

WarpNet: Weakly Supervised Matching for Single-view Reconstruction

Angjoo Kanazawa David W. Jacobs
University of Maryland, College Park

Manmohan Chandraker
NEC Labs America

Abstract

We present an approach to matching images of objects in fine-grained datasets without using part annotations, with an application to the challenging problem of weakly supervised single-view reconstruction. This is in contrast to prior works that require part annotations, since matching objects across class and pose variations is challenging with appearance features alone. We overcome this challenge through a novel deep learning architecture, WarpNet, that aligns an object in one image with a different object in another. We exploit the structure of the fine-grained dataset to create artificial data for training this network in an unsupervised-discriminative learning approach. The output of the network acts as a spatial prior that allows generalization at test time to match real images across variations in appearance, viewpoint and articulation. On the CUB-200-2011 dataset of bird categories, we improve the AP over an appearance-only network by 13.6%. We further demonstrate that our WarpNet matches, together with the structure of fine-grained datasets, allow single-view reconstructions with quality comparable to using annotated point correspondences.

1. Introduction

Reconstructing an object from a single image is a significant challenge, that can be tackled by matching keypoints to other instances in a fine-grained dataset. However, such datasets exhibit large intra-class shape variations or inter-class appearance variations, which cannot be handled by traditional features such as SIFT [20]. Recently, methods have been proposed to match instances across categories, relying on supervision in the form of part (keypoint) annotations [7, 15, 29] or 3D CAD models [3, 8] to augment appearance information with shape priors. Such annotations are labor-intensive, thus, too sparse for reconstruction and not scalable. Further, it can be quite difficult to obtain human-labeled annotations for parts that are not nameable. In contrast, this paper presents a framework to match images of objects with some degree of non-rigidity and articulation, across category and pose variations, without requiring supervised annotations. We then present an approach to the challenging novel problem of unsupervised single-view object reconstruction.

We postulate that the structure of fine-grained datasets,

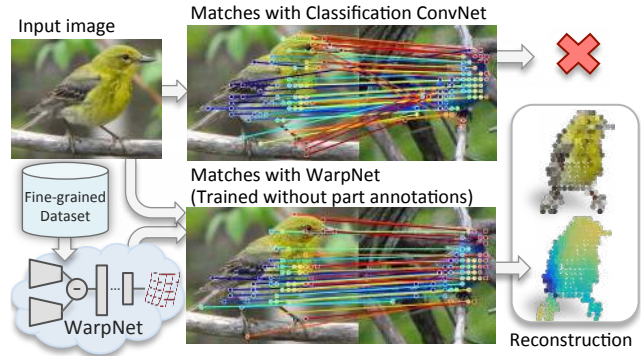


Figure 1: Given a single image of an object, we propose a novel deep learning framework for obtaining keypoint matches to other objects in a fine-grained dataset, without using any part annotations. The output of our network is used as spatial prior for matching across variations in appearance, pose and articulation (bottom), which is not possible with appearance features alone (top). Our match quality is high enough to be propagated across images to be used for single-view reconstruction without using any manually annotated keypoints (right).

combined with the power of convolutional neural networks (CNNs), allows matching instances of different categories without supervised annotation. Fine-grained datasets for objects such as birds can be analyzed along two dimensions – appearance and shape. Instances within the same category that are imaged in different poses can be matched by appearance similarity, while instances with similar pose or viewpoint from different categories can be matched through similarity in global shape. Instances with both appearance and shape variations may then be matched by propagation (Fig. 3). In Section 3, we demonstrate a practical realization of this intuition by introducing a deep learning architecture, WarpNet, that learns to warp points on one object into corresponding ones on another (from a possibly different category or pose), without requiring supervised annotations.

WarpNet is a Siamese network that accepts two images as input (Section 3.2). To overcome the absence of annotated keypoints, our training presents an image and a warped version related by a known thin-plate spline (TPS) transformation, which yields artificial correspondences. We assume the object bounding box and foreground segmentation are known, which can be obtained through state-of-the-art segmentation [10] or co-segmentation methods [17]. We experi-

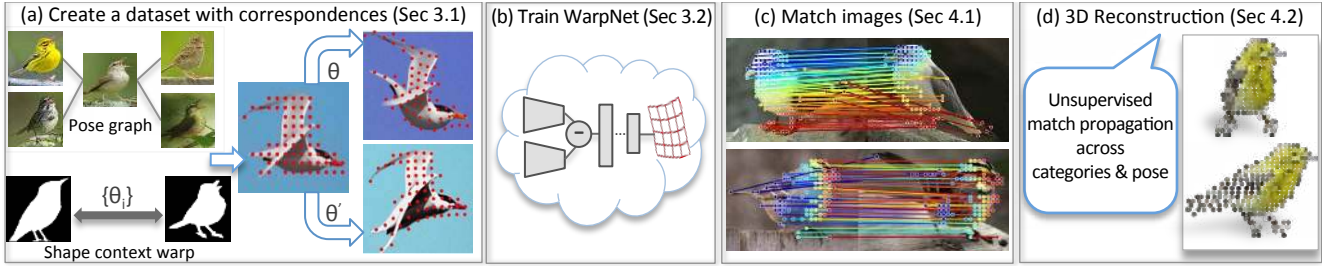


Figure 2: Overview of our framework. (a) Lacking part annotations, we exploit the fine-grained dataset to create artificial correspondences. (b) These are used to train our novel deep learning architecture that learns to warp one object into another. (c) The output of the network is used as a spatial prior to match across appearance and shape variations. (d) Our high-quality matches can be propagated across the dataset. We use the WarpNet output and the structure of fine-grained categories to perform single-view reconstruction without part annotations.

ment using both ground truth and co-segmentation outputs. In Section 3.1, we exploit neighborhood relationships within the dataset through the pose graph of Krause *et al.* [17] to compute exemplar TPS transformations between silhouettes, from which our artificial transformations are sampled. A point transformer layer inspired by [14] is used to compute the warp that aligns keypoints without supervision, which provides a spatial prior for matching (Section 4). We show that WarpNet generalizes well to match real images with distinct shapes and appearances at test time. In particular, it achieves matching accuracy over 13.6% higher than a baseline ILSVRC CNN [9].

