

Canonical Surface Mapping via Geometric Cycle Consistency

Nilesh Kulkarni

Abhinav Gupta*

Shubham Tulsiani*

Carnegie Mellon University

Facebook AI Research

{nileshk, abhinavg}@cs.cmu.edu

shubhtuls@fb.com


<https://nileshkulkarni.github.io/csm/>

Figure 1: We study the task of Canonical Surface Mapping (CSM). This task is a generalization of keypoint estimation and involves mapping pixels to canonical 3D models. We learn CSM prediction without requiring correspondence annotations, by instead using geometric cycle consistency as supervision. This allows us to train CSM prediction for diverse classes, including rigid and non-rigid objects.

Abstract

We explore the task of Canonical Surface Mapping (CSM). Specifically, given an image, we learn to map pixels on the object to their corresponding locations on an abstract 3D model of the category. But how do we learn such a mapping? A supervised approach would require extensive manual labeling which is not scalable beyond a few hand-picked categories. Our key insight is that the CSM task (pixel to 3D), when combined with 3D projection (3D to pixel), completes a cycle. Hence, we can exploit a geometric cycle consistency loss, thereby allowing us to forgo the dense manual supervision. Our approach allows us to train a CSM model for a diverse set of classes, without sparse or dense keypoint annotation, by leveraging only foreground mask labels for training. We show that our predictions also allow us to infer dense correspondence between two images, and compare the performance of our approach against several methods that predict correspondence by leveraging varying amount of supervision.

1. Introduction

Plato famously remarked that while there are many cups in the world, there is only one ‘idea’ of a cup. Any particular instance of a category can thus be understood via its relationship to this platonic ideal. As an illustration, consider an image of a bird in Figure 1. When we humans see this image, we can not only identify and segment the bird but also go further and even map pixels to an abstract 3D representation of the category. This task of mapping pixels in an image to locations on an abstract 3D model (which we henceforth call **canonical surface mapping**) is generalization and densification of keypoint estimation and is key towards rich understanding of objects. But how do we learn to do this task? What is the right data, supervision or models to achieve dense rich understanding of objects?

One way to learn the canonical surface mapping task is to collect large-scale labeled data. Specifically, we can label hundreds or thousands of keypoints per image for thousands

* the last two authors were equally uninvolved.

of images. As each keypoint location defines which pixel corresponds to a specific location on the 3D surface, this approach of manually labeling the keypoints can provide dense supervision for learning. This approach has in fact been shown to be quite successful for specific categories such as humans [2]. But of course collecting such labeled data requires enormous manual labeling effort, making it difficult to scale to generic categories.

Is there an alternative supervisory signal that can allow one to learn without reliance on such labelled data? Interestingly, we note that this task of canonical surface mapping is an inverse graphics task. Any such mapping is constrained by the geometry operating on the underlying 3D, and any predicted mapping should also respect this structure. In particular, for the pixels that belong to the object given by the object mask, the CSM function maps these pixels onto the 3D shape. These points on the 3D shape, when projected back using (known/predicted) camera, should map back to the same pixels. Our key insight is that one can complete the cycle (pixels \rightarrow 3D \rightarrow pixels) and use the consistency loss as an objective. The gradients from the loss can be propagated back to the CSM function prediction function, thereby allowing us to learn this mapping without reliance on strong forms of supervision.

In this paper, we present an approach to learn the task of canonical surface mapping from the set of images belonging to semantic category, their input masks and an abstract 3D model which represents the semantic category. Additionally, we show that predicting a canonical surface mapping for images allows us to infer dense correspondence across images of a category, and our approach enables recovery of dense correspondences without any correspondence supervision! In comparison to approaches that use dense supervision for this task [2], or approaches that leverage keypoints for the related tasks of semantic correspondence [7], or 3D reconstruction [18], this is significant decrease in supervision. This allows us to train our CSM model for a diverse set of classes: birds, zebras, cars and more (See Figure 1). We believe our approach can pave the way for large-scale internet-driven 3D understanding and correspondence inference since both semantic imagesets and masks are easy to obtain (and automatic approaches can be used as well).

2. Related Work

Dense Semantic Correspondences. A fundamental task that is equivalent to pursuing canonical surface mapping is that of inferring dense semantic correspondence – given two images, the goal is to predict for each pixel in the former, the corresponding pixel in the latter. Methods prior to the recent resurgence of deep learning [22, 24] demonstrated that matching using features such as SIFT could allow recovering correspondence across instances, and later work showed similar results using CNN features [13, 25]. While these

generic features allow recovering correspondence, learning specifically for the task using annotated data can improve results [7]. However, collecting such annotation can be tedious, so several approaches have attempted to relax the supervision for learning correspondence.

