSL model  $\phi : \Lambda \rightarrow \mathcal{Z}_{\phi}$  between different instances of an object category. For simplicity, we will refer to the SSL embedding space  $\mathcal{Z}_{\phi}$  simply as  $\mathcal{Z}$  in the rest of the text.

While those approaches use flat atlases  $\mathbb{R}^2 \rightarrow \mathcal{Z}$ , we build a *spherical prototype*  $S_{\mathcal{Z}} : S \rightarrow \mathcal{Z}$  that maps a point on the surface of the sphere  $s \in S^2$  to a feature vector  $\mathbf{z} \in \mathcal{Z}$ . This formulation has two important benefits. First, the input domain of  $S_{\mathcal{Z}}$  is simply  $S$ , meaning that it does not depend on  $I$ , and it is shared across all images, constituting a category-based map. Second, and in contrast to [16, 39], because we chose the output of  $f_S$  to lie on the surface of a sphere, we can use it to enforce 3D priors during training (Sec. 3.3).

Intuitively, for a given category, the spherical prototype  $S_{\mathcal{Z}}$  should be able to reconstruct any self-supervised feature map  $\phi(\mathbf{I})$  when queried on the coordinates predicted by  $f_S$ , i.e.  $S_{\mathcal{Z}}(f_S(\mathbf{I}, x))$  must closely approximate  $\phi(\mathbf{I}, x)$  for any  $\mathbf{I}, x$ . Even if  $S_{\mathcal{Z}}$  encodes a perfect category prototype, it will be unable to encode information about the image background. Hence we use instance masks  $M(\mathbf{I})$  at training time to restrict the loss computation to only the pixels on an object’s surface. In practice, these can be obtained using pretrained segmentation models [21]. Thus, given a training image  $\mathbf{I}$  and a pixel location  $x$ , our primary training objective to learn the parameters of  $f_S$  and  $S_{\mathcal{Z}}$  can be formulated as the minimization of

$$\mathcal{L}_{rec} = \frac{1}{|\Lambda|} \sum_{x \in \Lambda} M(\mathbf{I}, x) \times \Gamma(\phi(\mathbf{I}, x), S_{\mathcal{Z}}(f_S(\mathbf{I}, x))), \tag{1}$$

where  $\Gamma$  is the cosine distance between the two inputs. Our spherical prototypes are implemented via a neural network that maps points in  $\mathbb{R}^3$  to  $\mathcal{Z}$ . Importantly, each point  $s \in S^2 \subset \mathbb{R}^3$  must be mapped independently, meaning that it should correspond to a unique point  $\mathbf{z} \in \mathcal{Z}$ , irrespective of which instance it comes from or where in the image this part is located. To model multiple categories, and to encourage information sharing between closely related ones, we use a single network conditioned on a category embedding. In practice,  $S_{\mathcal{Z}}$  is a visual transformer [13] *without*

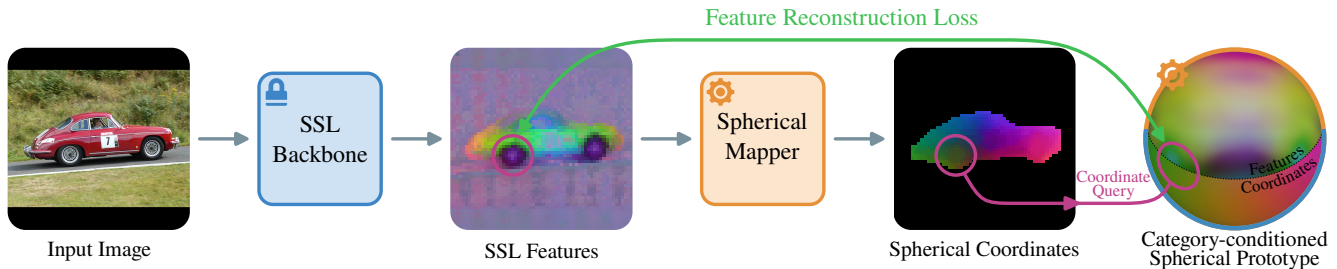


Figure 2. Overview of our semantic correspondence estimation approach. We begin by extracting features from a frozen self-supervised backbone, and further use them to predict spherical coordinates via a learned module. Each predicted point is used to query a jointly learned prototype, providing the supervision signal (Sec. 3.2). The sphere is used to enforce weak geometric priors (Sec. 3.3). During inference, SSL features are combined with spherical coordinates (Sec. 3.4). A blue outline indicates a fixed module, while orange indicate learned parameters.

self-attention, but with cross-attention between image tokens and a single category token  $c$  which is a one-hot embedding of the category. As learning  $f_S$  from scratch can be challenging, we further leverage the SSL backbone by simply adding a prediction head  $f'_S$  to it to produce spherical maps, *i.e.*  $f_S = f'_S \circ \phi$ . An overview of our architecture is presented in Fig. 2.

