

Hi-LASSIE: High-Fidelity Articulated Shape and Skeleton Discovery from Sparse Image Ensemble

Chun-Han Yao^{1*} Wei-Chih Hung² Yuanzhen Li³ Michael Rubinstein³
 Ming-Hsuan Yang^{1,3,4} Varun Jampani³

¹UC Merced ²Waymo ³Google Research ⁴Yonsei University

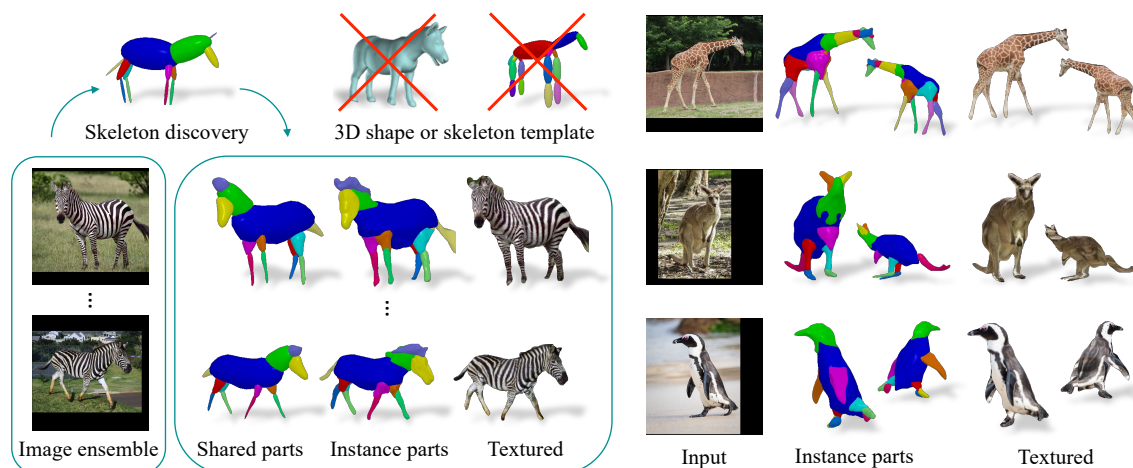


Figure 1. **Hi-LASSIE overview and sample reconstructions.** Given 20-30 images of an articulated animal class, we first discover a generic 3D skeleton, then jointly optimize the camera viewpoints, skeleton articulations, as well as shared and instance-specific neural part shapes. Hi-LASSIE is able to produce high-fidelity shapes and texture without any pre-defined shape model or 3D skeleton annotations. The part-based representation also allows applications like animation and motion re-targeting.

Abstract

Automatically estimating 3D skeleton, shape, camera viewpoints, and part articulation from sparse in-the-wild image ensembles is a severely under-constrained and challenging problem. Most prior methods rely on large-scale image datasets, dense temporal correspondence, or human annotations like camera pose, 2D keypoints, and shape templates. We propose Hi-LASSIE, which performs 3D articulated reconstruction from only 20-30 online images in the wild without any user-defined shape or skeleton templates. We follow the recent work of LASSIE that tackles a similar problem setting and make two significant advances. First, instead of relying on a manually annotated 3D skeleton, we automatically estimate a class-specific skeleton from the selected reference image. Second, we improve the shape reconstructions with novel instance-specific optimization strategies that allow reconstructions to faithfully fit on each instance while preserving the class-specific priors learned

across all images. Experiments on in-the-wild image ensembles show that Hi-LASSIE obtains higher fidelity state-of-the-art 3D reconstructions despite requiring minimum user input. Project page: chhankyao.github.io/hi-lassie/

1. Introduction

3D assets of articulated animals enable numerous applications in games, movies, AR/VR, etc. However, building high-fidelity 3D models of articulated shapes like animal bodies is labor intensive either via manual creation or 3D scanning. Recent advances in deep learning and 3D representations have significantly improved the quality of 3D reconstruction from images. Much of the success depends on the availability of either rich 3D annotations or multi-view captures, both of which are not always available in a real-world scenario. A more practical and scalable alternative is automatic 3D reconstruction from online images as it is straightforward to obtain image ensembles of any animal category (e.g., image search results). In this work,

*Work done as a student researcher at Google.




