

Shelf-Supervised Mesh Prediction in the Wild

Yufei Ye¹ Shubham Tulsiani² Abhinav Gupta^{1,2}
¹Carnegie Mellon University ²Facebook AI Research
 yufeiy2@cs.cmu.edu {shubhtuls, gabhinav}@fb.com

<https://judyye.github.io/ShSMesh/>



Figure 1. Given a single image, we predict a mesh with textures (rendered from the predicted view and a novel view). The models can learn directly from collections of images with only foreground masks, without supervision of mesh templates, multi-view association, camera poses, semantic annotations, etc.

Abstract

We aim to infer 3D shape and pose of object from a single image and propose a learning-based approach that can train from unstructured image collections, supervised by only segmentation outputs from off-the-shelf recognition systems (i.e. ‘shelf-supervised’). We first infer a volumetric representation in a canonical frame, along with the camera pose. We enforce the representation geometrically consistent with both appearance and masks, and also that the synthesized novel views are indistinguishable from image collections. The coarse volumetric prediction is then converted to a mesh-based representation, which is further refined in the predicted camera frame. These two steps allow both shape-pose factorization from image collections and per-instance reconstruction in finer details. We examine the method on both synthetic and the real-world datasets and demonstrate its scalability on 50 categories in the wild, an order of magnitude more classes than existing works.

1. Introduction

We live in a 3D world where 3D understanding plays a crucial role in our visual perception. Yet most computer vision systems in the wild still perform 2D semantic recognition (classification/detection). Why is that? We believe the key reason is the lack of 3D supervision in the wild. Most recent advances in 2D recognition have come from supervised learning but unlike 2D semantic tasks, obtaining supervision for 3D understanding is still not scalable.

While some recent approaches [10, 49] have attempted to build supervised 3D counterpart of 2D approaches, the concerns about scalability still remain. Instead, a more promising direction is to learn models of single image 3D reconstruction by minimizing the amount of manual supervision needed. Early approaches in this direction focused on using multi-view supervision [48, 53]. However, obtaining multiple views of the same objects/scene is still not easy for the data in the wild. Therefore, recent ap-

proaches [17, 23, 33] have attempted to learn single-image 3D reconstruction models from image collections. These approaches have targeted use of category templates, pose supervision and keypoints to provide supervision (See Table 1). However, such supervision still limits the scalability to hundreds of categories.

Our work is inspired by recent approaches that forgo supervision by exploiting meta-supervision from the category structure and geometric nature of the task. More specifically, the two common supervisions used are: (a) **rendering supervision** ([23, 26]): any given image of an instance in a category is merely a rendering of a 3D structure under a particular viewpoint. We can therefore enforce that the inferred 3D shape be consistent with the available image evidence when rendered; (b) **adversarial supervision** ([33]): in addition, the availability of an image collection also allows us to understand what renderings of 3D structures should look like in general. This enables us to derive supervisory signal not just from renderings of predictions in the input view, but also from novel views, by encouraging the novel-view renderings to look realistic. Prior work has exploited these supervisions but individually they pose several limitations for scaling 3D reconstruction models. For example, [11, 23] still requires template models. Similarly, [33] exploits the adversarial supervision and ignores the explicit geometric supervision. Therefore, such an approach only works on categories with strong structure and curated image collections.

This paper attempts to build upon the very recent successes in meta-supervision and provide an approach to scale learning of single image 3D reconstruction in the wild. We present a two-step approach: the first step relies on category-level understanding for coarse 3D inference (learned via meta-supervision). The second step specializes coarse models to match the details in the input image. Our approach can learn using only unannotated image collections, without requiring any ground-truth 3D [3, 10, 46], multi-view [44, 53], category templates [11, 22], or pose supervision [18, 48]. This not only allows our approach to infer accurate 3D, but also enables it to do so beyond the synthetic settings, using in-the-wild image collections in a ‘shelf-supervised’ manner: with only approximate instance segmentation masks obtained using off-the-shelf recognition systems as supervision. Yet our biggest contribution is the demonstration of scalability – we show results on order of magnitude more classes than existing papers.

2. Related Work

Supervised Single-view 3D Reconstruction. When considering object categories that exhibit limited shape variation, a common approach is to leverage predefined 3D deformable templates to either fit the input [3, 13, 36, 51] or to regress the model parameters [39, 40, 46]. While they

Table 1. Comparing ours to other image-based supervised works in terms of supervision and outputs.