Establishing matches between a given instance and other objects in the dataset opens the door to a novel problem – weakly supervised reconstruction in fine-grained datasets. Several sub-problems must be solved to achieve this goal, such as match propagation and image subset selection. Prior works such as [7, 29] approach these sub-problems, but the absence of supervised annotations poses new challenges. In Section 4.2, we suggest ways to overcome them through the use of matches from our WarpNet, the pose graph and heuristics that exploit the structure of fine-grained datasets. We demonstrate reconstructions that are nearly as good as those obtained using supervised annotations and better than those from appearance-only CNNs or unsupervised baselines such as deformable spatial pyramids [16].

To summarize, our key contributions are:

- A novel deep learning architecture, WarpNet, that predicts a warp for establishing correspondences between two input images across category and pose variations.
- A novel exemplar-driven mechanism to train WarpNet without requiring supervised keypoint annotations.
- An approach to unsupervised single-view object reconstruction that exploits the structure of the fine-grained dataset to yield reconstructions of birds nearly on par with the method that uses supervised part annotations.

2. Related Work

Supervised matching Several recent approaches use deep learning to learn a similarity metric between image patches

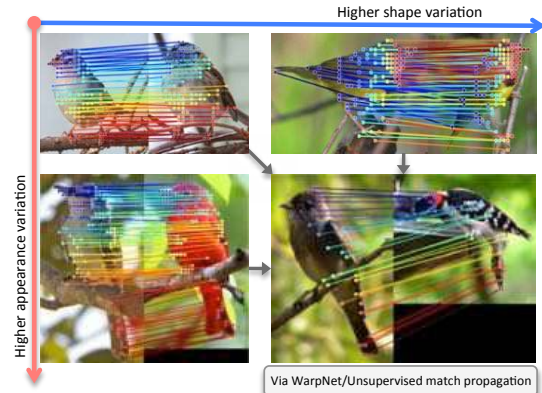


Figure 3: Intuition for matching in fine-grained datasets without supervised point annotations. Matching within a category exploits appearance similarity, while matching instances across related categories is possible through global shape similarity. By propagation, one may match across variations in both appearance and shape.

in a supervised manner [13, 33]. These works focus on matching images of the same instance (for example, the Statue of Liberty [31]) from various viewpoints, while we match deformable objects of different instances exhibiting a wide variety of appearances. Our task requires semantic understanding of object shape, beyond just local appearance. A CNN framework to predict dense optical flow on general scenes is proposed by [12], but in a supervised manner.

Matching or keypoint localization may be improved by augmenting appearance similarity with spatial priors. Supervised methods often use a dataset with labeled parts to obtain a non-parametric prior on keypoint locations [4, 19, 28]. These priors may be learned from data [26], but require supervised part annotations during training. Such annotation is laborious and consequently available only for a few nameable parts, which might be too sparse for reconstruction.

Unsupervised matching Also related to our approach are methods that use unsupervised spatial priors for dense matching [16, 27]. Unlike our work, these methods are purely geometric and do not learn category-specific semantic properties. Recently, [34] proposes an unsupervised approach for dense alignment of image sets. But while their focus is

global consistency, our emphasis is on pairwise matching through the WarpNet framework (for which they use flow). Thus, our contribution is complementary and may be used by their framework. We evaluate quantitatively on deformable bird categories, while they use rigid categories on PASCAL.

Single-view reconstruction A new challenge in computer vision is to reconstruct a target object from a single image, using an image collection of similar objects. The seminal work of [29] demonstrates the possibility of a solution, but relies on ground truth part annotations to establish correspondences. The subsequent works of [15, 7] take a step further in using part annotations only during training. In contrast, we do not require part annotations at either train or test time.

CNNs for learning transformations Similar to the recent work of [2], we use a Siamese network to predict transformations. The key difference is that predicting the ego-motion transformation in [2] is a pretext for feature learning, while we directly use the predicted transformation as well as its appearance features for matching. Further, they require ground truth transformation parameters in order to train their network, while we use the structure of the fine-grained dataset to generate artificial correspondences and implicitly optimize the parameters. Finally, rigid transformations in [2] are discretized in bins and the task is posed as classification, while our network outputs continuous thin-plate spline transformation parameters with a matching objective.

Our architecture is inspired by the recent spatial transformer network of Jaderberg *et al.* [14], which introduces a deep learning module to predict a spatial transformation. This acts as an attention mechanism driven by a classification objective. We extend the idea further to predict a warping function that aligns two object instances in an unsupervised manner. Our approach is in line with the recent work of [1], which demonstrates that CNNs can be trained without supervised labels by treating an image patch and its transformed versions as a “surrogate” class. However similar to [2], the unsupervised training objective of classifying the surrogate class is geared towards learning good features, while we show that the output of our network trained by an artificial dataset actually generalizes to matching real image pairs.