Figure 2. **User inputs across different techniques for articulated animal reconstruction.** Contrary to prior methods that leverage detailed 3D shapes or skeletons, Hi-LASSIE only requires the user to select a reference image where most animal body parts are visible.

we tackle a practical problem setting introduced in a recent work, LASSIE [38], where the aim is to automatically estimate articulated 3D shapes with only a few (20-30) in-the-wild images of an animal species, without any image-level 2D or 3D annotations.

This problem is highly under-constrained and challenging due to a multitude of variations within image ensembles. In-the-wild images usually have diverse backgrounds, lighting, and camera viewpoints. Moreover, different animal instances can have distinct 2D appearances due to pose articulations, shape variations, and surface texture variations (skin colors, patterns, and lighting). As shown in Fig. 2, early approaches in this space try to simplify the problem with some user-defined 3D templates, hurting the generalization of those techniques to classes where such templates are not always readily available. In addition, most of these methods (except LASSIE [38]) assume either large-scale training images of each animal species [16, 17, 43] or per-image 2D keypoint annotations [44], which also limits their scalability.

In this work, we propose a more practical technique that does not require any 3D shape or skeleton templates. Instead, as illustrated in Fig. 2, the user simply has to select a reference image in the ensemble where all the animal parts are visible. We achieve this by providing two key technical advances over LASSIE: 1) 3D skeleton discovery and 2) instance-specific optimization strategies. Our overall framework, named Hi-LASSIE, can produce **H**igher-fidelity articulated shapes than **L**ASSIE [38] while requiring minimal human input. Our key insight is to exploit the 2D part-level correspondences for 3D skeleton discovery. Recent works [1, 32] observe that the deep features extracted from a self-supervised vision transformer (ViT) [11] like DINO-ViT [7] can provide good co-part segmentation across images. We further exploit such features to reason about part visibility and their symmetry. At a high level, we first obtain a 2D skeleton using the animal silhouette and part clusters [1] obtained from DINO features [7]. We then uplift this 2D skeleton into 3D using symmetric part information that is present in the deep DINO features. Fig. 1 shows the skeleton for zebra images discov-

ered by Hi-LASSIE. Similar to LASSIE [38], we leverage 3D part priors (learned from and shared across instances) and the discovered 3D skeleton to regularize the articulated shape learning. Furthermore, we design three novel modules to increase the quality of output shapes: 1) High-resolution optimization by zooming in on individual parts, 2) Surface feature MLPs to densely supervise the neural part surface learning, and 3) Frequency-based decomposition of part surfaces for shared and instance-specific components. Note that Hi-LASSIE can generalize to diverse animal species easily as it does not require any image annotations or category-specific templates.

We conduct extensive experiments on the Pascal-Part [8] and LASSIE [38] image ensembles, which contain in-the-wild images of various animal species like horse, elephant, and penguin. Compared with LASSIE and other baselines, we achieve higher reconstruction accuracy in terms of key-point transfer, part transfer, and 2D IOU metrics. Qualitatively, Hi-LASSIE reconstructions show considerable improvement on 3D geometric and texture details as well as faithfulness to input images. Finally, we demonstrate several applications like animation and motion re-targeting enabled by our 3D part representation. Fig. 1 (right) shows some Hi-LASSIE 3D reconstructions for different animal species. The main contributions of this work are:

- To our best knowledge, Hi-LASSIE is the first approach to discover 3D skeletons of articulated animal bodies from in-the-wild image ensembles without using any image-level annotations. We show that the discovered 3D skeleton can faithfully fit all instances in the same class and effectively regularize the 3D shape optimization.
- Hi-LASSIE includes several novel optimization strategies that makes the output shapes richer in 3D details and more faithful to each image instance than prior methods.
- Extensive results on multiple animal classes and datasets demonstrate the state-of-the-art performance of Hi-LASSIE while requiring less user inputs than prior works.