	[15]	[22]	[17]	[18]	[44]	[50]	[33]	[8]	[26]	ours
pose	✓		✓	✓						
template		✓	✓	✓						
semantic			✓							(✓)
multi-view					✓					
mask	(✓)	✓	✓	✓	✓			✓	✓	✓
3D recon.		✓	✓	✓	✓	(✓)		✓	✓	✓
topology				✓				✓		✓
texture	✓		✓	✓	✓	✓	✓		✓	✓

achieve remarkable performance to reconstruct details, the predefined shape model may not be available for an arbitrary category or one template is not sufficient to fit all instances with varying topologies and shapes (*e.g.* chairs). Another line of works [1, 9, 35, 38, 49] learns a manifold of shape by first mapping the input to a latent space from which 3D shape is generated. However, these methods suffer from losing finer details as the reconstruction only rely on less expressive latent code. Similar to our approach, Gkioxari *et al.* [10] recently addressed these concerns by first inferring a coarse shape, and then refining it to match the given image. While all methods above have presented impressive results, they crucially require 3D supervision. In contrast, our work aims for similar inference where neither 3D nor pose annotation is available.

Unsupervised learning from image-based supervision.

With a similar motivation as ours to relax the need of supervision, several approaches study the reconstruction task with only multi-view or even single-view supervision. The key is to ensure reprojection consistency of the predicted 3D with available observations. While this relaxes the requirement for tedious 3D supervision, manual annotations are still required in different forms, such as semantic keypoints [17, 20], multi-view association [6, 29, 44, 48], categorical template [11, 17, 22, 23], or camera pose annotation [14, 15, 18, 29, 53]. Some recent works use self-supervised semantic co-part segmentation [26], foreground masks [8, 16], or symmetry [50] to further relax the manual annotation. Our work has similar setup while ours does not require semantic in training, and reconstructs textured full 3D meshes with various topology and shapes. Table 1 summarizes the differences of our method with others in terms of supervision and outputs.

Neural renderer for view synthesis and 3D representation.

Neural rendering [42] is a rapidly emerging field aiming to generate photo-realistic images or videos in a controllable manner by rendering its internal 3D representation with various geometric inspirations. In particular, many works [30, 31, 32, 37] present an approach to render a representation based on volumetric feature. While the initial applications are in multi-view settings, subsequent works [33, 38] have shown inference of such 3D representation from a single image. More relevant to our work, HoloGAN [33] presents a way to infer this from single single image in

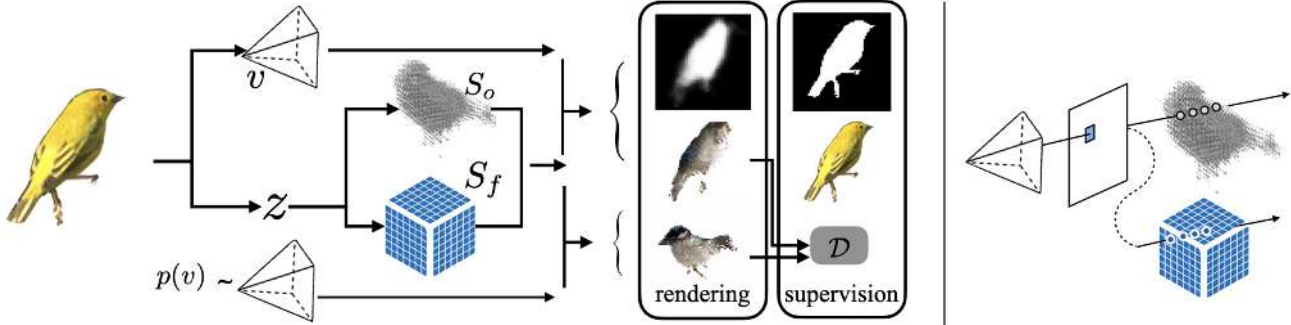


Figure 2. **Volumetric Representation Prediction and Rendering.** Left: Given an input image, the encoder-decoder network infers a semi-implicit volumetric representation (S_o, S_f) and a camera pose v . The semi-implicit volume is then projected from the predicted camera pose to obtain foreground image and mask. The semi-implicit volume is also projected from a novel view v' . The projections are required to both match the 2D observation and appear realistic. Right: Projection process mimicking ray marching.

an unsupervised setting by generating realistic novel views. However, their approach does not recover an explicit 3D representation. While we also adopt a neural renderer to render our volumetric representation, we explicitly incorporate reasoning about occupancy to improve 3D estimation.