3. Learning without Part Annotations

We present a deep learning framework, *WarpNet*, that learns the correspondence from one image to another without requiring part annotations. Given two images I_1 and I_2 , our network outputs a function that takes points in I_1 to points in I_2 . We parameterize this function as a thin-plate spline (TPS) transformation since it can capture shape deformations well [5]. Inspired by Dosovitskiy *et al.* [1], we generate artificial correspondences by applying known transformations to an image. However, our approach is distinct in using the structure afforded by fine-grained datasets and dealing with non-rigidity and articulations. Our network



Figure 4: Sample exemplar-TPS warped images used for training our WarpNet. Left: original images, right: artificial versions made by applying exemplar TPS warp + chromatic transformation. Notice changes in shape and articulations at the head and the tail.

generalizes well to instances of different categories at test time and we use its output as a spatial prior in computing a match between two objects. Figure 2 gives an overview of our approach. We discuss each step in detail below.

3.1. Generating Unsupervised Correspondences

Since we do not have annotated point correspondences, we create artificial ones by applying random spatial and chromatic transformations to images. The key requirement is that the spatial transformations applied are complex enough to learn meaningful correspondences, while producing transformed images that are reflective of actual image pairs to match at test time. For instance, affine transformations are not expressive enough to capture non-rigid deformations and articulations in birds. Instead, we use TPS transformations and exploit the fine-grained dataset to generate exemplar warps that span a realistic range of transformations.

We use the pose graph of Krause *et al.* [17], whose edge weights are determined by the cosine distance of the fourth-layer of a pre-trained ILSVRC CNN, which captures abstract concepts such as class-independent shape. We compute shape context TPS warps [5] between the silhouettes of images that are within 3 nearest-neighbors apart on the pose graph. We sort the TPS warps using the mean of their bending and affine energy, retaining only those between the 50th and 90th percentiles to avoid warps that are too trivial or too drastic. We create m transformed versions of every image by sampling from this set of TPS warps. We sample n points uniformly on the foreground, which we use as correspondences. Figure 4 shows the effect of transformations sampled from the exemplar-TPS warps. The images on the left are the originals and the ones on the right are transformed versions. Notice how the transformation induces changes in shape and articulations around the head and the tail, which validates the utility of our exemplar TPS warps.

3.2. WarpNet Architecture

Our proposed WarpNet is a Siamese network [11] that takes two images related by an exemplar TPS transformation, I_1 and I_2 , along with the corresponding n keypoint locations,

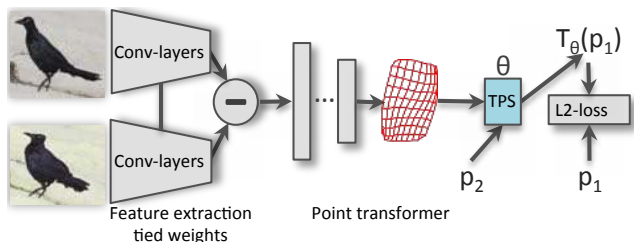


Figure 5: WarpNet architecture. Visual features are extracted from two input images using a Siamese CNN. They are combined to predict a deformed grid that parameterizes a TPS transformation. The network objective is to minimize the distance between corresponding points p_1 and p_2 of the image pair after applying the predicted transformation to p_2 .

as inputs during training (at test time, the input consists only of two images from possibly different categories and poses that must be matched). The main objective of WarpNet is to compute a function that warps points p_2 in I_2 to image coordinates in I_1 , such that after warping the L2 distance to the corresponding points p_1 in I_1 is minimized. Figure 5 illustrates the architecture of WarpNet.

First, the input images are passed through convolution layers with tied weights. The extracted features are then combined by element-wise subtraction of the feature maps. We subtract rather than concatenate the feature maps along the channels, since concatenation significantly increases the number of parameters in the network making it unstable to train. The combined feature maps are passed through a point transformer, similar to [14], which regresses on the (x, y) coordinates of a deformed $K \times K$ grid. The output grid, normalized to a range of $[-1, 1] \times [-1, 1]$, acts as the control points for computing a grid-based TPS transformation from I_2 to I_1 . This involves solving a system of linear equations, handled by the TPS layer. Please see the supplementary materials for details. The predicted TPS transformation is applied to the keypoints of I_2 generating the transformed version $T_\theta(p_2)$, which finally gets sent to the L2 loss layer along with p_1 . Since every step consist of linear operations, the whole network can be trained with backpropagation.

We implicitly train the warp parameters in terms of distance between corresponding points rather than direct supervision against the TPS warp coefficients. This provides a natural distance between warps, where we can train the network without knowing the exact transformation parameters used.

Figure 6 illustrates the output of the trained network given two real images as input, denoted source and target. Despite the fact that the network has never seen objects of different instances, it is able to compute warps between the two objects. Note that WarpNet accounts for variations in shape (fat to skinny, small to large birds), articulation (such as the orientation of the head or the tail) and appearance.

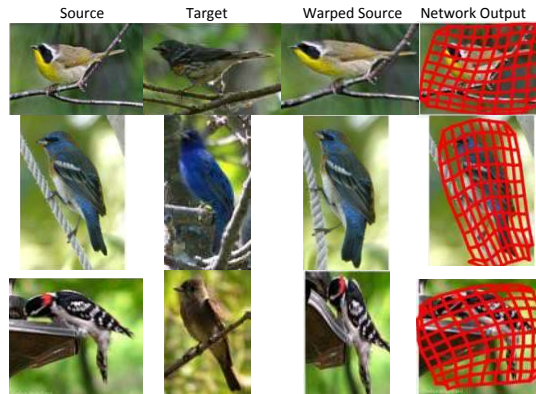


Figure 6: Visualizations of the network output. WarpNet takes two images, source and target, as inputs and produces a 10×10 deformed lattice (last column) that defines a TPS warp from target to source. The third column shows the warped source image according to the network output. Notice how the network accounts for articulations at the tail and the head as well as differences in shape of the birds. WarpNet is trained in an unsupervised manner and none of these images were seen by the network during training.