Among these, a common paradigm is to learn correspondence by self-supervision, where random perturbations of images are used as training pairs. This allows predicting parametric warping [17, 31, 32] to relate images, or learn equivariant embeddings [38] for matching. However, these methods are fundamentally restricted to training data of the same instance, with no change in the visible content, thereby limiting the performance for different instances with viewpoint changes. While for certain categories of interest *e.g.* humans, some approaches [27, 30, 35, 36, 42] show that it is possible to use calibrated multi-view or motion capture to generate supervision, this form of supervision is slightly tedious to collect for all classes. An alternate form of supervision can come via synthetic data, where synthetic image pairs rendered using the same pose as a real image pair, can help learn a correspondence function between real images that is cycle-consistent [52]. However, this approach relies on availability of large-scale synthetic data and known pose for real images to generate the supervisory signal, and we show that both these requirements can be relaxed.

Learning Invariant Representations. Our work is broadly related to methods that learn pixel embeddings invariant to certain transforms. These approaches leverage tracking to obtain correspondence labels, and learn representations invariant to viewpoint transformation [34, 49] or motion [44]. Similar to self-supervised correspondence approaches, these are also limited to training using observations of the same instance, and do not generalize well across instances. While our canonical surface mapping is also a pixel-wise embedding invariant to certain transforms, it has a specific geometric meaning *i.e.* correspondence to a 3D surface, and leveraging this is what allows learning without the correspondence supervision.

Category-Specific 3D Reconstruction. A related line of work pursued in the community is that of reconstructing the instances in a category using category-specific deformable models. Dating back to the seminal work of Blanz & Vetter [4], who operationalized D’Arcy Thompson’s insights into the manifold of forms [39], morphable 3D models have been used to model faces [4], hands [21, 37], humans [3, 26] and other generic classes [5, 18, 19, 46].

In conjunction with known/predicted camera parameters, this representation also allows one to extract a pixelwise canonical mapping. However, these methods often rely on 3D training data to infer this representation. Even approaches that relax this supervision [18, 19, 46] crucially

rely on (sparse or dense) 2D keypoint annotations during training. In contrast, we show that learning a canonical surface mapping is feasible even without such supervision. Further, we demonstrate that directly learning the mapping function leads to more accurate results than obtaining these via an intermediate 3D estimate.

Consistency as Meta-Supervision. Ours is not the only task where acquiring direct supervision is often infeasible, and the idea of leveraging some form of consistency to overcome this hurdle has been explored in several domains. Recent volumetric reconstruction [12, 29, 41, 48] or depth prediction [10, 11, 50] approaches use geometric consistency between the predicted 3D and available views as supervision. Similarly, the notion that when learning some transformations, their composition often respects a cyclical structure has been used for image generation [23, 53], correspondence estimation [50, 51] *etc.* In our setup, we also observe that the approach of using consistency as meta-supervision allows bypassing supervision. We do so by leveraging insights related to both, geometry and cycle consistency – given a surface mapping, there is a geometrically defined inverse transform with which the canonical surface mapping predictions should be cycle-consistent.

3. Approach

Given an image, our goal is to infer for each pixel on the object, its mapping onto a given canonical template shape of the category. We do so by learning a parametrized CNN f_θ , which predicts a pixelwise canonical surface mapping (CSM) given an input image. We show that our method, while only relying on foreground masks as supervision, can learn to map pixels to the given category-level template shape. Our key insight is that this mapping function we aim to learn has a geometric structure that should be respected by the predictions. We operationalize this insight, and learn a CSM predictor using a geometric cycle consistency loss, thereby allowing us to bypass the need for supervision in the form of annotated (sparse or dense) keypoints.

We first present in Section 3.2 our training setup in a scenario where the camera pose for each training image is given. We then show how we can relax this requirement of known camera in Section 3.3. Learning a CSM predictor implicitly allows us to capture the correspondence across instances, and we describe in Section 3.4 the procedure to recover dense semantic correspondence given two images.

3.1. Preliminaries

Surface Parametrization. The template shapes we learn mappings to are in fact two-dimensional surfaces in 3D space. The surface S of the template shape can therefore be parametrized via two parameters $u \in (0, 1)$ and $v \in (0, 1)$ (or equivalently a 2D vector \mathbf{u}). This parametrization im-

plies that we can obtain a mapping ϕ such that $\phi(\mathbf{u})$ represents a unique point on the surface S .

