

# SHIC: Shape-Image Correspondences with no Keypoint Supervision

Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi

Visual Geometry Group, University of Oxford  
{sunny, chrisr, vedaldi}@robots.ox.ac.uk  
[robots.ox.ac.uk/vgg/research/shic/](https://robots.ox.ac.uk/vgg/research/shic/)



**Fig. 1: Unsupervised canonical maps.** We show predictions from our *fully unsupervised* method SHIC, which finds correspondences between a rigid 3D template and a natural image. Correspondences are color-coded by assigning a distinct color to each template surface point. Our approach is highly data-efficient; the elephant, T-Rex, and Appa models above are trained on only 2800, 480, and 180 images, respectively.

**Abstract.** Canonical surface mapping generalizes keypoint detection by assigning each pixel of an object to a corresponding point in a 3D template. Popularised by DensePose for the analysis of humans, authors have since attempted to apply the concept to more categories, but with limited success due to the high cost of manual supervision. In this work, we introduce SHIC, a method to learn canonical maps without manual supervision which achieves *better* results than supervised methods for most categories. Our idea is to leverage foundation computer vision models such as DINO and Stable Diffusion that are open-ended and thus possess excellent priors over natural categories. SHIC reduces the problem of estimating image-to-template correspondences to predicting image-to-image correspondences using features from the foundation models. The reduction works by matching images of the object to non-photorealistic renders of the template, which emulates the process of collecting manual annotations for this task. These correspondences are then used to supervise high-quality canonical maps for any object of interest. We also show

that image generators can further improve the realism of the template views, which provide an additional source of supervision for the model.

## 1 Introduction

Correspondences play an important role in computer vision, with applications to pose estimation, 3D reconstruction, retrieval, image and video editing and many more. In this paper, we consider the problem of learning dense keypoints for any given type of objects without manual supervision. Keypoints identify common object parts, putting them in correspondence, and providing a key abstraction in the analysis of the objects’ geometry and pose. While keypoints are usually small in number, dense keypoints [9] are a generalization that considers a continuous family of keypoints indexed by the surface of a 3D template of the object. Dense keypoints provide more nuanced information than sparse ones and have found numerous applications in computer vision and computer graphics.

Despite their utility, learning keypoints, especially dense ones, remains labour-intensive due to the need to collect suitable manual annotations. Because of this, most keypoint detectors are limited to specific object classes of importance in applications, such as humans [9, 16, 33, 49]. Methods that generalize to more categories either have limited performance [17, 18], or require a significant amount of manual annotations for each class [24, 25]. They cannot scale to learning (dense) keypoints for the vast majority of object types in existence.

In contrast, foundation models such as DINO [4], CLIP [31], GPT-4 [27], DALL-E [32], and Stable Diffusion [35] are trained from billions of Internet images and videos with almost no constraints on the type of content observed. While these models do not provide explicit information about the geometry of objects, we hypothesise that they may do so *implicitly* and may thus be harnessed to generalize geometric understanding to more object types.

In this paper, we test this hypothesis by utilizing off-the-shelf foundation models to learn automatically high-quality dense keypoints. Given a single template mesh for an object class (*e.g.*, a horse or a T-Rex) to define the index set for the keypoints, and as few as 1,000 masked example images of the given class, we learn a high-quality image-to-template mapping.

Our method builds on recent advances in self-supervised image-to-image matching algorithms which, by using features from DINO [4] and the Stable Diffusion encoder [35], can generalize surprisingly well across images of different modalities or styles, such as natural images, animations or abstract paintings. Our idea is to reduce the problem of matching images to the 3D template to the one of matching images to *rendered views of the template*. Namely, we render a view of the 3D template and, given a query location in the source image, we find the corresponding vertex as a visual match on the rendered images. The template renders are *not* photorealistic, so the matching process emulates the process of manually annotating dense keypoints in prior works [9, 24]. We contribute several ideas to robustly pool information collected from different renders of the template, including accounting for visibility.


The approach we have described so far is training-free, as it uses only off-the-shelf components, but it is slow and the resulting correspondences lack spatial smoothness as they are established greedily. Our second step is thus to use these initial correspondences to supervise a more traditional dense keypoint detector in the form of a canonical surface map [9, 18, 38]. We utilize the Canonical Surface Embedding (CSE) representation of [25], which was designed to learn a mapping for several proximal object classes together (*e.g.*, cow, dog and horse), and can also efficiently represent image-to-template and image-to-image mappings by learning cross-modal embeddings. The most important result is that we can *outperform* the original manually-supervised model of [25] on their animal classes *without* using any supervision. This means that we can also learn maps for entirely new classes, such as T-Rex or Appa (a flying bison from a TV show), essentially at no cost (Fig. 1).

Finally, we note a further use of foundation models for our application: the generation of photorealistic synthetic images of the object. In particular, we show that a version of Stable Diffusion conditioned on depth can be used to texture the 3D images of the template, significantly narrowing the synthetic-to-real gap. These images are good enough to be used to supervise the dense pose map *directly*, with full synthetic supervision. We show that, while this is no substitute for utilizing real images as described above, it does improve the final performance further.

## 2 Related work

*Unsupervised image-to-image correspondences.* Many authors have sought to establish correspondences between images without manual supervision to address the cost of obtaining labels for this task. Early methods generate training data by applying synthetic warps to images [5, 22, 34, 40, 41], or use cycle consistency losses [14, 36, 42, 43]. GANs have also been used to supervise dense visual alignment [29]. Recent advances in self-supervised representation learning [4, 28] and generative modelling [35] have boosted the quality of unsupervised semantic correspondences significantly. For instance, [1] establish correspondences by seeking matches between DINO features, and [13, 19, 21, 37, 50] use Stable Diffusion instead. Similarly, [7, 23] use diffusion features to find mesh-to-mesh correspondences. [50] show that DINO and Stable Diffusion features are complementary, the first capturing precise but sparse correspondences and the second the general layout, and propose to combine them. In our work, we use the SD-DINO [50] features for matching images.