### 3.3. Enforcing geometric priors

The above spherical reconstruction formulation alone offers little benefit over matching the self-supervised features from  $\phi$  directly, especially considering the low dimensionality of the intermediate spherical space, which is not guaranteed to be well-behaved. However, this spherical structure enables the computation of simple geometric priors that can be enforced during training.

**Viewpoint regularization.** Many object categories exhibit some form of symmetry which can make it particularly challenging for image-based models to distinguish between different sides of an object, creating spurious correlations between distinct parts (*e.g.* see cars in Fig. 1). One potential way of addressing this challenge, without using dense groundtruth correspondence, is to exploit viewpoint information of the object instance.

An interesting property of using a spherical prototype to represent an object category is that it can easily be used to infer an object’s viewpoint with respect to the camera. We posit that the average coordinate of a spherical map of an image, which we present as  $\mu(f_S(\mathbf{I})) = 1/|\Lambda| \sum_{x \in \Lambda} f_S(\mathbf{I}, x)$ , can be viewed as a coarse approximation of the camera viewpoint under which the object is seen, *e.g.* right side views should be approximately mapped to the right side of the sphere, and left side views to the left. Intuitively, thinking of  $f_S$  as a mapping from image pixels to the object’s surface, the visible parts of an object should roughly be those closest to the camera, while the others should be hidden by self-occlusion.

By using the groundtruth camera viewpoint  $v_I$ , we en-

force the average direction of the spherical map  $\mu(f_S(\mathbf{I}))$  to align with  $v_I$ . When this is satisfied, symmetric views must be mapped to different parts of the sphere. While precise viewpoint information is hard to obtain, we observe it is sufficient to use *coarse* relative viewpoint supervision between images. In practice, viewpoints are discretized into a small number of bins, and  $v_I$  is represented by the median value of the corresponding bin. Then, we simply enforce that the dot product between viewpoints from two images equals that of their average spherical maps,

$$\mathcal{L}_{vp} = \sum_{\mathbf{I}, \mathbf{I}'} \|v_I \cdot v_{I'} - \mu(f_S(\mathbf{I})) \cdot \mu(f_S(\mathbf{I}'))\|^2. \quad (2)$$

**Relative distance loss.** As noted earlier, current self-supervised extracted features can confuse repeated object repeated parts. Depending on the target goal, this can be desirable, if the parts do indeed belong to the same part category. However, this is detrimental for dense correspondence estimation as the parts may actually be different part instances, *e.g.* left *vs.* right leg. It can be hard to distinguish between these based on their appearance alone and additional geometric context is required.

One trivial way of encoding this context is by making the assumption that similar parts that appear in different locations on an image (*i.e.* not beside each other) are in fact different instances of the same semantic part. Therefore, we would like to enforce the property on our predicted spherical maps that *distant pixels in the image should be mapped far apart on the sphere*. Due to scale variation, distance on the image plane is hard to link to distances on the surface of the sphere. However, we can reformulate this property as a relative distance comparison using a triplet. Given three points  $a$ ,  $b$ , and  $c$  on an image  $\mathbf{I}$ , and their corresponding positions  $s_a = f_S(\mathbf{I}, a)$ ,  $s_b$ , and  $s_c$  on the sphere, the ordering between distances should be preserved, that is:

$$\|a - b\| \leq \|a - c\| \iff \Gamma(s_a, s_b) \leq \Gamma(s_a, s_c). \quad (3)$$

In reality, this does not hold for distances along an object’s surface, which the sphere is supposed to represent, partly




Figure 3. Illustration of our geometry losses  $\mathcal{L}_{rd}$  and  $\mathcal{L}_o$ . The left image shows a spherical map from which a triplet of points is sampled.  $\mathcal{L}_{rd}$ : as the anchor patch  $a$  is closer to the positive  $b$  on the image compared with the negative  $c$ , its corresponding position  $s_a$  on the sphere must also be closer to  $s_b$  than  $s_c$ .  $\mathcal{L}_o$ : after projecting  $s_b$  and  $s_c$  to the plane tangent to the sphere at  $s_a$ , we ensure orientation is preserved by enforcing positive colinearity between  $u_b \times u_c$  and the normal vector  $n$ .

due to perspective. For example, in the case of a long train seen from a near frontal view, parts at the back of the train might appear very close to the front in the image, even though they are very far in 3D space. Nevertheless, we find that this is still an effective prior for discouraging repeated part instances from being mapped to the same location.

During training, point triplets  $(a, b, c)$  are randomly sampled among pixels that belong to the object. Among each such triplet, we chose an anchor location  $anc = a$ , a positive location  $pos = \arg \min_{x \in \{b, c\}} \|a - x\|$ , and a negative location  $neg = \arg \max_{x \in \{b, c\}} \|a - x\|$  (see Fig. 3). We can then define a triplet margin loss as:

$$\mathcal{L}_{rd} = \max(\Gamma(f_S(\mathbf{I}, anc), f_S(\mathbf{I}, pos)) - \Gamma(f_S(\mathbf{I}, anc), f_S(\mathbf{I}, neg)) + \delta, 0), \quad (4)$$

where  $\delta$  is a hyperparameter which specifies the margin. A useful side product of this formulation is that it encourages the spherical maps to be smooth, as strong discontinuities would produce a large anchor to positive distance.