2. Related Work

Animal pose and shape estimation. 3D pose and shape estimation of animal bodies from in-the-wild images is quite challenging considering the diverse 2D appearance

across different instances, viewpoints, and articulations, among others. Most mesh reconstruction methods [13, 17, 20, 30, 39] mainly deal with objects with simple or rigid shapes (*e.g.*, birds and cars), which cannot be applied to the highly-articulated animal bodies. Recent articulation-aware approaches address a different (often simplified) scenario by leveraging a pre-defined statistical model [28, 43–45], 3D shape/skeleton templates [16, 38], or dense temporal correspondence in videos [34–36]. Without assuming any of these annotations/data, Hi-LASSIE introduces a novel and practical framework to estimate high-fidelity animal bodies by discovering 3D skeleton from a reference image.

Skeleton extraction and part discovery. 2D skeleton/outline extraction has been widely studied and used as geometric context for shape/pattern recognition [2, 3, 14, 29, 33]. However, these methods often fail to identify separate skeleton bones/parts when they overlap or self-occlude each other in an image. In this work, we propose to lift 2D skeletons to 3D for articulated shape learning by jointly considering the geometric and semantic cues in 2D images. For part discovery, deep feature factorization (DFF) [10] and follow-up works [1, 9, 15, 18] show that one could automatically obtain 2D corresponding part segments by clustering deep semantic features. In the 3D domain, the object parts can be discovered by using volumetric cuboids [31], clustering 3D point clouds [23, 24, 26, 27], or learning part prior [37]. These methods mainly assume some form of supervision like 3D shapes or camera viewpoints. In this paper, we discover 3D parts of articulated animals from image ensembles without any 2D or 3D annotations/templates, which, to the best of our knowledge, is unexplored in the literature.

3D reconstruction from sparse images. Optimizing a 3D scene or object from multi-view images is a fundamental problem in computer vision. The majority of recent breakthroughs are based on the powerful Neural Radiance Field (NeRF) [25] representation. Given a set of multi-view images, NeRF learns a neural volume from which one can render high-quality novel views. NeRS [40] introduces a neural surface representation to learn compact 3D shapes from a sparse image collection, which we find suitable to model individual animal body parts. Another line of work [4, 6, 41] propose neural reflectance decomposition on image collections captured in varying illuminations, but they operate on rigid objects with ground-truth segmentation and known camera poses. SAMURAI [5] jointly reasons about camera pose, shape, and materials from image collections of a single rigid object. In contrast, our input consists of in-the-wild animal images with varying textures, viewpoints, and pose articulations captured in different environments.

3. Approach

Given a sparse image ensemble of an animal species, Hi-LASSIE first discovers a class-specific 3D skeleton that

specifies the initial 3D joint coordinates and part connectivity. Then, it jointly optimizes the camera viewpoint, pose articulation, and part shapes for each instance. Before describing our approach, we briefly review the LASSIE [38] method, where several parts are adopted in this work.

3.1. Preliminaries: 3D Skeleton & Parts in LASSIE

Given a user-provided 3D skeleton template that specifies the 3D joints and bones, LASSIE [38] represents an overall articulated shape by assembling several 3D parts. In particular, each part is defined as a deformable neural surface [40] wrapped around a skeleton bone. The part surfaces are parameterized by multi-layer perceptron networks (MLPs) which predict the deformation of any 3D point on a surface template (*e.g.* unit sphere). Formally, let $X \in \mathbb{R}^{3 \times m}$ denote the m uniformly sampled 3D points on a unit sphere, one can deform the 3D surface of the i -th part in the canonical space using the part MLP as $X \mapsto \mathcal{F}_i(X)$. The part surface is then rigidly transformed using the optimized 3D skeleton with scaling $s_i \in \mathbb{R}$, rotation $R_i \in \mathbb{R}^{3 \times 3}$, and translation $t_i \in \mathbb{R}^3$. The final part surface points V_i in the global coordinate frame can be expressed as:

$$V_i = s_i R_i \mathcal{F}_i(X) + t_i. \quad (1)$$

The rigid transformation of each part is defined by its corresponding bone length (scaling), orientation (rotation), and centroid (translation). This skeleton-based representation ensures that the connectivity of 3D parts under arbitrary articulations and easy to repose. Compared to explicit mesh representation, the neural surfaces also enable efficient mesh regularization while producing high-resolution surfaces since the MLPs take continuous surface coordinates as input. A key innovation of LASSIE is the use of part prior within the part MLPs $\{\mathcal{F}_i\}$ that regularizes the part shapes to be close to convex primitive geometric shapes like spheres, cones, etc. Please refer to [38] for further details. To take advantage of these properties, we adopt a similar part-based representation in this work. A key difference from LASSIE is that we alleviate the need for user-provided 3D skeleton template and instead discover the 3D skeleton automatically from a reference image chosen from the ensemble. In addition, we propose several improvements to the neural surfaces and the optimization process to address the limitations of LASSIE resulting in higher-fidelity reconstructions.

3.2. Discovering 3D Skeleton

2D skeleton extraction. Fig. 3 illustrates our 3D skeleton discovery process from a reference image. We first extract a set of 2D skeleton points from a reference silhouette using a skeletonize/morphological thinning algorithm [42]. Specifically, we iteratively remove pixels from the borders until none can be removed without altering the connectivity. These 2D points can be seen as candidates of under-




Figure 3. **3D Skeleton Discovery.** Given a reference image and its rough silhouette, we first extract, filter, and connect the 2D candidate points to form a 2D skeleton graph. Then, we find and split the symmetric parts by leveraging the 2D geometry and semantic cues. Finally, we design a simple heuristic to initialize the 3D joints, part shapes, and surface features.

lying skeleton joints and bones through camera projection. Similar to prior 2D shape matching methods [2, 33], we categorize the skeleton points into junctions, endpoints, and connection points. Assuming that each skeleton curve is one pixel wide, a skeleton point is called an *endpoint* if it has only one adjacent point; a *junction* if it has three or more adjacent points; and a *connection point* if it is neither an endpoint nor a junction. We sort the junctions and endpoints by distance transform (2D distance to closest border) and identify the ‘root’ junction with the largest distance to the border. Next, we find the shortest path (sequence of skeleton points) between root and each endpoint by traversing through the graph. From these paths, we can build a skeleton tree with 2D joints (junctions or endpoints) and bones (junction-junction connection or junction-endpoint connection). Since the silhouettes and skeleton points are sometimes noisy, we filter those noisy 2D joints using non-maximal suppression. That is, we remove a point if it lies within the coverage (radius calculated by distance transform) of its parent joint.

3D joint and part initialization. From the extracted 2D skeleton, one can roughly initialize the simple part shapes and surface features using the reference image. However, all the joints and bones would be initialized on a 2D plane (image plane) which is not sufficient for later 3D shape and camera pose optimization across instances. To obtain a better 3D skeleton initialization, we propose to find symmetric parts and separate them in the 3D space w.r.t. the symmetry plane. Our key insight is that symmetric parts (*e.g.* left and right legs/ears) share similar geometric and semantic features. Therefore, we design a heuristic to calculate the symmetry score of each pair of joints/bones based on their geometric and semantic feature distance. Specifically, we compute the following features for each joint/bone: length, average radius, and the average DINO features [7] along the paths to their common ancestor in the skeleton tree. Then, we identify pairs of joints that share similar features, which

usually correspond to the symmetric animal parts (*e.g.*, left and right leg). To uplift our 2D skeleton onto 3D, we set the z -coordinates of the each joint pairs to be on opposite sides of the symmetry plane ($z = 0$) and offset by their average radius (so that the initial 3D parts do not overlap). To better deal with overlapping parts in 2D silhouettes, we also split the parent of two symmetric parts if it has only two children. See Fig. 3 (right) for an example of estimated 3D skeleton.

3.3. Learning High-fidelity Articulated Shapes

Optimization setting. Our optimization framework takes a sparse set of n (typically 20-30) in-the-wild images $\{I_j\}_{j=1}^n$ as well as the discovered 3D skeleton $P \in \mathbb{R}^{p \times 3}$ with p joints and b bones/parts as input. The only image-level annotations like silhouettes and semantic clusters are obtained from the clustering results of self-supervisory DINO [7] features. All the images in an ensemble contain instances of the same species, but each instance could vary in pose, shape, texture, camera viewpoint, etc. We assume that each animal body is mostly visible in an image without severe truncation or occlusions by other objects although self-occlusion is allowed. We use j to denote the index over images $j \in \{1, \dots, n\}$ and i the index over parts $i \in \{1, \dots, b\}$. For each instance, Hi-LASSIE optimizes the camera viewpoint $\pi^j = (R_0, t_0)$, part rotations $R^j \in \mathbb{R}^{b \times 3 \times 3}$, and neural part deformation MLPs. The overall framework is shown in Fig. 4, where we first perform pose articulation then progressively add more geometric details (from the shared part shapes to instance-specific deformations).