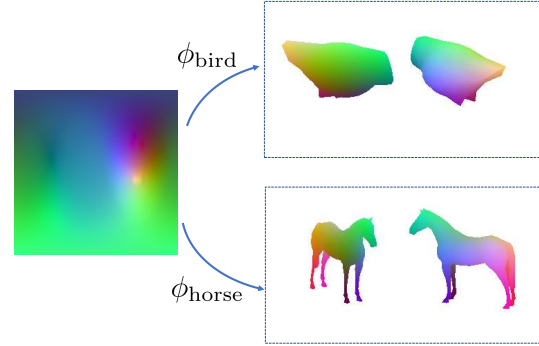


Figure 2: Surface Parametrization. We show the mapping from (u, v) space to the surface of the 3D model for two categories.

While there are several ways to construct such a mapping, one intuitive way is to consider \mathbf{u} to represent the polar angles to parametrize points on the surface of a hollow sphere, which can be mapped to a surface S by pushing it inward [28]. Given a template shape with a surface S , we use this approach to obtain the parametrization ϕ . We show some visualizations in Figure 2 for the mapping from a 2D square to template 3D shapes for two categories.

Canonical Surface Mapping. A canonical surface mapping C for an image I is a mapping from pixels onto the template 3D shape. Given a pixel $\mathbf{p} \equiv (x, y)$, $C[\mathbf{p}]$ represents the corresponding point on the surface. As the surface has a two-dimensional parametrization, C is equivalently an image of the same size as I , with a two-channel value at each pixel. Our parametrized CNN f_θ that predicts this mapping from an input image, therefore learns a per-pixel prediction task – given an RGB input image, it outputs a 2 dimensional vector for each pixel.

Camera Projection. We model the camera as a weak perspective (scaled orthographic) transformation. We represent the camera for every image I as π , parameterized by the scale $s \in \mathcal{R}$, translation $t \in \mathcal{R}^2$ and rotation r are three euler angles. We denote by $\pi(P)$ as the projection of a point P to the image coordinate frame using the camera parameters $\pi \equiv (s, t, r)$.

3.2. Learning via Geometric Cycle Consistency

We aim to learn a per-pixel predictor f_θ that outputs a canonical surface mapping given an input image I . We present an approach to do so using only foreground masks as supervision. However, for simplicity, we first describe here how we can learn this CSM predictor assuming known camera parameters for each training image, and relax this requirement in Section 3.3.

Our approach is to derive learning signal from the geometric




Figure 3: Geometric Cycle Consistency Loss. A pixel mapped to \mathbf{u} by CSM function f_θ gets mapped onto the 3D template via ϕ . Our loss enforces that this 3D point, when projected back under the camera π , should map back to the pixel.

nature of this task. In particular, as the 3D shapes underlying instances of a category are often similar (and therefore similar to the template shape), a pixel-wise mapping onto the 3D surface should be (approximately) cycle-consistent under reprojection. We capture this constraint via a geometric cycle consistency loss. This loss, in conjunction with an objective that allows the prediction to respect certain visibility constraints, allows us to learn f_θ .

Geometric Cycle Consistency Loss. Given an image I with associated camera π and foreground mask I_f , we wish to enforce that the predicted canonical surface mapping $C \equiv f_\theta(I)$, respects the underlying geometric structure. Concretely, as the instances across a category bear resemblance to the template shape, given a pixel \mathbf{p} on the object foreground, we would expect that its corresponding point on the 3D surface $\phi(C[\mathbf{p}])$ to (approximately) project back under the camera π which we denote as $\bar{\mathbf{p}}$. We define a geometric consistency loss (see Figure 3) that penalizes this inconsistency for all foreground pixels, thereby encouraging the network to learn pixel \rightarrow 3D mapping functions that are cycle-consistent under the 3D \rightarrow pixel reprojection.

$$L_{cyc} = \sum_{\mathbf{p} \in I_f} \|\bar{\mathbf{p}} - \mathbf{p}\|_2^2 \quad ; \quad \bar{\mathbf{p}} = \pi(\phi(C[\mathbf{p}])) \quad (1)$$

Incorporating Visibility Constraints. Enforcing that the pixels when lifted to 3D, project back to the same location is desirable, but not a sufficient condition. As an illustration, for a front facing bird, both the beak and tail project at similar locations, but only the former would be visible. This implies that points on the surface that are self-occluded under π can also result in minimizing L_{cyc} . Our solution is to discourage f_θ from predicting \mathbf{u} values that map to self-occluded regions under camera π .