*Animal pose estimation.* While most works on pose estimation focus on humans [2, 3, 8, 9, 26, 45], several authors have attempted to estimate the pose of animals by detecting [53], matching [15] or reconstructing [46, 47] them, or predicting the parameters of parametric models [54–56]. However, these methods do not scale well as they need annotations for each type of animal considered. Our method is most similar to [17, 18] in that we only require a template shape and a collection



**Fig. 2: Image-to-template correspondences using 2D renderings.** Using an unsupervised semantic correspondence method, we can find correspondences between an image of an object and a rendering of its 3D template. Here we show the similarity heatmap from the source location (annotated in red) to all pixel locations in the target image using SD-DINO [50].

of images to learn image-to-shape correspondences. However, our method achieves much better performance, while still using fewer images for training.

*Image-to-template correspondences.* Finding correspondences between images and a 3D template is useful for understanding the geometry of deformable objects, with several applications. For instance, it is used in biology to study the behaviour of animals [30, 44]. Most prior works focus on humans [9, 16, 33, 49] due to the availability of large-scale datasets of densely annotated image-template pairs, such as DensePose-COCO [9]. Similar datasets exist for animals, such as DensePose-LVIS [25], but are much smaller and still only cover a handful of animal classes. To learn image-to-shape correspondences, [17, 18] parametrise a 3D shape as a 2D  $uv$  map, and use cycle consistency to try to learn correspondences automatically, whereas [17] also learn to predict articulation. Similarly to [25], we use the CSE representation for learning the correspondences. However, differently from [25], our method does not rely on any human-annotated data.

### 3 Method

In this section, we describe SHIC, our method for learning dense keypoints without manual supervision. First, in Sec. 3.1 we recall the notion of dense keypoints, canonical surfaces and canonical surface maps. Then, in Sec. 3.2 we discuss using self-supervised features to establish dense semantic correspondences between pairs of images, lift those to dense keypoints in Sec. 3.3, and use the latter to supervise a canonical map in Sec. 3.4. Finally, in Sec. 3.5, we show how an image generator can produce realistic views of the template, which further improves results.

### 3.1 Canonical surface maps

Let  $I \in \mathbb{R}^{3 \times \Omega}$  be an image supported by the grid  $\Omega = \{1, \dots, H\} \times \{1, \dots, W\}$ . The image contains an object of a given type, such as a cat, and the goal is to assign an identity to each pixel  $u \in U_I$  of the object, where  $U_I \subset \Omega$  is the object mask in image  $I$ . The identification is carried out by a mapping  $f_I : U_I \rightarrow M$  that assigns each pixel  $u$  to a corresponding index  $f_I(u)$  in a set  $M$ . The set  $M \subset \mathbb{R}^2$  is a 2D surface embedded in  $\mathbb{R}^3$ , and is interpreted as a (fixed and rigid) 3D template of the object. The template  $M$  is also called a *canonical surface* and the function  $f_I$  a *canonical surface map*. The same canonical surface  $M$  is shared by all objects of that category. In this way, by mapping two images  $I$  and  $J$  to the same template, one can also infer a mapping between the images.

In practice, we approximate the surface  $M$  by a mesh supported by a finite set of  $K$  vertices  $V = \{x_1, \dots, x_K\} \subset M$  and triangular faces  $F$ . Hence the canonical map is a function  $f_I : \Omega \rightarrow V \subset M$ . This slightly simplifies the formulation as both index sets  $\Omega$  and  $V$  are finite. We also note that the value  $f_I(u)$  is undefined if pixel  $u \in \Omega - U_I$  does not belong to the object.

In prior works, learning the canonical map  $f$  often requires hundreds of thousands of manually specified image-to-template correspondences. In the next sections, we will show how to learn this mapping *automatically* instead.

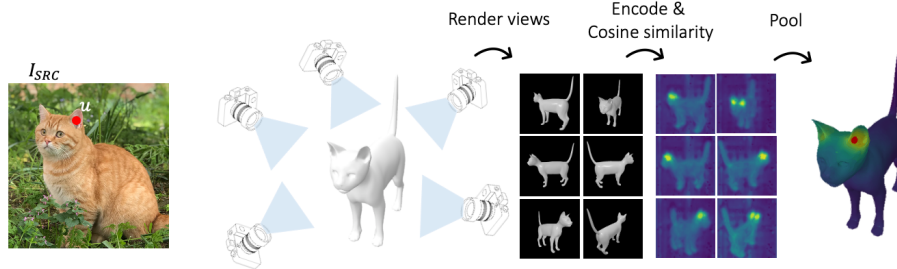
### 3.2 Unsupervised image-to-image correspondences

In order to learn the canonical map  $f$  automatically, we start by establishing correspondences between pairs of images  $I$  and  $J$  in an unsupervised fashion. We do so by first computing  $D$ -dimensional dense features  $\Phi \in \mathbb{R}^{D \times \Omega}$  using a pre-trained network. Then, we associate each query location  $u$  in the source image  $I$  to the location  $v_u$  in the target image  $J$  with the most similar feature vector based on the cosine similarity, *i.e.*,

$$v_u = \operatorname{argmax}_{v \in \Omega} S_{IJ}(u, v) \quad \text{where} \quad S_{IJ}(u, v) = \frac{\Phi_I(u) \cdot \Phi_J(v)}{\|\Phi_I(u)\|_2 \|\Phi_J(v)\|_2}.$$

The quality of the correspondences depends on the quality of the feature extractor  $\Phi$ . In particular, by using the unsupervised features by [50], it is possible to establish good correspondences between a (real) image  $I$  of the object and a *rendering* of the 3D template  $M$ .