Figure 7: Sample matches obtained by ILSVRC trained CNN versus WarpNet, where WarpNet’s relative robustness to variations in appearance, pose and articulation may be noted.

4. Matching and Reconstruction

4.1. Matching with WarpNet

Given two images I_i and I_j , a match for a point u_i in I_i is the most similar point v_j in I_j using the similarity score consisting of an appearance term and a spatial term:

$$s(u_i, v_j) = \exp\left(\frac{-d_f(u_i, v_j)}{\sigma_f}\right) + \lambda \exp\left(\frac{-d_w(u_i, v_j)}{\sigma_w}\right), \quad (1)$$

where $d_f(u, v)$ is the L2 distance of appearance features extracted at u_i and v_j , while d_w is a symmetric spatial prior:

$$d_w(u, v) = \frac{1}{2}(\|\mathbf{x}_i^u - T_{\theta_{ij}}(\mathbf{x}_j^v)\| + \|\mathbf{x}_j^v - T_{\theta_{ji}}(\mathbf{x}_i^u)\|). \quad (2)$$

We use WarpNet to compute $T_{\theta_{..}}$ in both directions.

The matches are then ranked by the ratio-test strategy [20], which allows discarding points in I_i that are similar to many other points in I_j . Since the keypoints are extracted densely on the foreground, we compute the similarity score ratio between the first and second nearest neighbors that are at least 10 pixels away. Figure 7 shows a few qualitative matching results comparing the baseline CNN and WarpNet.

4.2. Single-View Object Reconstruction

Obtaining good matches is a critical first step towards 3D reconstruction. While single-view 3D reconstruction methods in the past have relied on expensive supervised inputs such as part annotations or CAD models, our matching enables a first approach towards a challenging new problem, namely, part annotation free single-view reconstruction. We discuss initial approaches to variants of existing supervised methods or structure from motion (SFM) pipelines that may be used to solve this problem without requiring annotations.

Propagating correspondences In the CUB-200-2011 dataset, there are only 60 images for each category. Moreover, birds are often imaged from preferred viewpoints, but it is critical for reconstruction to obtain matches across a well-distributed set of viewpoints. On the other hand, deformations may be very high even within a category (open wings as opposed to closed), which makes straightforward matching within a category challenging. Inspired by the work of Carreira *et al.* [7], we use a shortest path method to propagate matches across objects of similar shapes in the dataset, in order to obtain a denser set of tracks. However, note that we lack the initial set of point annotations as well as the camera poses obtained through part annotations in [7, 29], who also manually select a subset of keypoints to eliminate articulations. Instead, we determine unsupervised matches purely through our WarpNet and rely on the pose graph to determine nearest neighbors for propagation.

Choosing a subset for reconstruction A key problem we encounter is the choice of images for reconstruction. In previous works on reconstruction within PASCAL VOC [7, 29], it has been possible to use the entire dataset since it contains less than 1000 images for birds. In contrast, CUB-200-2011 contains nearly 12000 images, which poses computational challenges and requires greater vigilance against outliers. Moreover, annotations in [7, 29] preclude the need for algorithmic considerations on baseline or shape variations in choosing the image set. For instance, to reconstruct a sitting bird imaged from a frontal view, we must propagate matches to side views of sitting birds in other categories to ensure a good baseline, while avoiding images of flying birds.

Given a collection of images, several heuristics have been proposed for selecting the right subset or order for multi-view rigid-body reconstruction [24, 25]. However, those are not directly applicable for single-view reconstruction of deformable objects. Instead, we propose three heuristics that utilize the structure of fine-grained bird datasets:

- Use images from categories that share a keyword (for example, all “warblers”, or all “sparrows”).
- Use images from categories that are related by an ornithological taxonomy, as defined by [23].
- Use images from the five nearest neighbor subcategories on a similarity tree of bird species [6].

The above heuristics perform comparably and address the same goal – introduction of matched keypoints from more than one subcategory to ensure good viewpoint coverage.

Reconstruction Given an image of a target object from one particular class, we consider images from several other categories using one of the above heuristics. We compute pairwise matches at 85% precision threshold between all pairs of images whose distance on the pose graph is less than 4. We ignore pairs that have less than 50 surviving matches. We then set up a virtual view network [7] to propagate matches across all the selected images. We use scores from (1), bounded between $[0, 1]$, as weights on the graphs connecting the keypoints. After propagation, we discard as spurious any propagated matches with shortest path distance more than 0.4 and remove all images that have less than 30 matches with the target object. We then create the measurement matrix of tracked keypoints of the target object. We only consider keypoints visible in at least 10% of the images as stable enough for reconstruction. We finally send the observation matrix to the rigid factorization method of [21], which robustly handles missing data, to obtain 3D shape.¹

5. Experiments

We perform experiments on the CUB-200-2011 dataset which contains 11788 images of 200 bird categories, with 15 parts annotated [30]. We reconstruct without part annotation, assuming objects are localized within a bounding box. We quantitatively evaluate our matches using and extending the part annotations. Next, we evaluate the effectiveness of WarpNet as a spatial prior and analyze the choice of transformations for creating the artificial training dataset. Finally, we demonstrate the efficacy of our framework with several examples of unsupervised single-view reconstruction.