**Orientation loss.** A final advantage of using  $S$  as a latent space is that it is an orientable 2D space, just like images. Therefore, enforcing that  $f_S$  preserves orientation prevents spurious matches between symmetric views. Intuitively, if a point  $a$  appears on the right side of  $b$  in the image, then  $s_a$  must appear in the same side of  $s_b$ . Given a point triplet as defined earlier, this can be enforced by making sure that the determinant of the image triplet and sphere triplet positively correlate. Formally, defining  $P_a$  as the orthogonal projection to the plane tangent to  $S$  at  $s_a$  and  $u_b$  as  $P_a(s_b) - s_a$ , we have  $d_I = \det(b - a, c - a)$  and  $d_S = \det(u_b, u_c)$ , and want  $sign(d_I) = sign(d_S)$  (see Fig. 3).

As objects can have complex geometry that our model does not have access to, there exists no exact relation linking the two determinants together, as for instance perspective distortions might alter angles. However, we can as-

sume that image triplets of large determinants should also have large determinant on the sphere, as a change of sign in determinant indicates the relative orientation on the sphere has been inverted. In practice, we randomly sample image triplets  $(a, b, c)$  and select those whose image determinant  $d_I$  is higher than a threshold  $d_\tau = 0.7$ , while swapping  $b$  and  $c$  if  $d_I$  is negative. Then, we enforce the determinant of the corresponding sphere triplet to be at least  $d_\tau$ ,

$$\mathcal{L}_o = \begin{cases} 0 & \text{if } d_I < d_\tau \\ \max(d_\tau - d_S, 0) & \text{if } d_I \geq d_\tau. \end{cases} \quad (5)$$

### 3.4. Correspondence via combined representations

During inference, it is possible to directly query the spherical maps  $f_S$  of two images to obtain correspondences using the cosine distance. That is, for an image pair  $\mathbf{I}, \mathbf{I}'$  and a query location  $q$  on  $\mathbf{I}$ , its location in  $\mathbf{I}'$  can be computed as

$$p^* = \arg \min_p \Gamma(f_S(\mathbf{I}, q), f_S(\mathbf{I}', p)). \quad (6)$$

While this protocol is sound, it comes with two important drawbacks. First,  $f_S$  produces spherical maps for the whole image, including the background, meaning segmentation masks would be required at inference time to prevent the emergence of spurious matches. Second, the spherical map is designed to be a smooth parameterization of the object surface, making it susceptible to missing fine details. A slightly incorrect mapping  $f_S(\mathbf{I}, q) + \epsilon$  has a high probability of finding a higher match in a nearby region of the correct match  $f_S(\mathbf{I}', p^*)$  if both maps are smooth. In comparison, SSL feature maps can exhibit strict feature separation, e.g. wheel features are very different from other nearby non-wheel features, making it unlikely that a wheel pixel finds a nearest neighbor among nearby non-wheel pixels.

To mitigate this issue, we make use of the feature maps that are already computed by the self-supervised backbone network  $\phi$ , in a similar way to supervised correspondence approaches that aggregate features from a network [35, 53, 59]. Existing literature has shown that DINO features are particularly effective at differentiating between the foreground and background [1], removing the need for inference-time segmentation. To leverage these DINO features, we reformulate Eq. (6) as the combination of self-supervised features and spherical locations:

$$p^* = \arg \min_p (1 - \alpha) \Gamma(\phi(\mathbf{I}, q), \phi(\mathbf{I}', p)) + \alpha \Gamma(f_S(\mathbf{I}, q), f_S(\mathbf{I}', p)), \quad (7)$$

where  $\alpha$  is a hyperparameter that balances each term.