3. Method

Our goal is to learn a model that, given an input image segmented with object mask, outputs a 3D shape in the form of a triangle mesh with texture and the corresponding camera pose. We use a two-step approach. First, we predict a canonical-frame volumetric representation and a camera pose to capture the coarse 3D structure which is consistent with categorical priors. We then convert this coarse volume to a memory-efficient mesh representation which is refined to better match the instance-level details.

We propose to learn the category level model from image collections using geometric and adversarial meta-supervision signals. More specifically, our key insight is that the projection of the predicted 3D should explain the observed images and masks, while also appearing realistic from a novel view.

3.1. Volumetric Reconstruction Model

First we define the volumetric reconstruction model which is learned separately for each category. Given an image this model predicts a volumetric representation in a canonical frame with corresponding camera pose. Note that unlike approaches that use a deformable category-level shape space, a volumetric representation allows us to capture larger shape and topology variations.

Concretely, we adopt a semi-implicit representation comprised of an explicit occupancy grid S_o , with an implicit 3D feature S_f , i.e. $S = (S_o, S_f)$. The latter can help capture appearance, texture, material, lighting, etc. This allows synthesizing both mask and appearance from a query view, and thereby lets us use both RGB images and foreground masks as supervision. The overall method is depicted in Figure 2.

Encoder-Decoder Architecture. We learn an encoder-decoder style network ϕ to predict this semi-implicit representation $(S, v) \equiv \phi(I)$ where v is the camera pose. The encoder maps the input image to a low-dimensional latent variable z and predict the camera pose, i.e. $(z, v) \equiv \phi_E(I)$. The latent variable z is then decoded to the volumetric representation, $S \equiv \phi_D(z)$. The key here is that the view-independent decoder learns to predict the shape in a canonical pose across all instances in the category. To further regularize the network, we leverage the observation that many objects exhibit reflection symmetry, and enforce a fixed symmetric plane ($x = 0$) via averaging predicted features in symmetrically related locations.

Volumetric Rendering. Our goal is to supervise the volumetric model using only 2D observations. Therefore, what we need is a rendering function (π) which projects volumetric representation to obtain 2D images and masks from a query view i.e. $(I, M) \equiv \pi(S_f, S_o, v)$. Similar to other volumetric neural renderers [30, 44], we use a geometrically informed projection process by mimicking ray marching.

For a given pixel \mathbf{p} , we use D samples along the ray to obtain a ‘rendered’ feature and mask value. Let us denote the coordinate of the d -th sample on the ray as $C_v + \lambda_d e_{\mathbf{p}}$ where $e_{\mathbf{p}}$ is the corresponding ray direction. We sample both implicit feature and occupancy at these locations, denoted as $S_f[C_v + \lambda_d e_{\mathbf{p}}]$ and $S_o[C_v + \lambda_d e_{\mathbf{p}}]$. We then composite these samples to obtain a per-pixel feature $s_f^{\mathbf{p}}$ and mask $s_m^{\mathbf{p}}$, by using the expected value with respect to ray stopping probability [43]:

$$s_f^{\mathbf{p}} = \sum_{d=1}^D (S_o[C_v + \lambda_d e_{\mathbf{p}}] \prod_{h=1}^{d-1} (1 - S_o[C_v + \lambda_h e_{\mathbf{p}}])) \cdot S_f[C_v + \lambda_d e_{\mathbf{p}}]$$

The pixelwise mask value $s_m^{\mathbf{p}}$ is similarly rendered by setting S_f to constant 1. While this process lets us directly compute the rendered mask M , we use a few upconvolutional layers to transform the rendered 2D feature image to

the output color image.

Training. We supervise this network with only unannotated images and foreground masks. We use three different kinds of supervision (or terms in the loss function):

Pixel consistency loss. Our first term is the simplest one. Any predicted volumetric representation when rendered in the same camera view should explain the input image and mask. This is performed in color space, mask space [28], and perceptual space [54].

$$\begin{aligned} \mathcal{L}_{rgb} &= \|\hat{I} - I\|_1 \\ \mathcal{L}_{mask} &= 1 - \frac{\|\hat{M} \otimes M\|_1}{\|\hat{M} \oplus M - \hat{M} \otimes M\|_1} \\ \mathcal{L}_{perc} &= \|h(\hat{I}) - h(I)\|_2^2 \end{aligned}$$

where \hat{I}, \hat{M} are rendered image and mask; h is the feature extracted by a pretrained AlexNet [21] and \oplus, \otimes are element-wise summation and multiplication respectively.