5.1. Experimental Details

We create the pose graph of [17] using the `conv4` feature of AlexNet trained on ILSVRC2012 [18]. For creating the artificial dataset, we only use the training data (~ 6000 images) and create $m = 9$ copies of each image using our exemplar-TPS. We resize all images to 224×224 . This results in approximately 120k image pairs, each with $n = 100$ point correspondences. Following [1], we apply spatial and chromatic data augmentation on-the-fly during training.

We use the VGG-M architecture of [9] until the `pool5` layer as the feature extraction component of WarpNet. The point transformer consists of C512-C256-F1024-D-Op using the notation of [2]. Both convolutional layers use 3×3 kernel, stride 1 with no padding, with ReLU non-linearity. The output layer is a regressor on the grid coordinates, with grid size $K = 10$. The feature extraction weights are initialized

¹A rigid factorization suffices to produce good reconstructions since the dataset is large enough, but non-rigid methods alternately could be used.

with weights pre-trained on the ILSVRC classification task, following prior state-of-the-art for correspondence [19].

For matching and reconstruction, images are resized with aspect ratio intact and the smallest side 224 pixels. We uniformly sample points on the foreground with a stride of 8 as keypoints for matching. For all experiments we use L2-normalized conv4 features extracted at the keypoints using the hole algorithm [10] for computing the appearance term in (1). Hyperparameters used for matching are $\sigma_f = 1.75$, $\sigma_w = 18$, $\lambda = 0.3$, tuned using the artificial dataset.

5.2. Match Evaluation

We compare our approach with ILSVRC pre-trained VGG-M conv4 [9], SIFT at radius 8 [20] and matches from the deformable spatial pyramid (DSP) [16]. Only the appearance term in (1) is used for computing matches with VGG-M conv4 and SIFT. For computing the matches with DSP, we mask out the background prior to extracting SIFT features following [7] and only keep matches of the keypoints. For this experiment, the set of keypoints to match includes the locations of annotated parts.

In order to evaluate WarpNet as a stand-alone learned spatial prior, we compare WarpNet with DSP by replacing the SIFT features in DSP with VGG features. We call this method *VGG+DSP*. We further evaluate WarpNet against the original DSP by using WarpNet as a spatial prior for SIFT matches, where the unary term d_f in (1) is computed with SIFT features. We call this method *SIFT+WarpNet*.

As discussed in Section 3.1, the only supervision required in training WarpNet is the segmentation mask to mine exemplar-TPS transformations. We also evaluate the robustness of WarpNet using co-segmentation outputs of [17], called *VGG+coseg*.

Test set We evaluate on 5000 image pairs that are within 3 nearest neighbors apart on the pose graph, comprising more than 50k ground truth matches.² Due to the unsupervised nature of the pose graph, these pairs exhibit significant articulation, viewpoint and appearance variations (see Figures 1, 6). We remove severely occluded pairs with less than 7 parts visible in both images and pairs whose TPS warp computed from part annotations have very high bending energy. None of the test images were used to train WarpNet.

Evaluation metrics We evaluate the accuracy of matches with the percentage of correct keypoints (PCK) metric [32], where a match is considered correct if the predicted point is within $\alpha * L$ of the ground-truth correspondence. Following [2], we chose L to be the mean diagonal length of the two images. We also compute the precision-recall (PR) curve adopting the procedure of [22]. A match is considered a true positive within a radius $\alpha = 0.05$, otherwise it is a false positive. In this setup, a recall of 1 is obtained only if all

²Please see supplementary materials for results on a test set with 1-nearest neighbors, where we observe similar trends but with higher PCKs.

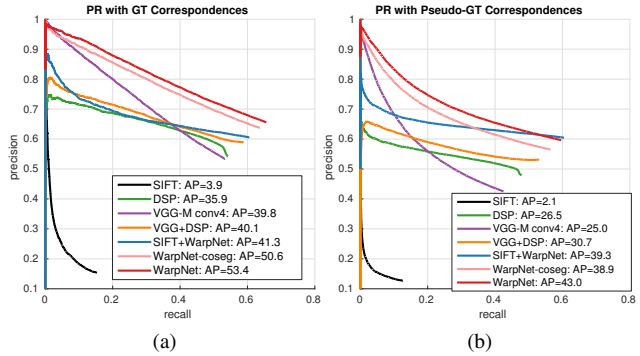


Figure 8: Precision-Recall curves for matching points between neighboring images on the pose graph. We evaluate points with (a) human-annotated correspondences and (b) expanded pseudo-ground-truth correspondences.

the matches retrieved are correct, that is, 100% α -PCK. We compute PR curves using the ratio-test values described in Section 4.1 for ranking the matches and report AP. For DSP, we use its matching cost for ranking instead of the ratios, since second closest matches are not available.

Results Figure 8(a) shows the obtained PR curves. WarpNet achieves an AP of 53.4%, an 13.6% increase over matches using just the appearance feature of VGG-M conv4. WarpNet achieves a much higher recall due to its spatial prior, learned without using any part annotations. As a side note, conv4 features of WarpNet alone achieve very similar performance to the VGG-M conv4. In all cases, WarpNet outperforms DSP as a spatial prior and changing SIFT to VGG features yields around 5% improvement in the final recall. *WarpNet-coseg* still outperforms the baseline VGG-M by 10.8%, showing our approach is applicable even without ground truth segmentations.