A point on the 3D shape is self-occluded under a camera π , its z-coordinate in camera frame is larger than the rendered depth at the corresponding pixel. We use Neural Mesh Renderer (NMR) [20] to render a depth map D_π for the template shape S under camera π , and define a visibility loss for each

pixel \mathbf{p} by checking if the z-coordinate (say $z_{\mathbf{p}}$) of its corresponding point $\phi(C[\mathbf{p}])$ on the 3D shape, when projected under π , has a larger z-coordinate.

$$L_{vis} = \sum_{\mathbf{p} \in I_f} \max(0, z_{\mathbf{p}} - D_\pi[\bar{\mathbf{p}}]) \quad (2)$$

Network Details. We implement f_θ as a network with UNet [33] style architecture. This network takes as input an image of size 256 x 256 and outputs a unit vector per pixel representing a point on surface of sphere which is then converted to a (u, v) coordinate analogous to latitude and longitude. We train our network to minimize the cycle-consistency and visibility objectives:

$$L_{consistency} = L_{vis} + L_{cyc} \quad (3)$$

Even though we do not have direct supervision for the mappings, as we train a shared predictor across instances, the explicit priors for geometric consistency, and the implicit inductive biases in CNNs for spatial equivariance are sufficient for us to learn a meaningful predictor.

Foreground Mask Prediction. While the training procedure described above encourages cycle-consistent predictions at pixels belonging to the object, the learned CNN f_θ also predicts some (possibly spurious) values at other pixels. To allow us to ignore these background pixels for inferring correspondence (see Section. 3.4), as well as for generating visualizations, we train an additional per-pixel mask predictor using standard cross-entropy loss L_{fg} against the ground-truth masks. To do so, we simply modify f_θ to yield an additional per-pixel foreground probability as output.

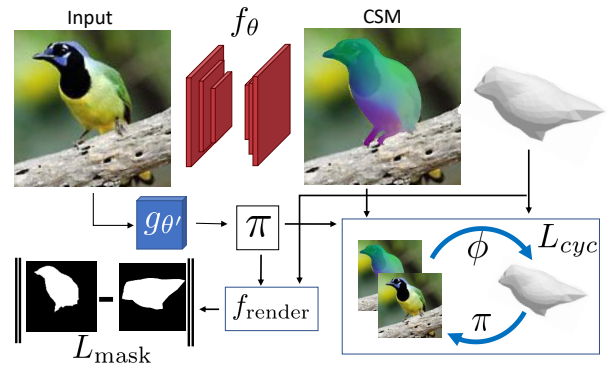


Figure 4: Overview of Training Procedure. We train a network to predict, for each pixel on the foreground, its mapping to the canonical shape. We also jointly learn to predict camera pose, and the geometric cycle-consistency loss L_{cyc} along with foreground supervision, provides learning signal to train our system.

3.3. Learning without Pose Supervision

We have presented our approach to learn a canonical surface mapping predictor f_θ assuming known cameras π for each training image. We note that our training objective is also differentiable w.r.t. the camera parameters, and we can therefore simply use predicted cameras instead of known cameras, and jointly learn pose and CSM prediction. This joint training can allow us to bypass the requirement of even camera supervision, and learn CSM prediction using only foreground mask annotations and a given template shape.

We therefore learn an additional camera-prediction CNN $g_{\theta'}$, and use the predicted cameras to learn the CSM predictor via the geometric consistency training objectives. However, to overcome certain trivial solutions, we also add a mask reprojection error, and following [16, 40] use a multi-hypothesis camera predictor to avoid local minima. Our overall training setup is depicted in Figure 4.

Mask Re-projection Loss. If the only learning objective comprises of the self-consistency between camera predictions and the predicted CSMs, the networks can learn some trivial solutions *e.g.* always predict a ‘frontal’ camera and corresponding CSM. To avoid this we enforce that the template shape, when viewed under a predicted camera π , should approximately match the known foreground image I_{fg} . To implement this loss, we use (NMR) [20] to obtain a differentiable render f_{render} , that given the template shape S and a camera π , renders a mask. While the poses may still be ambiguous *e.g.* front and back facing cars, this additional mask reprojection loss allows us to circumvent the mentioned trivial solutions. This reprojection loss is defined as follows:

$$L_{mask} = \|f_{render}(S, \pi) - I_f\|^2 \quad (4)$$

Multi-Hypothesis Pose Prediction. Instead of predicting a single camera $\pi \equiv g_{\theta'}(I)$, we follow previous methods [16, 40] and predict multiple hypotheses to overcome local minima. Our pose predictor outputs $\{(\pi_i, c_i)\} \equiv g_{\theta'}(I)$ - a set of $N_c = 8$ pose hypotheses π_i , each with an associated probability c_i . We initialize the camera predictor $g_{\theta'}$ using a pre-trained ResNet-18 network [15].