View synthesis adversarial loss. A degenerate solution could arise such that the shape is only plausible from the predicted view. To avoid it, we require the projection of predicted shape to appear realistic from a random novel view. Specifically, we sample another camera pose from a fixed prior to render the novel view: $I' = \pi(\phi_D(z), v'), v' \sim p(v)$. We then present this generated image to an adversarial discriminator with an objective to fool it. We similarly encourage photo-realism when rendering from the predicted camera pose. The loss is minimized in a vanilla GAN [12] scheme.

$$L_{adv} = \log \mathcal{D}(I) + \log(1 - \mathcal{D}(\pi(S, v))) + \log(1 - \mathcal{D}(\pi(S, v')))$$

Content consistency loss. To further regularize the network we build on a insight that the encoder and decoder networks should be self-consistent. Given a synthesized image from the decoder, the encoder should predict the actual content (latent variable with camera pose) that generated that image. Formally,

$$\begin{aligned} \mathcal{L}_{content} &= \|\phi_E(\pi_S(S, v)) - (z, v)\|_2^2 \\ &\quad + \|\phi_E(\pi_S(S, v')) - (z, v')\|_2^2 \end{aligned}$$

Empirically, we found $\mathcal{L}_{content}$ important to stabilize training.

Optimization. The neural renderer and decoder are trained to minimize all of the above objectives. But the encoder is not optimized with the adversarial loss, as in [25].

3.2. Instance-Level Specialization

The volumetric representation captures general category-level structure to hallucinate a full 3D shape. However, this shape is coarse as it is: a) limited by the volumetric resolution, and b) generated only from a

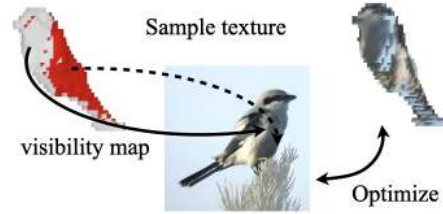


Figure 3. **Instance-level specialization.** We convert the inferred volumetric occupancy to an initial mesh. The mesh geometry and textures are then iteratively refined to better match the given input.

low-dimensional latent variable. On the other hand, a mesh representation is more flexible and can allow capturing the finer shape details. We therefore go beyond this coarse volumetric prediction, and capture details specific to this instance by converting the volume to an initial mesh, which is then adjusted to better match the input image.

Volumetric to Mesh. We first obtain an initial mesh from the predicted volumetric occupancy. This is done similar to Mesh-RCNN [10] by binarizing the occupancy grid S_o and extracting its surfaces. Next, every vertex is projected to the image to obtain visibility and texture at the vertex. At this step we leverage the symmetry of the mesh to fuse the textures from its reflective symmetric vertex. The final associated texture for each vertex is an average of itself and its visible symmetric neighbors.

Mesh refinement. We optimize the geometry and refine the texture of the mesh iteratively. Given a posed textured input mesh, we first optimize the vertex location and the camera pose such that the projection of the mesh matches the observation. After every step of mesh geometry update, vertex textures are re-sampled from the image given the adjusted projected location. More specifically, we use a mesh-based differential renderer [28] to project and render. The rendered images and masks $(\hat{I}, \hat{M}) \equiv \pi_G(G, v)$ are encouraged to be consistent with the input image and foreground mask. We regularize the optimization by penalizing large vertex displacement $\|\delta X\|_2^2$ and encourage Laplacian smoothness $\|\Delta X\|_2^2$.

4. Experiments

Our goal is to highlight how our approach learns to predict 3D meshes from image collections in the wild. Specifically, we show 3D reconstruction for 50 object categories from OpenImages dataset [24]. Note that this diverse set of reconstructions is an order of magnitude larger than those of any existing approaches.

However, there is no ground truth for OpenImages. Also, most baseline approaches fail to work on uncurated image collections. In order to provide comparisons, we perform two additional experiments. First, we compare on data drawn from 3D Warehouse[47], using rendered images as image collection. Using synthetic data allows us to provide




Figure 4. Visualization of categorical volumetric representation across different methods.