Figure 9(a) shows the PCK as a function of α , where WarpNet consistently outperforms other methods. We observe that VGG-M conv4 and DSP perform similarly, showing that while DSP obtains low recall at high precision, its overall match quality is similar to CNN features, an observation in line with [7]. Since only high precision matches are useful for reconstruction where outliers need to be avoided, we show the same curves thresholded at 85% precision in Figure 9(b) for VGG-M and our method. Note that some methods in black have zero recall at this precision. The growing gap between WarpNet and VGG-M conv4 as α increases suggests that, unlike WarpNet, appearance features alone make grossly wrong matches (see Figures 1 and 7).

Expanding the set of part annotations A caveat of the CUB-200-2011 for our task is that part annotations are sparse and concentrated on semantically distinct parts such as eyes and beaks around the head region, with only four points on the bird body that are often not all visible. To investigate matching performance more densely, we carefully expand

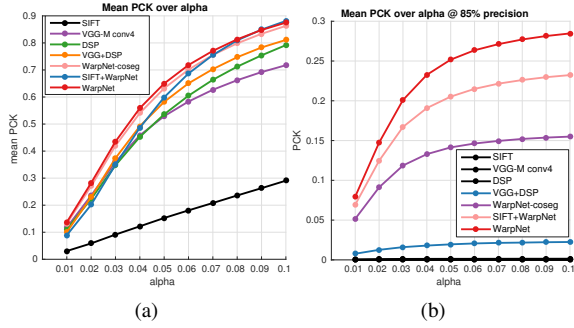


Figure 9: PCK (higher the better) over varying definition of correctness α . (a) Mean PCK of all retrieved matches regardless of ratio score. (b) Mean PCK with matches thresholded at 85% precision, which are the matches used for reconstruction.

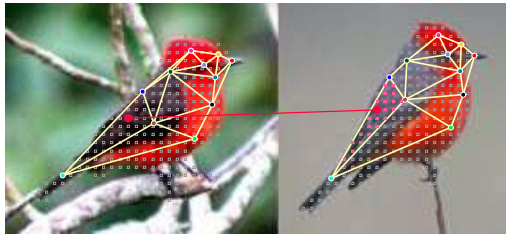


Figure 10: Illustration of the pseudo-gt correspondences. We triangulate each image using the annotated keypoints (colored points). The match for the big red dot in the left image is found by looking at points within the same triangle (small pink dots) in the right image and picking the closest point in terms of barycentric coordinates.

the ground-truth matches using the annotated parts. This process is illustrated in Figure 10. Given a pair of images I_1 and I_2 , we Delaunay triangulate each image independently using the parts visible in both as vertices. For a point u within a triangle in I_1 , we consider points in I_2 that are within the *same* triangle as possible candidates (shown as pink dots in Figure 10), find the point that is closest to u in terms of barycentric coordinates and accept this as a new pseudo ground-truth match if the distance is less than 0.1. Figure 8(b) shows the PR curve obtained using the pseudo-ground truth matches (in addition to the annotated parts). We see the same trends as Figure 8(a), but with a wider gap between the baselines and our method. This is reasonable given that bird bodies usually consist of flat or repeated textures that are challenging to match with local appearances alone, highlighting the efficacy of WarpNet’s spatial prior.

5.3. Choice of Transformations

We now analyze the choice of exemplar TPS transformations for creating the artificial dataset. We train another WarpNet under the same settings, but on an artificial dataset created using only affine spatial transformations, which we refer to as AffineNet. Note that AffineNet’s output is still a TPS transformation, thus, it has the same capacity as the original WarpNet. Figure 11(a) shows the PR curve of AffineNet

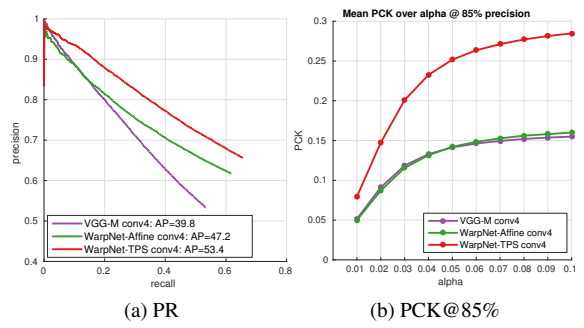


Figure 11: Comparing results for WarpNet trained on artificial data created using affine-spatial transformations with (a) PR curves and (b) PCK over α . WarpNet trained with exemplar-TPS is more effective in terms of recall and precision.

in comparison to WarpNet and VGG-M conv4. WarpNet outperforms AffineNet in all aspects. While AffineNet has a higher final recall (that is PCK of all matches) than VGG-M conv4, its recall at high precision is slightly lower than that of VGG-M conv4. This is highlighted in Figure 11(b), which shows PCK of matches at 85% precision over α , where AffineNet performs on par with VGG-M conv4. This indicates that the warps predicted by AffineNet are helpful in a general sense, but not precise enough to improve the recall at high precision. This experiment shows that using exemplar-TPS transformations for creating the artificial dataset is critical for training a useful WarpNet.

5.4. Single-view Object Reconstruction

We compare our method with three other matching methods. One is a supervised matching approach similar to [7], where the network predicted TPS warp T_θ in (2) is replaced by the supervised TPS warp computed using the annotated keypoints. We call this approach supervised and it is an upper-bound to our method since ground-truth part annotations are used for reconstruction. We also perform reconstructions with VGG-M conv4 features alone and DSP. We do not include the mirrored image as another viewpoint of the target object, since bilateral symmetry does not hold for articulated objects. For post-processing we use the xy-snapping method proposed in [7], which only uses the z -component from the reconstructed shape, while fixing the x , y coordinates. We do not resample the target objects multiple times prior to factorization since it did not seem to make a difference.