Overall Training Objective. As our pose predictor yields multiple pose hypotheses π_i , each with an associated probability c_i , we can train our networks by minimizing the expected loss. We denote by $L_{cyc}^i, L_{vis}^i, L_{mask}^i$ the corresponding losses under the camera prediction π_i . In addition to minimizing the expected loss over these terms, we also use an additional diversity prior L_{div} to encourage diverse pose hypotheses (see appendix for details). The overall training objective using these, is:

$$L_{tot} = L_{div}(g_{\theta'}(I)) + \sum_{i=1}^{N_c} c_i (L_{cyc}^i + L_{vis}^i + L_{mask}^i) \quad (5)$$

This framework allows us to learn the canonical surface mapping function f_θ via geometric cycle consistency, using only foreground mask annotations in addition to the given template shape. Once the network f_θ is learned, we can infer a canonical surface map from any unannotated image.

3.4. Dense Correspondences via CSM

We described an approach for predicting canonical surface mappings without relying on pose or keypoint annotations. This allows us to infer dense semantic correspondences given two images of the same semantic object category, because if pixels across images correspond, they should get mapped to the same region on the canonical surface. Given a (source, target) image pair (I_s, I_t) , let us denote by $(C_s, C_t, I_{fg}^s, I_{fg}^t)$ the corresponding predicted canonical surface mappings and foreground masks. Given these predictions, for any pixel \mathbf{p}_s on I_s , we can infer its corresponding pixel $T_{s \rightarrow t}[\mathbf{p}_s]$ on I_t by searching for the (foreground) pixel that maps closest to $\phi(C_s[\mathbf{p}_s])$.

$$T_{s \rightarrow t}[\mathbf{p}_s] = \arg \min_{\mathbf{p}_t \in I_{fg}^t} \|\phi(C_s[\mathbf{p}_s]) - \phi(C_t[\mathbf{p}_t])\| \quad (6)$$

Not only does our approach allow us to predict correspondences for pixels between two images, it *also allows us to infer regions of non-correspondence i.e.* pixels in source image for which correspondences in the target image do not exist (*e.g.* most pixels between a left and right facing bird do not correspond). We can infer these by simply denoting pixels for which the minimum distance in Eq. 6 is above a certain threshold as not having a correspondence in the target image. This ability to infer non-correspondence is particularly challenging for self-supervised methods that generate data via random warping [17, 31, 38] as the training pairs for these never have non-corresponding regions.

4. Experiments

Our approach allows us to predict canonical surface mappings across generic categories. However, due to lack of annotation for the task, which is in fact our motivation for learning without supervision, it is difficult to directly evaluate the predictions. Instead, as our approach also allows us to recover correspondences across any two images (Section 3.4), we can evaluate these using the task of keypoint transfer. This is a well-studied task by approaches that learn semantic correspondence, and we report comparisons to baselines that leverage varying degree of supervision while training. We first report these comparisons in Section 4.1, and then present results for additional generic categories (*e.g.* horses, sheep, cows) in Section 4.2, using Imagenet images with automatically obtained segmentation masks.

4.1. Evaluation via Keypoint Transfer

We use our learned CSM prediction models for the task of keypoint transfer – given a source and target image pair,




Figure 5: Keypoint transfer results. We show the quality of dense correspondence results by transferring ground-truth keypoints from **source** images in the top-row to **target** images in the bottom-row. It is interesting to note that method is able to transfer keypoints despite significant changes in the viewpoint.

where the source image has some annotated keypoints, the goal is to predict the location of these keypoints in the target image. We first describe the datasets used to train our model, and then briefly survey the various baselines we compare to and then present the evaluation results.