## 4. Experiments

**Implementation details.** We build our models on DINOv1-B/8 [6] and DINOv2-B/14 [40] backbones. Training masks




















																				avg
Custom	CATS [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	49.9
	MMNet+FCN [59]	55.9	37.0	65.0	35.4	50.0	63.9	45.7	62.8	28.7	65.0	54.7	51.6	38.5	34.6	41.7	36.3	77.7	62.5	50.4
	SCorSan [26]	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	<u>77.7</u>	<b>69.7</b>	54.4
DINOv1	DINOv1 [6]	44.3	26.8	57.6	22.0	29.3	32.8	19.7	54.0	14.9	40.1	39.3	29.3	29.0	37.0	20.0	28.2	40.6	21.1	32.6
	ASIC [16]	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	37.0
	Ours	47.1	26.0	70.9	21.8	37.5	34.9	32.4	60.0	23.2	53.6	48.5	42.5	28.3	42.7	21.1	41.9	39.7	41.7	39.7
SD	DIFT [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	<u>57.8</u>	57.1	71.1	<u>63.4</u>	57.7
	SD [57]	63.1	55.6	80.2	33.8	44.9	49.3	47.8	74.4	38.4	70.8	53.7	61.1	54.4	55.0	54.8	53.5	65.0	53.3	56.1
DINOv2	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	<u>68.2</u>	30.7	62.0	54.3	24.2	56.2
	DINOv2 + SD [57]	73.0	<u>64.1</u>	<u>86.4</u>	40.7	<b>52.9</b>	55.0	53.8	78.6	45.5	77.3	64.7	69.7	63.3	<b>69.2</b>	<b>58.4</b>	67.6	66.2	53.5	63.3
	Ours (sphere only)	46.7	28.8	66.3	33.0	36.5	66.6	59.1	74.9	25.4	65.7	50.1	52.7	27.1	13.7	15.8	46.6	73.5	36.7	45.5
	Ours	<b>76.9</b>	61.2	85.9	<u>42.1</u>	48.4	<u>73.3</u>	<u>67.2</u>	<u>80.0</u>	<u>46.3</u>	<u>80.2</u>	<u>66.7</u>	<u>71.2</u>	<u>66.0</u>	63.9	36.2	<u>68.6</u>	67.8	42.2	<u>63.6</u>
	Ours + SD	<u>74.8</u>	<b>64.5</b>	<b>87.1</b>	<b>45.6</b>	<u>52.7</u>	<b>77.8</b>	<b>71.4</b>	<b>82.4</b>	<b>47.7</b>	<b>82.0</b>	<b>67.3</b>	<b>73.9</b>	<b>67.6</b>	60.0	49.9	<b>69.8</b>	<b>78.5</b>	59.1	<b>67.3</b>

Table 1. Keypoint matching scores on SPair-71k evaluated using PCK@0.1 with *macro*-averaging for the summary scores. We present our approach using DINOv1 [6] features (middle rows) and DINOv2 [40] features (bottom rows). In both cases, we improve over the DINO only baselines and are superior to fully supervised methods (top rows). **Bold** entries are best per category and underlined are second-best.

are obtained using a pretrained Mask R-CNN model [21]. Our spherical mapper consists of a linear dimension reduction layer halving the number of features, a visual transformer block with self-attention, and a final linear layer with an output dimension of three. The resulting feature map is then normalized pixel-wise so that its output values lies on the surface of  $S^2$ . The dimensionality reduction layer, while not strictly necessary, is helpful considering the small number of images used during training, while the transformer block helps in disentangling sides and repeated parts by allowing global reasoning between image patches.

The final loss term is computed as  $\mathcal{L} = \mathcal{L}_{rec} + \lambda_{rd}\mathcal{L}_{rd} + \lambda_o\mathcal{L}_o + \lambda_{vp}\mathcal{L}_{vp}$ . We set  $\lambda_{rd} = \lambda_o = 0.3$ , while  $\lambda_{vp} = 0.1$  is set slightly lower as it can have the detrimental effect of pulling all predictions towards the average over image pixels - cf. Eq. (2). The relative distance margin is set to  $\delta = 0.5$  to encourage feature separation. To perform matching in Eq. (7),  $\alpha$  is set to 0.2 so that we leverage the highly discriminative DINO features, while using the sphere as a symmetry and repetition-breaking mechanism. Our model is trained for a total of 200 epochs using Adam [28].

**Datasets.** Following existing work [16, 26, 57], we evaluate our method primarily on the Spair-71k dataset [34], which contains images from 18 different object categories. The evaluation set contains between 600 and 900 image pairs, each of which are annotated with the keypoints that are visible in both images. We train our model on the Spair-71k training split on all categories simultaneously, without using any keypoint annotations. Additionally, we evaluate on the AwA-pose dataset [4], which contains images from 35 quadruped categories, annotated with keypoints. We also present results by training on the Freiburg cars dataset [45]. This dataset was collected by densely sampling approximately 100 images each around 48 different car instances. The higher image count, coupled with  $360^\circ$  dense sampling, enables us to learn a more detailed spherical prototype.

#### 4.1. Keypoint average precision

While issues with symmetries are clear in Fig. 4, we argue that the PCK metric fails to account for them as it is only computed between keypoints that appear in *both* the source and target image. Hence, models that predict high similarity between two visually similar but semantically different keypoints might not be penalized for it. PCK can partly account for repetition-related mistakes if the parts appear simultaneously in both images, but it does not penalize symmetry-related issues, like mistaking the left and right side of a car, as the corresponding keypoints do not appear simultaneously. PCK was initially proposed for simpler cases where keypoints were assumed to be always be visible, e.g. faces or different instances in the same pose. However, as benchmarks become more challenging, a more robust protocol should be used. To illustrate how much these failure cases impact SSL models, we propose an alternative evaluation metric, Keypoint Average Precision - KAP@ $\kappa$ . It is similar to PCK@ $\kappa$ , but accounts for these errors by evaluating average precision instead of accuracy and penalizes methods that predict matches when none exist.