Table 2. Quantitative results (3D IoU / F-score) on synthetic data comparing different methods for shape reconstruction.

	airplane	car	chair
HoloGAN [33]	0.28 / 0.31	0.43 / 0.44	0.26 / 0.25
PrGAN [8]	0.29 / 0.18	0.48 / 0.37	0.28 / 0.28
Ours	0.33 / 0.46	0.55 / 0.43	0.31 / 0.29
Ours (refined)	— / 0.49	— / 0.42	— / 0.31

quantitative evaluation and perform ablative analysis. Second, we also compare with some of the other curated common datasets used in the literature (CUB, Chair-in-Wild, ImageNet Quadrupeds). This helps us to qualitatively compare with some baseline approaches.

4.1. Synthetic Data

We first evaluate our method on models from 3D Warehouse [47], using the subset recommended by Chang *et al.* [4]. We select three categories which are commonly used to evaluate single-view reconstruction: aeroplane, car, and chair. Note that within a category, the shapes across instances can have a large variation and even different topology, especially for chairs. Each 3D model is rendered from 20 views, with uniformly sampled azimuth $[0^\circ, 360^\circ]$ and elevation elevation $[-60^\circ, 60^\circ]$. However, the network is not provided with multi-view associations in training.

Evaluation metrics. We report 3D IoU with resolution 32^3

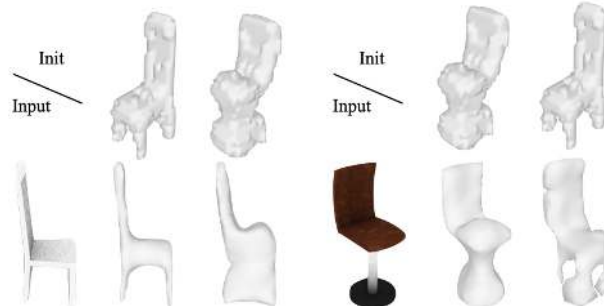


Figure 5. Ablation study: refining mesh initialized from the predicted volume (col 2/5) and another volumes (col 3/6).

Table 3. Quantitative results (3D IoU) on synthetic data to ablate the effect of each loss term.

	airplane	car	chair
Ours	0.33	0.55	0.31
Ours $-\mathcal{L}_{adv}$	0.25	0.44	0.22
Ours $-\mathcal{L}_{cont}$	0.24	0.54	0.23

and F-score in the canonical frame for volumetric reconstruction and report F-score [41] for mesh refinement. The F-score can be interpreted as the percentage of correctly reconstructed surface. As our predictions (and those of baselines) can be in an arbitrary canonical frame that is different from the ground truth frame, we explicitly search for azimuth, elevation for each instance and binarizing threshold for each category to align the predicted canonical space with the ground-truth.

Baselines. We compare our approach to [8, 33]. We adapt HoloGAN [33] by training their system on our data, and obtaining a 3D output by adding a read-off function from the learned volumetric feature to occupancy by enforcing the reprojection consistency with foreground masks. We implement PrGAN [8] using our encoder-decoder network. Our implementation provides a boost to original PrGAN.

Figure 4 visualizes the reconstructions in a canonical frame on 3 categories. HoloGAN is able to reconstruct a blobby shape, but as it does not explicitly represent 3D occupancies, it struggles to generate a coherent shape. PrGAN is able to capture the coarse shape layout but it lacks some details like flat body of aeroplanes. In contrast, we reconstruct the shape more faithfully to the ground-truth as we leverage information from both appearance and foreground masks. Quantitatively, we report the 3D IoU on these categories in Table 2 and, consistent with the qualitative results, observe empirical gains across all categories.

Mesh Refinement. Table 2 also reports the evaluation of the mesh refinement stage. Compared with the initial meshes converted from volumetric representation, our specialized meshes match the true shape better.