4.1.1 Experimental Setup

Datasets. We use bird images from the CUB-200-2011[43] and the car images from PASCAL3D+ [47] dataset for quantitative evaluation. CUB-200-2011 contains 6000 training and test images with 200 different species. Each bird has 14 annotated keypoints, a segmentation mask, and a bounding box. Note that we only use keypoint annotation at test time to evaluate our method on the task of dense correspondence as described earlier. We also train a model on the car category from PASCAL3D+ [47] which has over 6000 training and test images but evaluate only on cars from PASCAL VOC [9] with 12 keypoint annotations per instance. We downloaded a freely available mesh from [1] to serve as a bird template shape, used an average of 10 Shapenet [6] models to obtain a template shape for cars.

Baselines. We report comparisons to several methods that leverage varying amount of supervision for learning:

Category Specific Mesh Reconstruction (CMR) [18] learns to reconstruct the 3D shape and predict pose for a given instance, but relies on training time supervision of known keypoint locations and segmentation masks. Since a common morphable model is used across a category, we can compute the implied surface mappings via computing for each pixel, the coordinate of the mean shape that is rendered at its location (or nearest location in case of imperfect projection). We can then infer correspondences as in Section 3.4.

Zhou et al. [52] exploit a large collection of 3D synthetic models to learn dense correspondence via cycle-consistency. During training, they crucially rely on pose supervision (from PASCAL 3D+), as each cycle consists of

synthetic images rendered from the same view as the real image pair. Their method outputs dense correspondences in the form of a per-pixel flow, and infers non-correspondence using a ‘matchability’ score.

Dense Equivariance (DE) [38] is a self-supervised method to learn correspondences, and does not require any pose or keypoint annotations. We re-implement this baseline such that it can exploit the annotations for object masks (see appendix for details). DE learns a per-pixel feature vector, and enforces corresponding pixels to have a similar feature. The supervision for correspondences is obtained via applying known in-plane random warps to images. During inference, we can recover the correspondence for a source pixel by searching for the most similar feature in the target image.

VGG Transfer. Inspired by Long *et al.*’s [25] observation that generic learned features allow recovering correspondences, we designed a baseline which infers correspondence via nearest neighbours in this feature space. Specifically for a pixel in the source image we lookup its VGG feature from the conv4 layer and finds its corresponding nearest neighbour in the target image (we found these features to perform better than AlexNet used by Long *et al.* [25]).

4.1.2 Evaluation Metrics

We evaluate the various methods on two metrics: a) Percentage of Correct Keypoints (PCK), and b) Keypoint Transfer AP (APK). We use two separate metrics, because while the PCK metric evaluates the accuracy of keypoint transfer for keypoints that are visible in both, source and target image, it does not disambiguate if an approach can infer that a particular source keypoint does not correspond to any pixel on the target. So therefore, while PCK lets us evaluate correspondence accuracy, the APK metric also lets us measure accuracy at inferring *non-correspondence*.

Percentage of Correct Keypoints (PCK): Given a (source, target) image pair with keypoint annotations on the source,




Figure 6: Predicted Canonical Surface mapping for six different categories. The color at each image pixel depicts the color at the corresponding surface point on the 3D template shape in the left row. While the predictions are mostly accurate, some error modes include: a) inferring globally incorrect CSM due to pose ambiguity (*e.g.* third horse), or b) incorrect local predictions due to missing segmentation (*e.g.* the second sheep).

each method predicts a single estimate for the corresponding location in the target image. The PCK metric reports the mean accuracy of keypoint predictions across keypoints that are common across pairs. A prediction is considered correct only when the predicted location lies within $\alpha * \max(h, w)$ radius around the ground truth annotation for the transfer. We report results with $\alpha = 0.1$, and h, w refer to height and width of the image to which the keypoints were transferred.

Keypoint Transfer AP (APK): In addition to predicting a location in target image for each each keypoint in source

image, this metric requires a confidence in the estimate. Ideally, if a source keypoint does not correspond in a target image, the corresponding predicted confidence should be low, whereas it should be high in case of a keypoint visible in both. Our approach and CMR [18] can rely on the (inverse) distance on the template/mean shape as a confidence measure. Zhou *et al.* [52] produce a ‘matchability’ score, and the feature based methods ‘DE’ [38] and ‘VGG transfer’ [25] can leverage feature similarity as confidence.

Given these predictions, we vary the confidence thresholds,

Annotation	Method	Birds		Cars	
		PCK	APK	PCK	APK
KP + Seg. Mask	CMR [18]	47.3	22.4	44.1	16.9
Pose + Syn. Data	Zhou et. al [52]	-	-	37.1	10.5
Pose + Seg. Mask	CSM (ours) w/ pose	56.0	30.6	51.2	21.0
Seg. Mask	Dense Equi [38]	34.8	11.1	31.5	5.7
	VGG Transfer	17.2	2.6	11.3	0.6
	CSM (ours)	48.0	22.4	40.0	11.0

Table 1: PCK and APK. Percentage of correct keypoints (PCK) and Keypoint Transfer AP (APK) at $\alpha = 0.1$. See Section 4.1.2 for metric descriptions. All evaluations are on 10000 image pairs per category. Higher is better.