When performing evaluation on a pair of images, instead of restricting it to keypoints that simultaneously appear in both images, we consider *all* keypoints that appear in the source, irrespective of whether they are also visible in the target image. For each such keypoint, if it appears in the target image, we extract the highest similarity value within radius  $\kappa$  of the ground-truth and label it as positive. We also extract the highest similarity value *outside* the radius and label it as negative, penalizing overly high predictions on wrong, possibly repeated parts. If the source keypoint does not appear on the target image, we simply extract the highest similarity prediction across the whole target image and label it as negative to penalize incorrectly high similarity when none should exist. This effectively reformulates the task as a binary classification problem, where the embed-

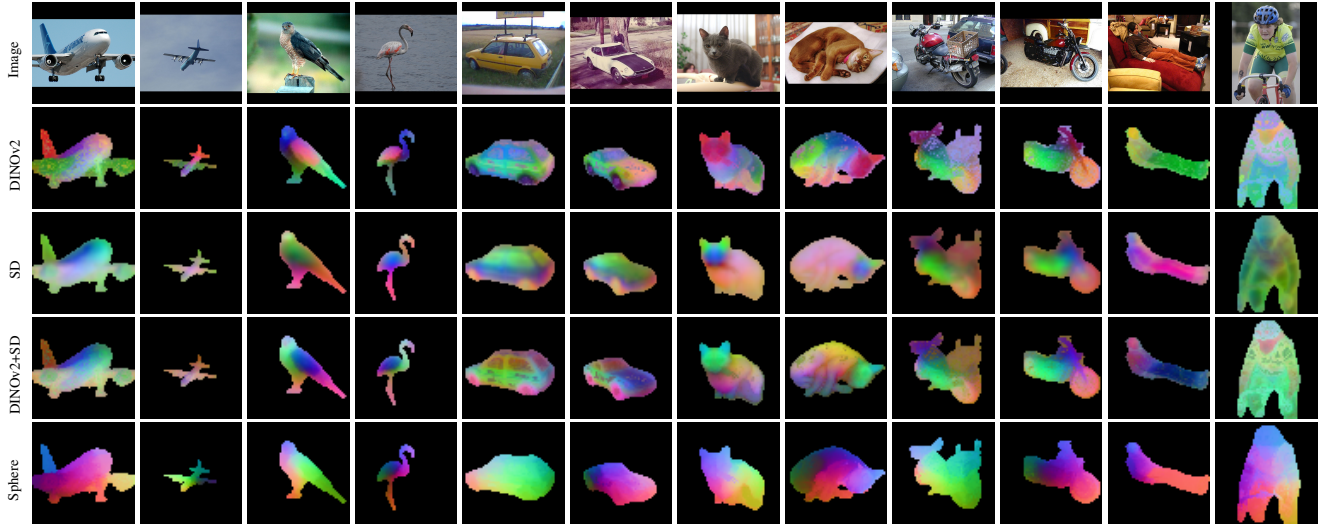


Figure 4. Qualitative comparison of dense correspondence maps. For DINOv2, SD, and DINO+SD features we perform PCA on the segmented object features independently for each category, then visualize the three main components. Note that the SD and DINO+SD features are not completely equivalent to the ones used to compute matches, but are provided here for illustration. Spherical maps from  $f_S$  (Sphere) for our approach are visualized directly. Our spherical maps correctly identify the different sides of objects, whereas other features fail to capture these differences.

ding of each source keypoint has to be close to, and only to, the corresponding target keypoint. Finally, we compute the average precision over all the extracted pairs, and report mean average precision per-category.

## 4.2. SPair-71k results

We first evaluate on the task of keypoint transfer between pairs of images using Spair-71k [34]. We present quantitative per-category results in Tab. 1, using the standard PCK metric along with *macro*-averaged scores as the number of per-category pairs varies greatly. We split the models according to their backbone: custom supervised models, DINOv1-based, SD-based, and DINOv2-based. A striking preliminary observation is the effectiveness of DINOv2, surpassing supervised approaches on most categories. While its success can partially be attributed to evaluation biases [2], this demonstrates the effectiveness of large scale learned self-supervised features.

Using a DINOv1 backbone, our model improves over its backbone on most categories, and over ASIC [16] on average. When combined with the DINOv2 backbone, our model provides improvements on all but two categories. A clear pattern can be identified, with strongest improvements being observed over blob-like symmetric objects with repeated parts (+21.0 on bus and +15.7 on car), while spherical mapping hinders performance on highly deformable categories (-4.3 on person), or high-genus objects for which a sphere is a poor surface prior (-0.8 on bicycle). Our model performs slightly better compared to DINO+SD [57], while requiring a fraction of the computational cost at inference time. Finally, following [57] we evaluate adding Stable Dif-

fusion (SD) [42] features to our method at inference time, which yields further improvements. This illustrates that spherical maps and SD features capture different facets of the correspondence problem.