This is illustrated in Fig. 2, where we show the cosine similarity heatmaps between a feature at a query location  $u$  of in the source image  $I$  and all locations in several 3D renders of the template. While the correspondences correctly identify the type of body part (paw), two problems are apparent: (i) there is left-right ambiguity, which is common for unsupervised semantic correspondence methods [51], (ii) when the correct match is not visible (as on the top of Fig. 2, where only the back paws are visible), the correspondence will always be wrong. In the next section, we lift these image-based correspondences into correspondences with the template  $M$ , which also alleviates these issues.



**Fig. 3: Zero-shot image-to-template correspondences.** From left to right: an image  $I$  with a selected pixel  $u$ ; several views  $J_i$  of the synthetic template; corresponding renderings and similarities  $S_{I, J_i}(u, v)$  as functions of the target locations  $v \in \Omega$ ; the final similarities  $\Sigma_I(u)$  visualized as a heatmap on top of the canonical surface  $M$ . The maximizer of the latter (red dot) identifies the vertex  $x_k$  that best corresponds to the selected pixel  $u$  in the source image  $I$  (*i.e.*, base of the left ear of the cat).

### 3.3 Unsupervised image-to-template correspondences

Given the source image  $I$  and a pixel  $u$ , we now consider the problem of finding the vertex  $x_k$  in the template  $M$  that best represents it. In order to do so, we develop a similarity measure between pixels and vertices, utilizing the image-to-image similarity metric of Sec. 3.2. We first generate  $N$  different views  $J_i = \text{Rend}(M, c_i)$ ,  $i = 1, \dots, N$ , of the canonical surface  $M$  by rendering it from viewpoints  $c_i$  (camera parameters). For each view  $J_i$ , we project each vertex  $x_k$  to its closest location in the mask  $U_{J_i}$ , defining

$$v_i(k) = \underset{v \in U_{J_i}}{\operatorname{argmin}} \|v - \pi(x_k, c_i)\|, \quad (1)$$

where  $\pi(x_k, c_i)$  is the camera projection function. We also denote by  $V_i \subset V$  the subset of vertices  $x_k$  that are visible in view  $J_i$ .

Given this notation, we can define a new score  $\Sigma$  measuring the compatibility between each location  $u$  in the source image  $I$  and each vertex  $x_k \in V$  in the canonical surface, and corresponding matches  $\tilde{x}$ , as follows:


$$\Sigma_I(u, x_k) = \underset{i: x_k \in V_i}{\operatorname{pool}} S_{I, J_i}(u, v_i(k)), \quad \tilde{x}(u, I) = \underset{x_k \in V}{\operatorname{argmax}} \Sigma_I(u, x_k). \quad (2)$$

The goal of the pooling operator is to assess the compatibility between pixel  $u$  and vertex  $x_k$  into a single score that consolidates the information collected from the different viewpoints  $c_i$ . Note that only the views where the vertex is visible are pooled. In practice, we set the pooling operator to average or max pooling.

*Illustration.* Figure 3 illustrates the similarity maps  $S_{I, J_i}$  between the source image  $I$  and various views  $J_i$  of the rendered 3D object, as well as the result  $\Sigma_I$  of mapping and pooling them on the canonical surface itself. We see that the correct semantic parts on the shape are identified (ears), and the base of the left ear is selected as the most similar to the query  $u$ .

*Template realism and rendering function.* The template  $M$  captures the typical shape of the object. Mathematically, its main purpose is to define the *topology* of the object’s surface, but the latter is usually topologically equivalent to a sphere [38]. In dense pose [24] there are two reasons for not using a sphere. The first is that

the *metric* of the template surface can be used to regularize correspondences (e.g., by capturing an approximate notion of how far apart physical points are). The second is that *renders* of the template are given to human annotators to establish correspondence with the template. Our method can be seen as automatizing the annotation step. Just like for manual annotation, it does *not* require a photorealistic rendition of the template. However, it does not mean that *all renditions are equally good* from the viewpoint of the matching network. Inspired by previous work on image generation [6], we find that rendering a *normal map* of the 3D template results in better matches than rendering a shaded version of the same (see the embedded figure). We discuss realistic rendering in Sec. 3.5.



### 3.4 Unsupervised canonical surface maps

Here, we show how to learn the canonical surface map  $f$  from the image-to-template correspondences constructed in Sec. 3.3. An overview is in Fig. 4.


*Continuous Surface Embeddings.* Following [24, 25], we represent the map  $f$  via Continuous Surface Embeddings (CSEs). CSE assign embedding vectors  $e_I(u), e(x_k) \in \mathbb{R}^D$  to each image pixel  $u$  and each mesh vertex  $x_k$  so that the correspondences are defined by maximizing their similarity:

$$f_I(u) = \operatorname{argmax}_{x_k \in V} p(x_k|u, I), \quad \text{where } p(x_k|u, I) = \frac{\exp(\langle e_I(u), e(x_k) \rangle)}{\sum_{t=1}^K \exp(\langle e_I(u), e(x_t) \rangle)}. \quad (3)$$

Learning the CSE model thus amounts to learning the vertex embeddings  $e(x_k)$  as well as a corresponding dense feature extractor  $e_I$ .