Per-instance deformation by frequency decomposition. Unlike LASSIE where the part shapes are shared across all instances, which limits the shape fidelity to input images, we propose to learn instance-specific part shapes that can better account for instance-varying and articulation-dependent deformations. Intuitively, instances of the same species should share similar base part shapes and only deviate in high-frequency details (*e.g.* ears and tail). Such detailed difference is usually caused by part articulations or

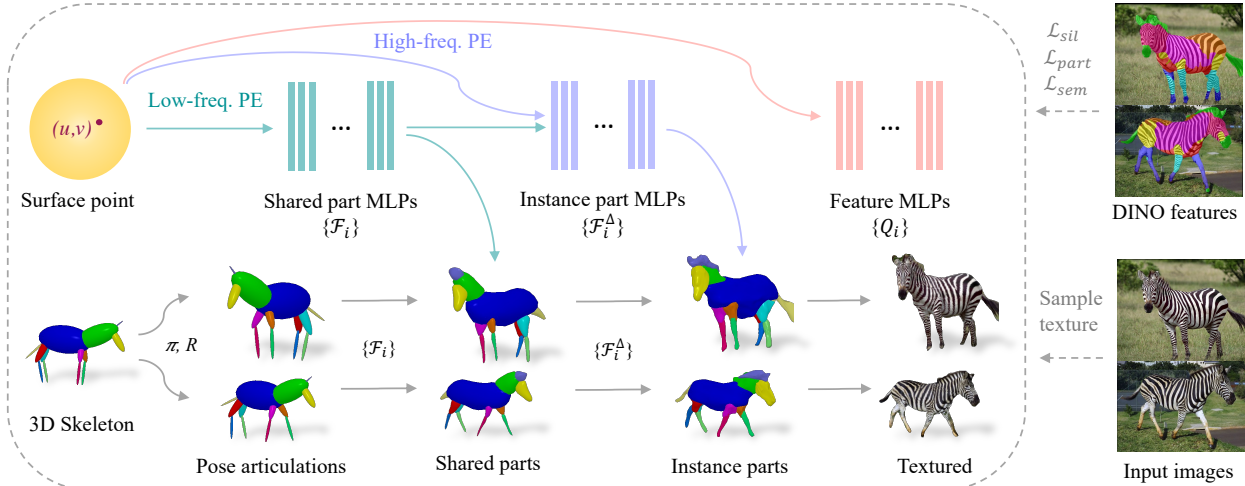


Figure 4. **Hi-LASSIE optimization framework.** Based on the discovered 3D skeleton, we reconstruct an articulated shape by optimizing the camera viewpoints, pose articulations, and part shapes. We represent the 3D part shapes as neural surfaces, which are decomposed into shared (low-frequency) and instance-specific (high-frequency) components via positional encoding (PE) of input surface coordinates. The only image-level annotation is from the self-supervisory DINO features, and the surface texture is sampled from the input images.




Figure 5. **Frequency-decomposed part MLP for per-instance shape deformation.** We design a MLP network to represent the neural surfaces composed of varying amount of details. By increasing the frequencies of input positional encoding, the outputs of deeper layers can express more detailed deformation. The output and hidden layers are linear layers.

instance variance. We observe that naively fine-tuning the part surfaces on one instance tends to overfit to that instance and nullify the 3D shape prior learned across all instances. For instance, the part shapes can overfit to the target silhouette (input view) but look unrealistic in novel views. To address this, we propose to decompose the surface deformations in the spatial frequency space so that each instance can have high-frequency detail variations while sharing the same low-frequency base part shapes. As shown in Fig. 5, we achieve this by designing a part deformation MLP with frequency-decomposed input and output components. In a forward pass through the MLP, we first apply positional encoding (PE) to an input coordinate $x \in \mathbb{R}^3$ using sines with varying frequencies ω and phases ϕ : $\text{PE}_i(x)$