We show the results of evaluating correspondences with KAP in Tab. 2. We observe particularly large discrepancies between PCK and KAP in highly symmetric objects with repeated parts, *e.g.* buses and cars, and the ranking between our model and DINOv2 is reversed on humans, possibly due to DINOv2’s symmetric predictions on left and right limbs. All models across all categories obtain lower results when using KAP highlight the symmetry issues they are facing. A limitation of KAP is that it might still count correct repeated parts mistakes if their distance in the image is less than  $\kappa$ . Nonetheless, we argue that KAP is more informative than PCK and is just as efficient to evaluate.

The qualitative visualization shown in Fig. 4 illustrates these improved results. Our model correctly maps the two opposite sides of the cars to different sphere regions, while other feature maps look extremely similar, even when looking at two opposite sides of an object, with the notable exception of SD on birds. The fact that such similarities exist between opposite views suggest that the utility of baselines such as DINOv2 would be quite limited in real scenarios.

## 4.3. Additional evaluation

**Animals with Attributes (AwA).** To further evaluate our model, we also report results on the AwA-pose dataset from [4], which contains images from 35 different categories of quadruped. While our spherical prototype needs category information during training to reconstruct




















																				avg
DINOv2 [40]	53.5	54.0	60.2	35.5	44.4	36.3	31.7	61.3	37.4	54.7	52.5	51.5	48.8	48.2	37.8	44.1	47.4	38.2	46.5	
SD [57]	44.4	48.5	54.5	31.5	45.2	32.7	30.0	68.4	35.8	55.2	47.9	48.1	44.8	42.3	<u>44.5</u>	39.2	52.7	<u>51.2</u>	45.4	
DINOv2 + SD [57]	52.0	<b>55.9</b>	59.2	34.7	<b>49.0</b>	36.0	32.5	70.3	39.8	59.8	53.1	52.4	50.6	<b>50.4</b>	<b>47.8</b>	46.2	53.3	49.8	49.6	
Ours (sphere only)	38.4	34.2	53.9	33.0	37.9	49.7	43.4	<u>71.7</u>	29.8	57.1	45.8	42.5	32.4	27.0	29.5	37.1	<u>57.4</u>	36.0	42.1	
Ours	<b>60.7</b>	51.2	<b>63.1</b>	<u>38.4</u>	45.0	<b>55.9</b>	<u>45.7</u>	69.7	40.4	<u>63.2</u>	<u>54.8</u>	<u>54.3</u>	<u>51.2</u>	<u>48.7</u>	38.8	<u>47.9</u>	55.5	42.2	<u>51.5</u>	
Ours + SD	<u>58.9</u>	<u>54.2</u>	<u>62.2</u>	<b>39.6</b>	<u>46.6</u>	<u>54.5</u>	<b>47.1</b>	<b>76.2</b>	<b>40.9</b>	<b>65.3</b>	<b>57.3</b>	<b>56.1</b>	<b>54.2</b>	47.4	43.7	<b>49.4</b>	<b>62.4</b>	<b>52.0</b>	<b>53.8</b>	

Table 2. Keypoint matching scores on SPair-71k evaluated using KAP@0.1 with macro-averaging for the summary scores.

category-specific features, only the sphere mapper is used during inference, meaning it can be readily applied to unseen categories. Five quadruped classes (cat, cow, dog, horse, and sheep) appear in both SPair and AWA-pose, and are therefore seen during training, but the rest are completely new for our sphere mapper. Nonetheless, results shown in Tab. 3 show that the sphere mapper generalizes well to unseen categories. We attribute this to the assistance from the strong category-agnostic features from DINOv2. Once again, the improvement of using our spherical method is much more apparent using KAP, *i.e.* the PCK improvement over DINOv2 alone is 2.2 when adding SD and 2.6 when adding spheres, while these become 1.8 and 3.9 respectively with KAP.

	Dv2	SD	Dv2+SD	Ours	Ours+SD
PCK@0.1	65.9	56.0	68.1	68.7	69.8
KAP@0.1	55.0	50.7	56.8	58.9	60.6

Table 3. Average scores when evaluating on AWA-pose using a random subset of 200 pairs per category. ‘Ours’ denotes our model trained on SPair-71k with a DINOv2 backbone.

**Model ablation.** We also analyze the impact of removing each individual loss term when training our model in an ablation study. The results in Tab. 4 indicate that all variants except the no  $\mathcal{L}_{rd}$  one do not significantly improve over the DINOv2 baseline as they also fail to identify opposite sides of objects and repeated parts.

DINOv2	no $\mathcal{L}_{vp}$	no $\mathcal{L}_{rd}$	no $\mathcal{L}_o$	full model
56.2	58.6	61.2	56.0	63.6

Table 4. Average PCK@0.1 scores when training and testing on SPair-71k using different ablated versions of our approach.

**Inference speed.** The results in Tab. 1 demonstrate how effective the addition of Stable Diffusion features are. However, computing correspondences through diffusion requires adding noise to images, running iterative diffusion, and performing joint decomposition between the source and target features. As a result, its inference speed is much slower compared to using DINOv2 features. In comparison, our approach only requires a few extra layers on top of DINOv2 and has negligible additional overhead. Timing results in Tab. 5 illustrate that adding diffusion-extracted features reduces throughput by an order of magnitude.