Figure 7: Keypoint Transfer PR Curves. We report the transfer precision vs recall curves for all the methods on the task of keypoint transfer. Dashed lines represent methods with pose or keypoint supervision. Solid lines denote approaches without such supervision. The area under the curve is reported in the legend for each of the plots (higher is better). The plot on left is for CUBS-Birds [43], and the one on the right is on cars and keypoints from PascalVOC [9]. See Section 4.1.2 for metric descriptions.

and plot ‘Transfer Precision’ vs ‘Transfer Recall’ and report the area under the curve as the AP. ‘Transfer Recall’ measures the fraction of correspondences in the ground-truth that have been recovered above the threshold (at the lowest confidence threshold, this value is similar to PCK). ‘Transfer Precision’ measures the fraction of correspondences above the threshold that are correct (a prediction for a non-corresponding keypoint is always deemed incorrect). For a high precision, a method should predict low confidence scores for non-corresponding keypoints. We explain these metrics in more detail in the appendix.

4.1.3 Results

In addition to reporting the performance of our method, without any pose supervision, we also evaluate our approach when using pose supervision (denoted as ‘CSM w/Pose’) to better compare to baselines that use similar [52] or more [18] annotations. However, note that all results visualization in the paper are in a setting *without* known pose. We report the PCK and APK results in Table 1, and observe

that our approach performs better than the alternatives. We also show the Transfer AP plots in Figure 7, and note large the relative performance boost (in particular over the self-supervised method [38]), indicating that our approach, in addition to inferring correspondences when they exist, can realize when regions do not correspond. We also visualize some qualitative results for keypoint transfer in Figure 5.

4.2. Learning from Unannotated Image Collections

As our method does not require keypoint supervision during training, we can apply it to learn canonical surface mappings for generic classes using just category-level image collections (with automatically obtained segmentation). We use images for various categories from ImageNet [8], obtain instance segmentation using an off-the-shelf system [14], and manually filter out instances with heavy occlusion. This results in about 1000 instances per category, and we train our CSM predictors using a per-category template model downloaded from the web (in fact, for zebras we use a horse model). We show qualitative results (on held-out images) in Figure 6 and observe that we learn accurate mappings that also respect correspondence across instance. Please see supplementary for additional visualizations.

5. Discussion

We present an approach to learn canonical surface mappings for generic categories using a geometric cycle consistency objective. Our approach allows us to do so without keypoint or pose supervision, and learn CSM prediction and infer dense correspondence while only relying on foreground masks as supervision. While this is an encouraging step towards understanding the underlying 3D structure and associations across images, several challenges still remain. In particular, as we seek to explain the per-pixel predictions via a reprojection of a single rigid template, our approach is not directly applicable to categories where the shapes across instances differ significantly or undergo large articulation. It would be interesting to extend our method to also allow for predicting the underlying deformation and articulation in addition to camera transforms. Additionally, while our approach allowed relaxing correspondence supervision, it would be desirable to take a step further, and learn from unannotated image collections without foreground mask supervision. Lastly, our approach leveraged geometric cycle consistency, and videos may provide an additional learning signal by enforcing consistency of predictions through time [45].

Acknowledgements. We would like to thank the members of the CMU Visual and Robot Learning group and the anonymous reviewers for helpful discussions and feedback. ST would also like to thank Alyosha Efros for his constant nudges to ‘work on cycles!’.