The vertex embeddings are optimized directly as there is a single template mesh. However, due to the large number of vertices, they are not assumed to be independent but to form a smooth (vector) function over the mesh surface. This way, the number of parameters required to express them can be reduced significantly. Collectively, all embeddings  $e(x_k)$ ,  $k = 1, \dots, K$ , form an embedding matrix  $E \in \mathbb{R}^{K \times D}$ . The latter is decomposed as the product  $E = UC$  where  $U \in \mathbb{R}^{K \times Q}$  is a smooth and compact functional basis (akin to Fourier components defined on the mesh) such that  $Q \ll K$ . Following [24, 25], we use the lowest eigenvectors of the Laplace-Beltrami operator (LBO) of the mesh  $M$  to form  $U$ . The only learnable parameters are  $C \in \mathbb{R}^{Q \times D}$ , which are few.

The other component is the feature extractor  $e_I(u)$ . For this, we encode the source image  $I \in \mathbb{R}^{3 \times H \times W}$  with a frozen self-supervised encoder (DINO [28]), before decoding it with a CNN back to the original resolution to the required feature tensor  $e_I \in \mathbb{R}^{D \times H \times W}$ .




**Fig. 4: CSE dense pose predictor.** We jointly train a deep network  $\Phi$  and a matrix  $C$ , that transforms LBO eigenvectors to a shared  $D$ -dimensional space. We use pseudo-ground truth, obtained as described in Sec. 3.3 for supervision. The image encoder is a frozen pre-trained DINO ViT, and the decoder we learn is a CNN.

*Training formulation.* We train our model with several losses. The first one simply uses the pseudo-ground-truth correspondences  $\tilde{x}(u, I)$  of Eq. (2) in Eq. (3). We pose this as a classification problem, where given a query location  $u$  in image  $I$  is matched probabilistically to the pseudo-ground-truth  $\tilde{x}(u, I)$  using the cross-entropy loss:  $\mathcal{L}_{\text{pseudo}}(I) = -\frac{1}{|U_I|} \sum_{u \in U_I} \log p(\tilde{x}(u, I)|u, I)$ . Additionally, we use the distance-aware loss of [25]:  $\mathcal{L}_{\text{dist}}(I) = -\frac{1}{|U_I|} \sum_{u \in U_I} \sum_{x \in V} d(x, \tilde{x}) p(x|u, I)$ , where  $d(x, \tilde{x})$  is the geodesic distance between the vertex  $x$  and the pseudo ground-truth  $\tilde{x}$ , which discourages placing probability mass far from  $\tilde{x}$ .

As noted in Sec. 3.3, there is some ambiguity because of the symmetry of most animals, where the matches are confused between left and right. To reduce this ambiguity, we use a cycle consistency loss [24], where given a starting location  $u$ , we match it to a vertex  $x_k$  in the mesh, and then matching that back to the image, results in the probability  $p(v|u, I) = \sum_{x_k \in V} p(v|x_k, I)p(x_k|u, I)$  of landing to a location  $v$ . Here  $p(v|x_k, I)$  is the same as  $p(x_k|v, I)$  from Eq. (3) up to renormalization. We close the image-shape-image cycle and supervise using  $\mathcal{L}_{\text{cyc}}(I) = \sum_{u \in U_I} \sum_{v \in U_I} \|u - v\| p(v|u)$ .

To further reduce the left-right ambiguity, we assume that the template  $V$  has a bilateral symmetry (true for most categories). Then, for each vertex  $x \in V$ , let  $x_F \in V$  be its symmetric one (for meshes which are not exactly symmetric, we let  $x_F$  be the closest approximation to the symmetric version of  $x$ ). Given an image  $I$  and a pixel  $u$ , denote by  $I_F$  and  $u_F$  their horizontal flips. Suppose that  $u$  is the pixel that corresponds to vertex  $x$  in image  $I$ . Then one can show [39] that pixel  $u_F$  must correspond to vertex  $x_F$  in image  $I_F$ , leading to the loss:  $\mathcal{L}_{\text{eq}}(I) = \frac{1}{|U_I|} \sum_{u \in U_I} \sum_{x \in V} |p(x|u, I) - p(x_F|u_F, I_F)|$ .





**Fig. 5: Realistic rendering of the template.** We create synthetic data for pixel-vertex correspondences by generating photorealistic images from depth renders. The corresponding vertices we obtain from the projections of vertices on the image.

### 3.5 Increasing the realism of synthetic data

The renders  $J_i$  of the 3D template  $M$  can be used to supervise the canonical map  $f$  directly because we *know* the 2D location  $v_i(k)$  of each vertex  $x_k$  in image  $J_i$  based on Eq. (1). With this, we can write the loss:

$$\mathcal{L}_{\text{syn}}(J_i, v_i) = -\frac{1}{K} \sum_{k=1}^K \sum_{i:V_k \in i} \log p(x_k | v_i(k), J_i) \quad (4)$$

The very limited diversity and realism of the renders  $J_i$  makes this loss uninteresting, but, as shown in Fig. 5, we can use a powerful image generator to significantly augment the realism of such renders.

To do this, we first render a depth image of the template  $M$  from a random viewpoint  $c$ . We also sample a random background image and predict its depth using [48], blend the foreground and background depth images, and use the depth-to-image ControlNet [52] of [48] to generate photorealistic image  $J_i$  of the template. We prompt the depth-to-image model (i) using the object’s class name, e.g., “horse”, and (ii) specifying the viewpoint (“front”, “side”, or “back”), which we heuristically obtain from the camera location w.r.t. the 3D template  $M$ .

The results are photo-realistic renders  $J_i$  of the template. Note that the appearance of different renders is not consistent, but this is a feature rather than an issue in our case because we need to learn an image-to-template map, which is invariant to details of the appearance. The main limitation is that the template is fixed, so there is no diversity in terms of pose and 3D shape. Hence, we expect loss Eq. (4) to be complementary rather than substitutive of the one above.