DINOv2	SD	DINOv2+SD	Ours	Ours+SD
3.4	0.38	0.35	3.3	0.34

Table 5. Descriptor computation throughput in pairs/second at inference time on a single A5000 GPU, where higher is better.

#### 4.4. Limitations

While our approach results in significant improvements over the DINOv2 baseline for many categories (see Tab. 1), there are still some limitations. For example, we do not perform as well on a small number of categories that cannot be well approximated by a sphere (*e.g.* humans). Furthermore, the very low-dimensional spherical encoding on its own is not as effective at discriminating parts and sub-parts and thus needs to be combined with self-supervised features. However, as we already densely compute these features at inference time for each test image, there is no additional cost in combining them with our spherical mapping when evaluating correspondence. Finally, our method relies on additional supervision at training time in the form of automatically computed segmentation masks and weak camera pose supervision. However, results in the supplementary material suggest that coarse camera viewpoints are enough to obtain good performance.

## 5. Conclusion

We presented a new approach for semantic correspondence estimation that is robust to issues resulting from object symmetries and repeated part instances. In addition, we proposed a new evaluation protocol that is more sensitive to these issues. Our approach leverages recent advances in the self-supervised learning of discriminative image features and combines them with a weak geometric spherical prior. By using weak pose supervision to represent objects categories on the surface of spheres, we are able to enforce simple geometric constraints during training that result in more geometric and semantically consistent representations. Results on the task of semantic keypoint matching between images from different object instances, a fundamental task in 3D computer vision, demonstrate that our approach alone is superior to existing strong baselines and can be combined with existing methods to further boost performance.

**Acknowledgments.** This project was supported by the EP-SRC Visual AI grant EP/T028572/1.



## References

- [1] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *ECCV Workshop - What is Motion For?*, 2022. 1, 2, 5
- [2] Mehmet Aygün and Oisín Mac Aodha. Demystifying unsupervised semantic correspondence estimation. In *ECCV*, 2022. 1, 2, 7
- [3] Mehmet Aygün and Oisín Mac Aodha. SAOR: Single-View Articulated Object Reconstruction. In *CVPR*, 2024. 2
- [4] Prianka Banik, Lin Li, and Xishuang Dong. A novel dataset for keypoint detection of quadruped animals from images. *arXiv:2108.13958*, 2021. 6, 7
- [5] Steven S. Beauchemin and John L. Barron. The computation of optical flow. *ACM computing surveys*, 1995. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 1, 2, 5, 6
- [7] Zhuoyuan Chen, Hailin Jin, Zhe Lin, Scott Cohen, and Ying Wu. Large displacement optical flow from nearest neighbor fields. In *CVPR*, 2013. 2
- [8] An-Chieh Cheng, Xueting Li, Min Sun, Ming-Hsuan Yang, and Sifei Liu. Learning 3d dense correspondence via canonical point autoencoder. *NeurIPS*, 2021. 2
- [9] Seokju Cho, Sunghwan Hong, Sangryul Jeon, Yunsung Lee, Kwanghoon Sohn, and Seungryong Kim. Cats: Cost aggregation transformers for visual correspondence. *NeurIPS*, 2021. 1, 2, 6
- [10] Christopher B Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. *NeurIPS*, 2016. 2
- [11] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005. 2
- [12] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *PAMI*, 2007. 2
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3
- [14] Vittorio Ferrari, Tinne Tuytelaars, and Luc Van Gool. Simultaneous object recognition and segmentation from single or multiple model views. *IJCV*, 2006. 2
- [15] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *ECCV*, 2020. 2
- [16] 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, 3, 6, 7
- [17] Bumsu Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow: Semantic correspondences from object proposals. *PAMI*, 2017. 2
- [18] Kai Han, Rafael S Rezende, Bumsu Ham, Kwan-Yee K Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. Sennet: Learning semantic correspondence. In *ICCV*, 2017. 2
- [19] Christopher G Harris and JM Pike. 3d positional integration from image sequences. *Image and Vision Computing*, 1988. 2
- [20] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 2
- [21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 3, 6
- [22] 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
- [23] Paul Henderson and Vittorio Ferrari. Learning to generate and reconstruct 3d meshes with only 2d supervision. In *BMVC*, 2018. 2
- [24] Heiko Hirschmüller. Accurate and efficient stereo processing by semi-global matching and mutual information. In *CVPR*, 2005. 2
- [25] Tao Hu, Liwei Wang, Xiaogang Xu, Shu Liu, and Jiaya Jia. Self-supervised 3d mesh reconstruction from single images. In *CVPR*, 2021. 2
- [26] Shuaiyi Huang, Luyu Yang, Bo He, Songyang Zhang, Xuming He, and Abhinav Shrivastava. Learning semantic correspondence with sparse annotations. In *ECCV*, 2022. 1, 2, 6
- [27] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. *NeurIPS*, 2018. 2
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [29] Nilesh Kulkarni, Abhinav Gupta, and Shubham Tulsiani. Canonical surface mapping via geometric cycle consistency. In *ICCV*, 2019. 2, 3
- [30] Ce Liu, Jenny Yuen, Antonio Torralba, Josef Sivic, and William T Freeman. Sift flow: Dense correspondence across different scenes. In *ECCV*, 2008. 1, 2
- [31] Jonathan L Long, Ning Zhang, and Trevor Darrell. Do convnets learn correspondence? *NeurIPS*, 2014. 2
- [32] David G Lowe. Object recognition from local scale-invariant features. In *ICCV*, 1999. 2
- [33] David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004. 2
- [34] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Spair-71k: A large-scale benchmark for semantic correspondence. *arXiv:1908.10543*, 2019. 6, 7
- [35] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Learning to compose hypercolumns for visual correspondence. In *ECCV*, 2020. 2, 5
- [36] Tom Monnier, Matthew Fisher, Alexei A Efros, and Mathieu Aubry. Share with thy neighbors: Single-view reconstruction by cross-instance consistency. In *ECCV*, 2022. 2
- [37] Natalia Neverova, David Novotny, Marc Szafraniec, Vasil Khalidov, Patrick Labatut, and Andrea Vedaldi. Continuous surface embeddings. In *NeurIPS*, 2020. 2
- [38] H K Nishihara. Prism: A practical realtime imaging stereo matcher. In *Intelligent Robots: 3rd International Conference on Robot Vision and Sensory Controls*, 1984. 2