In Figure 5, we visualize the refinement results with an interesting ablation to further highlight the importance of mesh initialization. Instead of initializing with our pre-

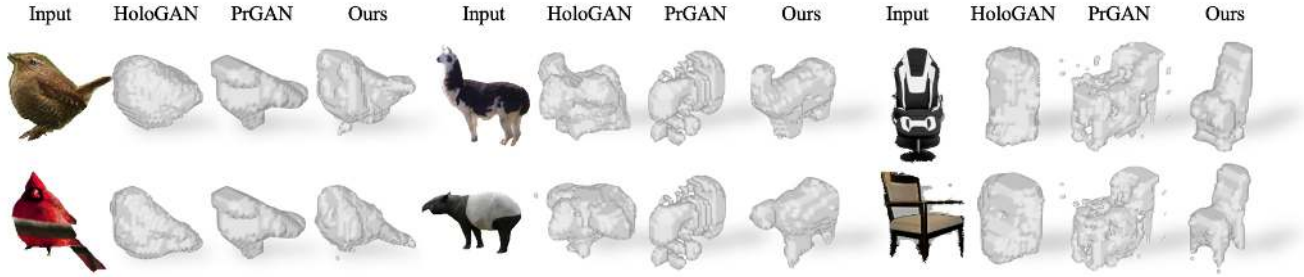


Figure 6. Visualizing categorical volumetric representation across different methods on CUB-200-2011, Quadrupeds, Chairs in the wild.

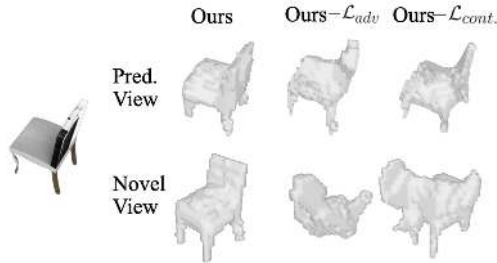


Figure 7. Ablation study: comparing reconstructed volumes when the model disables different loss terms.

dicted volume, we initialize the mesh from another chair consisting of different numbers of chair legs. The refinement fails to specialize well. This indicates that the meshes for all instances cannot be adjusted from one single shape especially when shapes have a large variance, and that our volumetric prediction, though coarse, provides an important initialization for the instance-level refinement.

Loss ablation. We provide quantitative (Table 3) and qualitative (Figure 7) results to show each loss term is necessary. Without adversarial loss, the model collapses to generate shapes only looking similar to the input from the predicted view. It does not even look like a chair from another view, since this degenerate solution is not penalized by other losses. Without content loss, the performance also drops, especially on categories with larger shape variance like chairs. The consistency loss is not ablated because it is needed for the task of reconstruction.

4.2. Curated Collections

We also examine our method on three real-world datasets that have been curated and used in the literature for the 3D reconstruction problem:

CUB-200-2011 [45]: The CUB dataset consists of 6k images of 200 bird species with annotated foreground masks.

Quadrupeds from ImageNet [7]: The Quadrupeds dataset consists of 25k images of different quadrupeds from ImageNet. Masks are provided by Kulkarni *et al.* [22], who use an off-the-shelf segmentation system [19] and manually filter out the truncated or noisy instances. Quadrupeds consists of multiple 4-legged animal species including buffalo, camels, sheep, dogs, etc. The animals also exhibit rich ar-

ticulation *e.g.* running, lying, heads up or heads down. This makes the underlying shape variance significantly larger than the CUB dataset.

Chairs in the wild [7, 34, 52]: For chairs in the wild, we combine chairs in PASCAL3D, ImageNet, and Stanford Online Products Dataset to get 2084 images for training and 271 for testing. Masks in [7, 34] are from segmentation systems [5, 19] and those in [52] are from annotations.

Figure 6 qualitatively compares our volumetric reconstruction to baselines on the 3 separate real-world datasets. Similar to results on synthetic data, HoloGAN reconstructs only coarse blobby volumes as it does not explicitly consider occupancy or geometric-informed projection. PrGAN collapses to shapes with little variance, since it does not use appearance cues. But the real datasets have noisier foreground masks and textures contain more information. In contrast, we are able to learn the coarse categorical shape just from the foreground images. Our reconstructions also capture subtle differences like the length of bird tails, articulated heads of the quadrupeds, the style of chairs.