References

- [1] Free3d.com. <http://www.free3d.com>.
- [2] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *CVPR*, 2018.
- [3] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. SCAPE: Shape Completion and Animation of PEOple. *SIGGRAPH*, 2005.
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, 1999.
- [5] Thomas J Cashman and Andrew W Fitzgibbon. What shape are dolphins? building 3d morphable models from 2d images. *TPAMI*, 2013.
- [6] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [7] Christopher B Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. In *NeurIPS*, 2016.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [9] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.
- [10] Ravi Garg, Vijay Kumar BG, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *ECCV*. Springer, 2016.
- [11] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017.
- [12] JunYoung Gwak, Christopher B Choy, Animesh Garg, Manmohan Chandraker, and Silvio Savarese. Weakly supervised 3d reconstruction with adversarial constraint. In *3DV*, 2017.
- [13] Bumsu Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow. In *CVPR*, 2016.
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [16] Eldar Insafutdinov and Alexey Dosovitskiy. Unsupervised learning of shape and pose with differentiable point clouds. In *NeurIPS*, 2018.
- [17] Angjoo Kanazawa, David W Jacobs, and Manmohan Chandraker. WarpNet: Weakly supervised matching for single-view reconstruction. In *CVPR*, 2016.
- [18] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018.
- [19] Abhishek Kar, Shubham Tulsiani, João Carneira, and Jitendra Malik. Category-specific object reconstruction from a single image. In *CVPR*, 2015.
- [20] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *CVPR*, 2018.
- [21] Sameh Khamis, Jonathan Taylor, Jamie Shotton, Cem Keskin, Shahram Izadi, and Andrew Fitzgibbon. Learning an efficient model of hand shape variation from depth images. In *CVPR*, 2015.
- [22] Jaechul Kim, Ce Liu, Fei Sha, and Kristen Grauman. Deformable spatial pyramid matching for fast dense correspondences. In *CVPR*, 2013.
- [23] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *ICML*, 2017.
- [24] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *TPAMI*, 2011.
- [25] Jonathan L Long, Ning Zhang, and Trevor Darrell. Do convnets learn correspondence? In *NeurIPS*, 2014.
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *SIGGRAPH Asia*, 2015.
- [27] Gerard Pons-Moll, Jonathan Taylor, Jamie Shotton, Aaron Hertzmann, and Andrew Fitzgibbon. Metric regression forests for correspondence estimation. *IJCV*, 2015.
- [28] Emil Praun and Hugues Hoppe. Spherical parametrization and remeshing. In *TOG*, 2003.
- [29] Danilo Jimenez Rezende, SM Ali Eslami, Shakir Mohamed, Peter Battaglia, Max Jaderberg, and Nicolas Heess. Unsupervised learning of 3d structure from images. In *NeurIPS*, 2016.
- [30] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised geometry-aware representation for 3d human pose estimation. In *ECCV*, 2018.
- [31] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *CVPR*, 2017.
- [32] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. End-to-end weakly-supervised semantic alignment. In *CVPR*, 2018.
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [34] Tanner Schmidt, Richard Newcombe, and Dieter Fox. Self-supervised visual descriptor learning for dense correspondence. *IEEE Robotics and Automation Letters*, 2017.
- [35] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *CVPR*, 2017.
- [36] Jonathan Taylor, Jamie Shotton, Toby Sharp, and Andrew Fitzgibbon. The vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation. In *CVPR*, 2012.
- [37] Jonathan Taylor, Richard Stebbing, Varun Ramakrishna, Cem Keskin, Jamie Shotton, Shahram Izadi, Aaron Hertzmann, and Andrew Fitzgibbon. User-specific hand modeling from monocular depth sequences. In *CVPR*, 2014.
- [38] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object frames by dense equivariant image labelling. In *NeurIPS*, 2017.
- [39] D’Arcy Thompson. *On Growth and Form*. Cambridge Univ. Press, 1917.
- [40] Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Multi-view consistency as supervisory signal for learning shape and pose prediction. In *CVPR*, 2018.

- [41] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *CVPR*, 2017.
- [42] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NeurIPS*, 2017.
- [43] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [44] Xiaolong Wang, Kaiming He, and Abhinav Gupta. Transitive invariance for self-supervised visual representation learning. In *ICCV*, 2017.
- [45] Xiaolong Wang, Allan Jabri, and Alexei A. Efros. Learning correspondence from the cycle-consistency of time. In *CVPR*, 2019.
- [46] Jiajun Wu, Tianfan Xue, Joseph J Lim, Yuandong Tian, Joshua B Tenenbaum, Antonio Torralba, and William T Freeman. Single image 3d interpreter network. In *ECCV*, 2016.
- [47] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *WACV*, 2014.
- [48] Xinchun Yan, Jimei Yang, Ersin Yumer, Yijie Guo, and Honglak Lee. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In *NeurIPS*, 2016.
- [49] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In *CVPR*, 2017.
- [50] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.
- [51] 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.
- [52] Tinghui Zhou, Philipp Krähenbühl, Mathieu Aubry, Qixing Huang, and Alexei A. Efros. Learning dense correspondence via 3d-guided cycle consistency. In *CVPR*, 2016.
- [53] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.