Figure 12 shows reconstructions for various types of birds using the four methods from three viewpoints: camera view, 45° azimuth and 45° elevation. The colors indicate depth values (yellow is close, blue is far), with range fixed across all methods. WarpNet produces reconstructions that are most consistent with the supervised approach. Reconstructions from VGG-M and DSP are noisy due to errors in matching and often produce extreme outlier points that

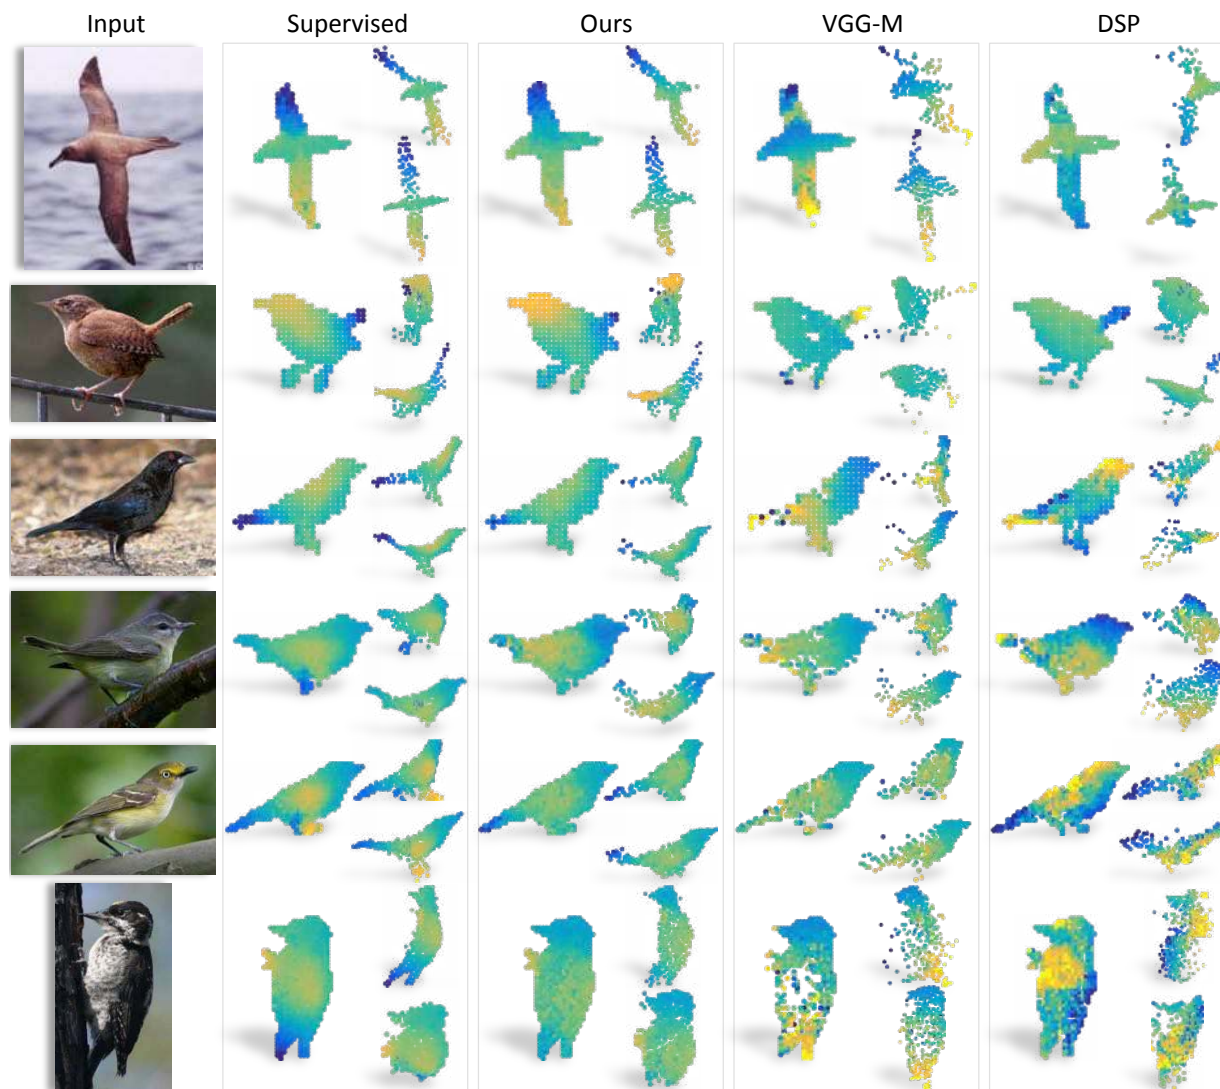


Figure 12: Sample reconstructions showing 3 views for each method: The camera viewpoint followed by the 45° azimuth in counter-clockwise direction (top right) and 45° elevation (bottom right). Colors show the depth where yellow is closer and blue is farther. The supervised method uses the spatial prior computed from annotated part correspondences, which can be seen as an upper bound. No part correspondences were used for the last three methods. WarpNet consistently obtains reconstructions most similar to the supervised method.

had to be clipped for ease of visualization. Articulated parts such as tails and wings are particularly challenging to match, where VGG-M and DSP often fail to recover consistent depths. A weakness of our method is that the TPS prior may sometimes hallucinate birds of similar pose even with wide baseline. This may be avoided by better choice of images for reconstruction. Please see supplementary material for more results, qualitative matches and reconstruction videos.