Mesh Refinement. Figure 8 visualizes our refined meshes from the predicted view and a novel view on these three dataset. We observe that we predict meaningful texture even for invisible regions and that the shape of the mesh also looks plausible from another view. On CUB-200-2011, our method captures the categorical shapes like blobby bodies, beaks and tails while captures subtle shape differences between birds such as the tail length, body width, neck bending, etc. On Quadrupeds, we are able to capture quadrupeds common traits such as torso with one head and front back legs. We can also depict their uniqueness such as the camel hump and longer legs, the tapir having stout neck, the sheep raising up its heads, the horse bending down its neck, *etc.* On Chairs in the wild, the learned common model differentiates one-leg and four-leg chairs respectively. The four legs and seat can be hallucinated even when occluded. The subtle differences such as a wide or a narrow chair back are also captured. Despite the challenges in the datasets, it is encouraging that our model can capture both, the common shapes and specialized details just by learning from these unannotated image collections.




Figure 8. Visualizing our refined shapes from the predicted view (2nd column in each quadruplet) and a novel view (3rd+4th / quadruplet) on CUB, Quadrupeds, Chairs in the wild. We are able to capture both the shared shapes in categories and instance-specific differences.

4.3. OpenImages 50 Categories

Finally, the highlight of our model is the ability to scale to images in the wild. We evaluate our model on 50 categories on Open Images including bagel, water tap, hat, *etc.* The size of each category ranges from 500 to 20k. The foreground masks [2] are from annotation and filtered by a fine-tuned occlusion classifier. Figure 10 visualizes the reconstructed meshes from the predicted view and a novel view. Our method works on a large number of categories, including thin (water taps, saxophone), flat (wheels, surfboards), blobby structures (Christmas trees, vases). We are able to reconstruct shapes with various topology such as bagels, mugs, handbags. The model captures the categorical shapes shared within classes and hallucinates plausible occluded regions (mushroom, mugs). We can also capture details at instance-level, such as the number of wheels of roller-skaters, styles of high-heels, hats, *etc.*

Integrating on COCO. We additionally show results of our models on COCO [27] without fine-tuning (Figure 9). We first detect and segment the objects with off-the-shelf segmentation [19] system. Based on the predicted classes, we then pass the segmented objects to our category-specific models which are trained on previous datasets. Despite more cluttered scenes and the dataset domain shift, our models can lift the 2D detection to 3D meshes for various categories while preserving instance details.

5. Conclusion

We presented an approach to predict 3D representations from unannotated images by learning a category-level volumetric prediction followed by instance-level mesh specialization. We found that both are important to infer an accurate 3D reconstruction. While we obtained encouraging results across diverse categories, our approach has several limitations. For example, our rendering model is simplistic and not incorporate lighting during rendering. Thus we cannot easily reason about concave structure. Additionally,

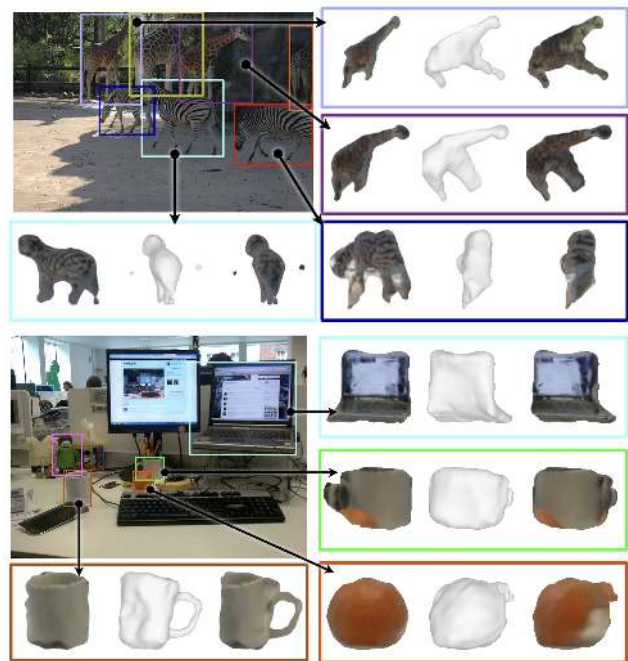


Figure 9. Test on COCO: visualization of lifting detection results to meshes via the shelf-supervised models.

while we only examined setups without annotated supervision like mesh templates, our system could potentially incorporate additional (sparse) supervision to improve the reconstruction quality. While these challenges still remain, we believe our work on inferring accurate reconstruction with limited supervising can provide a scalable basis towards the goal of reconstructing generic objects in the wild.