### 3.6 Learning formulation

Given a dataset  $\mathcal{D}$  of masked training images of the object, our loss is:

$$\mathcal{L} = \frac{1}{|\mathcal{D}|} \sum_{I \in \mathcal{D}} (\alpha \mathcal{L}_{\text{pseudo}}(I) + \beta \mathcal{L}_{\text{cyc}}(I) + \gamma \mathcal{L}_{\text{dist}}(I) + \delta \mathcal{L}_{\text{eq}}(I)) + \zeta \sum_{i=1}^N \mathcal{L}_{\text{syn}}(J_i, v_i),$$

where  $\alpha, \beta, \gamma, \delta$  and  $\zeta$  are coefficients set empirically.

Method	Supervision	Horse	Sheep	Bear	Zebra	Cow	Elephant	Giraffe	Average
CSE [25]	S	24.1	32.0	35.7	24.9	25.4	26.1	<b>18.0</b>	26.6
Zero-shot <sub>SD-DINO</sub>	U	37.2	41.4	48.2	32.0	32.3	36.0	26.3	36.2
SHIC (Ours)	U	<b>23.3</b>	<b>30.3</b>	<b>33.0</b>	<b>23.3</b>	<b>22.7</b>	<b>23.9</b>	18.1	<b>24.9</b>

**Table 1: Evaluation on DensePose-LVIS.** We compare the supervised (S) method of [25] to our unsupervised method (U) and our adaptation of SD-DINO to DensePose described in Sec. 3.3. We evaluate [25] using their published weights. We measure geodesic error (*lower is better*).

Method	Supervision	Cow	Sheep	Horse	Average
Rigid-CSM [18]	S	28.5	31.5	42.1	34.0
A-CSM [17]	S	29.2	39.0	44.6	37.6
CSE [25]	S	51.5	46.3	59.2	52.3
Rigid-CSM [18]	U	26.3	24.7	31.2	27.4
A-CSM [17]	U	26.3	28.6	32.9	29.3
SHIC (Ours) <sub>im2im</sub>	U	69.1	55.9	58.7	61.2
SHIC (Ours) <sub>im2m2im</sub>	U	<b>73.5</b>	<b>73.5</b>	<b>63.1</b>	<b>70.0</b>

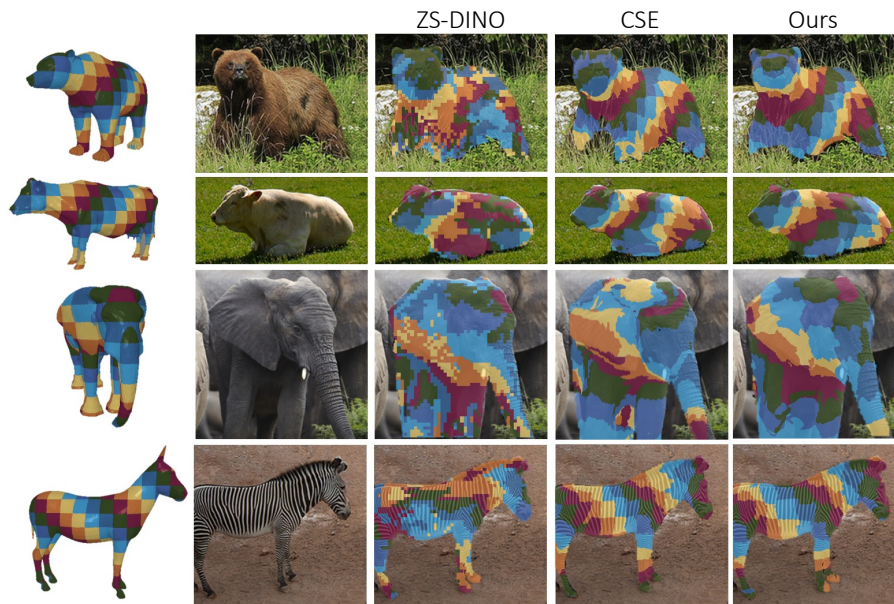
**Table 2: PCK-Transfer on PF-Pascal.** We compare against prior work on image-to-image semantic correspondences. We predict image-to-image correspondences either by directly predicting the correspondences or by performing image-to-vertex-to-image matching. We use the reported numbers from [17, 18, 25], and evaluate using PCK-0.1.

## 4 Experiments

We evaluate our method for learning canonical maps automatically, without manual keypoints supervision, against unsupervised and supervised prior work.

### 4.1 Implementation details

To obtain the image-to-shape similarities  $\Sigma$ , we use  $N = 72$  renderings of the template shape (using the surface normal style), render the surface normals, and compute the image-to-image similarities  $S$  using the features from [50]. For each rendered image, we automatically get the pixel-to-vertex matches from the camera projection function. Finally, we aggregate the similarities across all views using  $\max$  pooling. During training, we randomly select 100 foreground points from each image and their corresponding vertices from  $\Sigma$ . Following [25], we use the lowest  $Q = 64$  eigenvectors of the LBO of the template mesh and use dimensionality  $D = 16$  for the joint image-shape embedding space. For the image encoder  $\Phi$ , we use a frozen DINO-v2 [28] backbone, followed by a decoder consisting of 5 convolutional layers. We did not do any tuning of the parameters in the final formulation of the loss. We used values for the loss hyper-parameters that make the losses roughly of similar scale, namely  $\alpha = 0.1, \beta = 0.002, \gamma = 0.002, \delta = 0.001, \zeta = 0.1$ . All models and code will be released upon acceptance of the paper.