6. Conclusions and Future Work

We introduce a framework for matching and reconstruction in fine-grained datasets that avoids the expense and scalability challenges of part annotations. The core of our approach is a novel deep learning architecture that predicts a function to warp one object into another. We show that our

network can be trained without supervised part annotations by exploiting the structure of fine-grained datasets and use its output as a spatial prior for accurate matching. Our approach achieves significant improvements over prior state-of-the-art without using part annotations and we show reconstructions of similar quality as supervised methods. Key challenges for future work are to determine optimal subsets of images for reconstruction and a good order for adding images that allows incremental reconstruction with bundle adjustment.

Acknowledgments This work was part of A. Kanazawa’s internship at NEC Labs America, in Cupertino. A. Kanazawa and D. Jacobs were also supported by the National Science Foundation under Grant No. 1526234.

References

- [1] A. Dosovitskiy, J.T. Springenberg, M. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in Neural Information Processing Systems 27 (NIPS)*, 2014. 3, 5
- [2] P. Agrawal, J. Carreira, and J. Malik. Learning to see by moving. In *ICCV*, 2015. 3, 5, 6
- [3] S. Y.-Z. Bao, M. Chandraker, Y. Lin, and S. Savarese. Dense object reconstruction with semantic priors. In *CVPR*. IEEE, 2013. 1
- [4] P. N. Belhumeur, D. W. Jacobs, D. J. Kriegman, and N. Kumar. Localizing parts of faces using a consensus of exemplars. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(12):2930–2940, 2013. 2
- [5] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(4):509–522, 2002. 3
- [6] T. Berg and P. N. Belhumeur. How do you tell a blackbird from a crow? In *Proc. Int. Conf. Computer Vision (ICCV)*, December 2013. 5
- [7] J. Carreira, A. Kar, S. Tulsiani, and J. Malik. Virtual view networks for object reconstruction. In *CVPR*. IEEE, 2015. 1, 2, 3, 5, 6, 7
- [8] T. J. Cashman and A. W. Fitzgibbon. What shape are dolphins? building 3D morphable models from 2D images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(1):232–244, 2013. 1
- [9] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014. 2, 5, 6
- [10] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015. 1, 6
- [11] S. Chopra, R. Hadsell, and Y. L. L. Cun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, pages I: 539–546, 2005. 3
- [12] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. *ICCV*, 2015. 2
- [13] X. Han, T. Leung, Y. Jia, R. Sukthankar, and A. C. Berg. Matchnet: Unifying feature and metric learning for patch-based matching. In *CVPR*. IEEE, 2015. 2
- [14] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks. 2015. 2, 3, 4
- [15] A. Kar, S. Tulsiani, J. Carreira, and J. Malik. Category-specific object reconstruction from a single image. 2015. 1, 3
- [16] J. Kim, C. Liu, F. Sha, and K. Grauman. Deformable spatial pyramid matching for fast dense correspondences. In *CVPR*. IEEE, 2013. 2, 6
- [17] J. Krause, H. Jin, J. Yang, and F.-F. Li. Fine-grained recognition without part annotations. In *CVPR*, pages 5546–5555. IEEE, 2015. 1, 2, 3, 5, 6
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 5
- [19] J. L. Long, N. Zhang, and T. Darrell. Do convnets learn correspondence? In *Advances in Neural Information Processing Systems*, pages 1601–1609, 2014. 2, 6
- [20] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 1, 4, 6
- [21] M. Marques and J. P. Costeira. Estimating 3D shape from degenerate sequences with missing data. *Computer Vision and Image Understanding*, pages 261–272, Feb. 2009. 5
- [22] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(10):1615–1630, 2005. 6
- [23] C. L. of Ornithology. The birds of north america online (p. rodewald, ed.). <http://bna.birds.cornell.edu/BNA/>, August 2015. 5
- [24] N. Snavely, S. M. Seitz, and R. Szeliski. Modeling the world from Internet photo collections. *IJCV*, November 2008. 5
- [25] N. Snavely, S. M. Seitz, and R. Szeliski. Skeletal graphs for efficient structure from motion. In *Proc. Computer Vision and Pattern Recognition*, 2008. 5
- [26] J. J. Tompson, A. Jain, Y. LeCun, and C. Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *Advances in Neural Information Processing Systems*, pages 1799–1807, 2014. 2
- [27] Y. Tsuduki and H. Fujiyoshi. A method for visualizing pedestrian traffic flow using SIFT feature point tracking. In *PSIVT*, pages 25–36, 2009. 2
- [28] S. Tulsiani and J. Malik. Viewpoints and keypoints. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [29] S. Vicente, J. Carreira, L. de Agapito, and J. Batista. Reconstructing PASCAL VOC. In *CVPR*. IEEE, 2014. 1, 2, 3, 5
- [30] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 5
- [31] S. Winder, G. Hua, and M. Brown. Picking the best DAISY. In *CVPR*, pages 178–185, 2009. 2
- [32] Y. Yang and D. Ramanan. Articulated pose estimation with flexible mixtures-of-parts. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1385–1392. IEEE, 2011. 6
- [33] S. Zagoruyko and N. Komodakis. Learning to compare image patches via convolutional neural networks. *CVPR*, 2015. 2
- [34] T. Zhou, Y. J. Lee, S. X. Yu, and A. A. Efros. Flowweb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In *CVPR*. IEEE, 2015. 2