Acknowledgements. The authors would like to thank Nilesh Kulkarni for providing segmentation masks of Quadrupeds. We'd also thank Chen-Hsuan Lin, Chaoyang Wang, Nathaniel Chodosh and Jason Zhang for fruitful discussion and detailed feedback on manuscript. Carnegie Mellon Effort has been supported by DARPA MCS, DARPA SAIL-ON, ONR MURI and ONR YIP.



Figure 10. Visualizing our reconstructed meshes from the predicted view (2nd in each quadruplet) and a novel view (3rd and 4th in each quadruplet) trained on multiple categories on Open Images.

References

- [1] Ijaz Akhter and Michael J Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, 2015.
- [2] Rodrigo Benenson, Stefan Popov, and Vittorio Ferrari. Large-scale interactive object segmentation with human annotators. In *CVPR*, 2019.
- [3] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*, 2016.
- [4] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *TPAMI*, 2017.
- [6] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *ECCV*, 2016.
- [7] 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.
- [8] Matheus Gadelha, Subhransu Maji, and Rui Wang. 3d shape induction from 2d views of multiple objects. In *3DV*, 2017.
- [9] Rohit Girdhar, David F Fouhey, Mikel Rodriguez, and Abhinav Gupta. Learning a predictable and generative vector representation for objects. In *ECCV*, 2016.
- [10] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *ICCV*, 2019.
- [11] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. *ECCV*, 2020.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- [13] Marc Habermann, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Livecap: Real-time human performance capture from monocular video. *TOG*, 2019.
- [14] Paul Henderson and Vittorio Ferrari. Learning single-image 3d reconstruction by generative modelling of shape, pose and shading. *IJCV*, 2019.
- [15] Paul Henderson, Vagia Tsiminaki, and Christoph H Lampert. Leveraging 2d data to learn textured 3d mesh generation. In *CVPR*, 2020.
- [16] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Escaping plato’s cave: 3d shape from adversarial rendering. In *ICCV*, 2019.
- [17] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018.
- [18] Hiroharu Kato and Tatsuya Harada. Learning view priors for single-view 3d reconstruction. In *CVPR*, 2019.
- [19] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. *CVPR*, 2020.
- [20] Chen Kong and Simon Lucey. Deep non-rigid structure from motion. In *ICCV*, 2019.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [22] Nilesh Kulkarni, Abhinav Gupta, David Fouhey, and Shubham Tulsiani. Articulation-aware canonical surface mapping. In *CVPR*, 2020.
- [23] Nilesh Kulkarni, Abhinav Gupta, and Shubham Tulsiani. Canonical surface mapping via geometric cycle consistency. In *ICCV*, 2019.
- [24] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.
- [25] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. *ICML*, 2015.
- [26] Xueting Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Self-supervised single-view 3d reconstruction via semantic consistency. *ECCV*, 2020.
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [28] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *ICCV*, 2019.
- [29] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3d supervision. In *NeurIPS*, 2019.
- [30] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *TOG*, 2019.
- [31] Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rerendering in the wild. In *CVPR*, 2019.
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *ECCV*, 2020.
- [33] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *ICCV*, 2019.
- [34] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *CVPR*, 2016.
- [35] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019.
- [36] Bojan Pepik, Michael Stark, Peter Gehler, and Bernt Schiele. Teaching 3d geometry to deformable part models. In *CVPR*,

- 2012.
- [37] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *CVPR*, 2019.
 - [38] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *NeurIPS*, 2019.
 - [39] Edward J Smith, Scott Fujimoto, Adriana Romero, and David Meger. Geometrics: Exploiting geometric structure for graph-encoded objects. *ICML*, 2019.
 - [40] Hao Su, Charles R Qi, Yangyan Li, and Leonidas J Guibas. Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In *ICCV*, 2015.
 - [41] Maxim Tatarchenko, Stephan R Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, 2019.
 - [42] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. *CGF*, 2020.
 - [43] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *CVPR*, 2017.
 - [44] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. *TPAMI*, 2019.
 - [45] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
 - [46] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018.
 - [47] 3D Warehouse. <https://3dwarehouse.sketchup.com/>. Accessed: 2019-11-15.
 - [48] Olivia Wiles and Andrew Zisserman. Silnet: Single- and multi-view reconstruction by learning from silhouettes. *BMVC*, 2017.
 - [49] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *NeurIPS*, 2016.
 - [50] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *CVPR*, 2020.
 - [51] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *CVPR*, 2019.
 - [52] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *WACV*, 2014.
 - [53] 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.
 - [54] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.