**Fig. 6: Mapping a textured mesh.** We map a textured mesh over the image using the predicted dense correspondences.

## 4.2 Training data

For most of our evaluations, we consider the DensePose-LVIS dataset [25], which applies DensePose to a variety of animal classes. Of those, we consider the horse, sheep, bear, zebra, cow, elephant and giraffe classes<sup>1</sup>. The DensePose-LVIS data contains a total of 6k images of these categories, as well as a reference 3D template for each category. Knowledge of the 3D template is necessary to interpret the annotations in the dataset, as well as to compute the geodesic distances required for evaluation. Every animal has up to three manually annotated pixel-template correspondences to the corresponding template. We use these annotations only for evaluation. We only train our models on cropped animals from DensePose-LVIS [25]. The number of instances for each class varies from 2,899 for horses (most) to 735 for bears (least). In comparison, competing methods train on more images. [17] train on combined PASCAL and ImageNet images, and the supervised method of [25] is trained on DensePose-LVIS and DensePose-COCO, the latter of which consists of 5 million image-to-template annotations for humans. Likely due to the density of the pseudo-ground truth, our method only needs a much smaller amount of data and can be trained with as few as a few hundred

<sup>1</sup> For the classes cat and dog, we could not obtain the 3D templates from [25] and could therefore not use the annotated image-to-template correspondences for evaluation.

Ablation	DensePose-LVIS
Ours	<b>24.9</b>
Pooling <b>mean</b> (instead of <b>max</b> )	26.6
Data w/o Synth data	25.6
w/o LVIS	33.8
w/o pseudo GT	65.8

**Table 3: Data and other ablations.** First, we ablate the pooling function used to construct  $\Sigma$ , and evaluate **mean** instead of **max** pooling. Next, we compare the data we use — we remove (1) synthetic data, (2) natural images, (3) the pseudo ground truth for natural images. In (3), we still use natural images for  $\mathcal{L}_{\text{eq}}$  and  $\mathcal{L}_{\text{cyc}}$ , but do not use the pseudo-GT for  $\mathcal{L}_{\text{pseudo}}$  and  $\mathcal{L}_{\text{dist}}$ .

Ablation	DensePose-LVIS
Ours	<b>24.9</b>
Losses w/o $\mathcal{L}_{\text{pseudo}}$	27.3
w/o $\mathcal{L}_{\text{dist}}$	25.8
w/o $\mathcal{L}_{\text{eq}}$	25.6
w/o $\mathcal{L}_{\text{cyc}}$	25.8
w/o $\mathcal{L}_{\text{eq}}$ & $\mathcal{L}_{\text{cyc}}$	26.1

**Table 4: Ablating the losses.** We assess the contributions of the losses removing them one at a time. On the last row, we remove  $\mathcal{L}_{\text{eq}}$  and  $\mathcal{L}_{\text{cyc}}$  to show that although both address symmetry, they are complementary. Ablating  $\mathcal{L}_{\text{syn}}$  falls under Tab. 3 (w/o Synth data) as it uses different data. We show the average score over all classes.


images (*e.g.*, bear class from DensePose-LVIS, or the models we show in Fig. 1). Similarly to CSM [17, 18] and CSE [25], we use masks for training.

### 4.3 Evaluation

We evaluate our models and CSE [25] on DensePose-LVIS using the geodesic error. We normalize the maximum geodesic distance on each mesh to 228, following [10, 24], and use a heat solver to obtain all vertex-to-vertex geodesic distances. Additionally, we use PF-PASCAL [12] to evaluate our model on keypoint transfer. PF-PASCAL consists of pairs of images with annotated salient keypoints (*e.g.*, left eye, nose, *etc.*), and we evaluate image-to-mesh-to-image correspondences using the image-to-image annotations. We use the test split of [50] and evaluate using PCK 0.1, following prior work [17, 25, 50].

*Image-to-shape correspondences.* In Tab. 1 we compare the quality of the image-to-template correspondences established by SHIC, by our zero-shot method based on SD-DINO, and by CSE [25], which is supervised, on the DensePose-LVIS dataset. The most important result is that SHIC learns better canonical maps than CSE despite using no supervision. While the training images of SHIC and CSE are the same and the pseudo-ground truth is noisy ( $\text{Zero-shot}_{\text{SD-DINO}}$  is what we use for supervision), this result can be explained by the fact that our automated supervision is significantly denser than the manual labels collected by [25] (just three per image).

We show some qualitative results on this dataset in Fig. 6, where we color every point on the image according to the corresponding color on the mesh. The regularity of the remapped texture illustrates the quality of the correspondences. Once more, the learned canonical map (ours) is significantly better than the pseudo-ground truth (SD-DINO). Compared to CSE, SHIC performs similarly



**Fig. 7: Image-to-image correspondences.** We show image-to-image correspondences on PF-PASCAL, which we find using pixel-to-vertex-to-pixel matching. The heatmaps on the shape show the similarity from the source image location to every vertex.

and tends to have a more regular structure on the heads. We show more qualitative evaluations and failure cases in the Appendix.

*Ablation study.* We ablate the components of our method in Tab. 3 and Tab. 4. When using mean pooling instead of max pooling for obtaining the pseudo-ground-truths, we see a significant drop in performance. This is because image-to-image correspondences are more reliable when the objects are in similar poses, and with max pooling we only get contribution from the most similar view. Next, we look at the importance of different sources of supervision. When we remove the pseudo-ground-truth from our synthetic pipeline (Sec. 3.5), the model loses performance. When we exclude all natural images (w/o LVIS), and only train on synthetically generated data, we see a more pronounced drop in performance. Finally, we exclude the pseudo-ground-truth from SD-DINO (Sec. 3.3) and thus only train with the synthetic data and the cycle consistency and equivariance losses on natural images. In this case, the model learns a degenerate solution, where for natural images it only predicts vertices on one side of the shape (*i.e.*, left or right). Such degenerate solutions have been observed by [42] for unsupervised image-to-image matching when using cycle consistency losses. This shows the importance of using our pseudo-ground-truth. Finally, we ablate the losses we use in Tab. 4. All losses are necessary for the final performance. We find that  $\mathcal{L}_{\text{pseudo}}$ , where we frame pixel-to-vertex matching as a multi-class classification problem, has a bigger contribution than the distance-aware loss  $\mathcal{L}_{\text{dist}}$  of [25]. Additionally, we see that both losses that address symmetry,  $\mathcal{L}_{\text{eq}}$  and  $\mathcal{L}_{\text{cyc}}$ , improve performance, and are complementary to each other.