- [39] Dolev Ofri-Amar, Michal Geyer, Yoni Kasten, and Tali Dekel. Neural congealing: Aligning images to a joint semantic atlas. In *CVPR*, 2023. 2, 3
- [40] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dino2: Learning robust visual features without supervision. *arXiv:2304.07193*, 2023. 1, 2, 5, 6, 8
- [41] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. *NeurIPS*, 2018. 2
- [42] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 7
- [43] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *ICCV*, 2011. 2
- [44] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 1, 2
- [45] Nima Sedaghat and Thomas Brox. Unsupervised generation of a viewpoint annotated car dataset from videos. In *ICCV*, 2015. 6
- [46] Roman Shapovalov, David Novotny, Benjamin Graham, Patrick Labatut, and Andrea Vedaldi. Densepose 3d: Lifting canonical surface maps of articulated objects to the third dimension. In *ICCV*, 2021. 2
- [47] Aleksandar Shtedritski, Andrea Vedaldi, and Christian Rupprecht. Learning universal semantic correspondences with no supervision and automatic data curation. In *ICCV Workshops*, 2023. 2
- [48] Supasorn Suwajanakorn, Noah Snavely, Jonathan J Tompson, and Mohammad Norouzi. Discovery of latent 3d keypoints via end-to-end geometric reasoning. *NeurIPS*, 2018. 2
- [49] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent correspondence from image diffusion. In *NeurIPS*, 2023. 2, 6
- [50] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object frames by dense equivariant image labelling. *NeurIPS*, 2017. 1, 2, 3
- [51] James Thewlis, Samuel Albanie, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of landmarks by descriptor vector exchange. In *ICCV*, 2019. 3
- [52] Philip HS Torr and Andrew Zisserman. Feature based methods for structure and motion estimation. In *International workshop on vision algorithms*, 1999. 2
- [53] Nikolai Ufer and Bjorn Ommer. Deep semantic feature matching. In *CVPR*, 2017. 2, 5
- [54] Matthew Walmer, Saksham Suri, Kamal Gupta, and Abhinav Shrivastava. Teaching matters: Investigating the role of supervision in vision transformers. In *CVPR*, 2023. 1, 2
- [55] Shangzhe Wu, Ruining Li, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. MagicPony: Learning articulated 3d animals in the wild. In *ICCV*, 2023. 2
- [56] Lin Yen-Chen, Pete Florence, Jonathan T Barron, Tsung-Yi Lin, Alberto Rodriguez, and Phillip Isola. Nerf-supervision: Learning dense object descriptors from neural radiance fields. In *ICRA*, 2022. 2
- [57] Junyi Zhang, Charles Herrmann, Junhwa Hur, Luisa Polania 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. 2, 6, 7, 8
- [58] 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. 2
- [59] Dongyang Zhao, Ziyang Song, Zhenghao Ji, Gangming Zhao, Weifeng Ge, and Yizhou Yu. Multi-scale matching networks for semantic correspondence. In *ICCV*, 2021. 1, 2, 5, 6
- [60] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. In *ICLR*, 2022. 1, 2
- [61] Tinghui Zhou, Yong Jae Lee, Stella X Yu, and Alyosha A Efros. Flowweb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In *CVPR*, 2015. 2
- [62] Tinghui Zhou, Philipp Krahenbuhl, Mathieu Aubry, Qixing Huang, and Alexei A Efros. Learning dense correspondence via 3d-guided cycle consistency. In *CVPR*, 2016. 2, 3