*Keypoint transfer.* Next, we evaluate SHIC on the PF-PASCAL [56] in Tab. 2. In this dataset, one evaluates the quality of image-to-image correspondences

instead of image-to-template. There are two ways of using our method to induce image-to-image correspondences. The first is to use the learned canonical maps, transferring points from one image to the template and then back to the other image. The second is to directly match the image-based CSE embedding learned in Sec. 3.4. The key findings from our results are: (1) SHIC outperforms the supervised versions of Rigid [18] and Articulated [17] CSM, as well as CSE [25] and greatly outperforms all unsupervised approaches. (2) The image-to-image correspondences induced by the canonical maps are significantly better than the ones induced by the image-based CSE embeddings, once again illustrating the importance of the canonical maps. We show qualitative examples in Fig. 7.

*Novel classes.* SHIC can be trained on any class as long as there is a collection of a few hundred images and a suitable template mesh. In Fig. 1 we show qualitative results from a model we train on two classes — T-Rex and Appa, a six-legged flying bison from Avatar: The Last Airbender, using 480 and 180 manually collected images, respectively. We extract masks for training using the open vocabulary segmentation method of [20]. Although the 3D template for Appa is toy-like and does not resemble the images closely, and we only use 180 images, SHIC still manages to learn useful correspondences. This is a considerable advantage of our method over supervised previous work, as it can be trained from a small number of images and without human supervision. This allows the construction of general-purpose shape correspondence models for almost any category.

## 5 Conclusion

We have introduced an unsupervised method to learn correspondence matching between a 3D template and images. Critically, this model can be trained without supervision and from less than 200 images, which makes it applicable to a vast number of objects. This is a significant step beyond previous work that required lots of manually labelled correspondences. We hope that SHIC will enable many downstream tasks where learnt robust correspondence estimation was previously impossible.

*Ethics.* We utilize the DensePose-LVIS dataset [24] and PF-PASCAL [11] for evaluation in a manner compatible with their terms. The images may contain humans, but we only consider occurrences of animals and there is no processing of biometric information. For further details on ethics, data protection, and copyright please see <https://www.robots.ox.ac.uk/~vedaldi/research/union/ethics.html>.

*Acknowledgements.* We thank Orest Kupyn for helpful discussions and Luke Melas-Kyriazi for proofreading this paper. A. Shtedritski is supported by EPSRC EP/S024050/1. A. Vedaldi is supported by ERC-CoG UNION 101001212.

## References

1. Amir, S., Gandelsman, Y., Bagon, S., Dekel, T.: Deep ViT features as dense visual descriptors. CoRR **abs/2112.05814** (2021)
2. Bourdev, L.D., Malik, J.: Poselets: Body part detectors trained using 3D human pose annotations. In: Proc. ICCV (2009)
3. Cao, Z., Simon, T., Wei, S., Sheikh, Y.: Realtime multi-person 2D pose estimation using part affinity fields. In: Proc. CVPR (2017)
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9650–9660 (2021)
5. Chen, J., Wang, L., Li, X., Fang, Y.: Arbicon-net: Arbitrary continuous geometric transformation networks for image registration. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 32. Curran Associates, Inc. (2019)
6. Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. arXiv.cs **abs/2303.13873** (2023)
7. Dutt, N.S., Muralikrishnan, S., Mitra, N.J.: Diffusion 3d features (diff3f): Decorating untextured shapes with distilled semantic features. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4494–4504 (June 2024)
8. Felzenszwalb, P.F., McAllester, D.A., Ramanan, D.: A discriminatively trained, multiscale, deformable part model. In: Proc. CVPR (2008)
9. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7297–7306 (2018)
10. Güler, R.A., Neverova, N., Kokkinos, I.: DensePose: Dense human pose estimation in the wild. In: Proc. CVPR (2018)
11. Ham, B., Cho, M., Schmid, C., Ponce, J.: Proposal flow. In: Proc. CVPR (2016)
12. Ham, B., Cho, M., Schmid, C., Ponce, J.: Proposal flow: Semantic correspondences from object proposals. IEEE transactions on pattern analysis and machine intelligence **40**(7), 1711–1725 (2017)
13. Hedlin, E., Sharma, G., Mahajan, S., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: Unsupervised semantic correspondence using stable diffusion. arXiv.cs (2023)
14. Jeon, S., Kim, S., Min, D., Sohn, K.: Parn: Pyramidal affine regression networks for dense semantic correspondence. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 351–366 (2018)
15. Kanazawa, A., Jacobs, D.W., Chandraker, M.: WarpNet: Weakly supervised matching for single-view reconstruction. In: Proc. CVPR (2016)
16. Kreiss, S., Bertoni, L., Alahi, A.: Pifpaf: Composite fields for human pose estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11977–11986 (2019)
17. Kulkarni, N., Gupta, A., Fouhey, D.F., Tulsiani, S.: Articulation-aware canonical surface mapping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 452–461 (2020)
18. Kulkarni, N., Gupta, A., Tulsiani, S.: Canonical surface mapping via geometric cycle consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2202–2211 (2019)

19. Li, X., Lu, J., Han, K., Prisacariu, V.: Sd4match: Learning to prompt stable diffusion model for semantic matching. arXiv preprint arXiv:2310.17569 (2023)
20. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023)
21. Luo, G., Dunlap, L., Park, D.H., Holynski, A., Darrell, T.: Diffusion hyperfeatures: Searching through time and space for semantic correspondence. In: Advances in Neural Information Processing Systems (2023)
22. Melekhov, I., Tiulpin, A., Sattler, T., Pollefeys, M., Rahtu, E., Kannala, J.: Dgc-net: Dense geometric correspondence network. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1034–1042. IEEE (2019)
23. Morreale, L., Aigerman, N., Kim, V.G., Mitra, N.J.: Neural semantic surface maps. In: Computer Graphics Forum. vol. 43, p. e15005. Wiley Online Library (2024)
24. Neverova, N., Novotny, D., Szafraniec, M., Khalidov, V., Labatut, P., Vedaldi, A.: Continuous surface embeddings. Advances in Neural Information Processing Systems **33**, 17258–17270 (2020)
25. Neverova, N., Sanakoyeu, A., Labatut, P., Novotny, D., Vedaldi, A.: Discovering relationships between object categories via universal canonical maps. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 404–413 (2021)
26. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: Proc. ECCV (2016)
27. OpenAI: Chatgpt. <https://chat.openai.com/>
28. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
29. Peebles, W., Zhu, J.Y., Zhang, R., Torralba, A., Efros, A.A., Shechtman, E.: Gansupervised dense visual alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13470–13481 (2022)
30. Pereira, T., Aldarondo, D.E., Willmore, L., Kislin, M., Wang, S.S.H., Murthy, M., Shaevitz, J.W.: Fast animal pose estimation using deep neural networks. bioRxiv (2018)
31. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
32. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 (2022)
33. Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., Guibas, L.J.: Humor: 3d human motion model for robust pose estimation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 11488–11499 (2021)
34. Rocco, I., Arandjelovic, R., Sivic, J.: Convolutional neural network architecture for geometric matching. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6148–6157 (2017)
35. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2021)
36. Shtedritski, A., Rupprecht, C., Vedaldi, A.: Learning universal semantic correspondences with no supervision and automatic data curation. In: Proc. IEEE/CVF International Conference on Computer Vision (ICCV) Workshops (2023)



37. Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023)
38. Thewlis, J., Bilen, H., Vedaldi, A.: Unsupervised learning of object frames by dense equivariant image labelling. In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)* (2017)
39. Thewlis, J., Bilen, H., Vedaldi, A.: Modelling and unsupervised learning of symmetric deformable object categories. In: *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)* (2018)
40. Truong, P., Danelljan, M., Gool, L.V., Timofte, R.: Gocor: Bringing globally optimized correspondence volumes into your neural network. *Advances in Neural Information Processing Systems* **33**, 14278–14290 (2020)
41. Truong, P., Danelljan, M., Timofte, R.: Glu-net: Global-local universal network for dense flow and correspondences. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6258–6268 (2020)
42. Truong, P., Danelljan, M., Yu, F., Van Gool, L.: Warp consistency for unsupervised learning of dense correspondences. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10346–10356 (2021)
43. Truong, P., Danelljan, M., Yu, F., Van Gool, L.: Probabilistic warp consistency for weakly-supervised semantic correspondences. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8708–8718 (2022)
44. Waldmann, U., Chan, A.H.H., Naik, H., Nagy, M., Couzin, I.D., Deussen, O., Goldluecke, B., Kano, F.: 3d-muppet: 3d multi-pigeon pose estimation and tracking. *arXiv preprint arXiv:2308.15316* (2023)
45. Wei, S., Ramakrishna, V., Kanade, T., Sheikh, Y.: Convolutional pose machines. In: *Proc. CVPR* (2016)
46. Wu, S., Jakab, T., Rupprecht, C., Vedaldi, A.: DOVE: Learning deformable 3D objects by watching videos. In: *arXiv* (2021)
47. Wu, S., Li, R., Jakab, T., Rupprecht, C., Vedaldi, A.: MagicPony: Learning articulated 3D animals in the wild. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2023)
48. Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. *arXiv:2401.10891* (2024)
49. Zhang, H., Tian, Y., Zhou, X., Ouyang, W., Liu, Y., Wang, L., Sun, Z.: Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11446–11456 (2021)
50. Zhang, J., Herrmann, C., Hur, J., Cabrera, L.P., Jampani, V., Sun, D., Yang, M.H.: A Tale of Two Features: Stable Diffusion Complements DINO for Zero-Shot Semantic Correspondence. *arXiv preprint arxiv:2305.15347* (2023)
51. Zhang, J., Herrmann, C., Hur, J., Chen, E., Jampani, V., Sun, D., Yang, M.H.: Telling left from right: Identifying geometry-aware semantic correspondence. *arXiv.cs* (2023)
52. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models (2023)
53. Zhang, N., Donahue, J., Girshick, R.B., Darrell, T.: Part-based R-CNNs for fine-grained category detection. In: *Proc. ECCV* (2014)
54. Zuffi, S., Kanazawa, A., Berger-Wolf, T., Black, M.J.: Three-d safari: Learning to estimate zebra pose, shape, and texture from images" in the wild". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5359–5368 (2019)

55. Zuffi, S., Kanazawa, A., Black, M.J.: Lions and tigers and bears: Capturing non-rigid, 3d, articulated shape from images. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3955–3963 (2018)
56. Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3D menagerie: Modeling the 3D shape and pose of animals. In: Proc. CVPR (2017)