$= \sin(\omega_i x + \phi_i)$, where $i = 0, \dots, L$, and L is the number of layers in the network. The hidden activations z_i and final outputs y_i of layer i are computed as:

$$z_i = \text{PE}_i(x) \otimes (W_i^h z_{i-1} + b_i^h) \text{ with } z_0 = \text{PE}_0(x); \quad (2)$$

$$y_i = y_{i-1} + (W_i^o z_i + b_i^o) \text{ with } y_0 = \mathbf{0}; \quad (3)$$

where \otimes denotes the Hadamard (element-wise) product, (W^h, b^h) parametrize the hidden linear layers, and (W^o, b^o) represent the output linear layers. This formulation leverages a useful property that repeated Hadamard product of sines are equivalent to a sum of sines with varying amplitude, frequency, and phase [12]. Therefore, with increasing PE frequencies from shallow to deep layers, we can produce output shapes with increasing amount of high-frequency details. This allows us to perform instance-specific optimization while preserving the common base shapes by sharing the shallow layers of part MLPs across instances and optimizing the deep layers to be instance-specific as shown in Fig. 5. Note that our MLP network is inspired from band-limited coordinate networks (BACON) [21]. However, we add the cumulative sum of all previous outputs to the current output of each layer since the high-frequencies details should be deformations from low-frequency based shapes. Moreover, unlike BACON, that learns the frequencies of sine functions in PE together with other parameters, we pre-define and fix the input frequencies to improve the optimization stability in our ill-posed problem and better control the separation of shared and instance components.

Zoom in to parts for higher details. Similar to LASSIE, we perform analysis-by-synthesis to supervise the overall shape reconstruction since we do not have access to any form of 3D supervision. We render the part surfaces via




Figure 6. **High-resolution rendering and optimization by zooming in on parts.** Based on the initial estimates of 2D part localization, we crop and upsample each part region to perform shape optimization at higher resolution.

a differentiable renderer [22] and compare them with the pseudo ground-truth masks obtained from DINO feature clustering. The silhouette loss \mathcal{L}_{sil} is written as: $\mathcal{L}_{sil} = \sum_j \|M^j - \hat{M}^j\|^2$, where M^j and \hat{M}^j are the rendered silhouette and pseudo ground-truth of instance j , respectively. To capture more shape details in the images, we propose to render and compare the high-resolution silhouettes by zooming in on individual parts during optimization. Concretely, we crop the 2D part masks estimated by Hi-LASSIE, upsample the cropped regions to higher resolution, and calculate the part silhouette loss \mathcal{L}_{part} on the zoomed-in part masks as:

$$\mathcal{L}_{part} = \sum_i \sum_j \left\| \Gamma_i^j(M^j) - \Gamma_i^j(\hat{M}^j) \right\|^2, \quad (4)$$

where Γ_i^j denotes the crop-and-upsample operation for part i on instance j . We show some example outputs before and after zoomed-in optimization in Fig. 6.

2D-3D semantic consistency via feature MLPs. To densely enforce the 2D-3D semantic consistency of part surfaces, we further introduce feature MLPs to improve the semantic loss proposed in LASSIE. For each instance j , the semantic consistency loss is defined as the Chamfer distance between the foreground pixels $\{p | \hat{M}^j(p) = 1\}$ and 3D surface points $\{x \in X\}$ in a high-dimensional space:

$$\mathcal{L}_{sem} = \sum_j \left(\sum_p \min_x \mathcal{D}(p, x) + \sum_x \min_p \mathcal{D}(p, x) \right). \quad (5)$$

The high-dimensional distance \mathcal{D} is defined as:

$$\mathcal{D}(p, x) = \underbrace{\|\pi^j(x) - p\|^2}_{\text{Geometric distance}} + \alpha \underbrace{\|Q(x) - K^j(p)\|^2}_{\text{Semantic distance}}, \quad (6)$$

where α is a scalar weighting for semantic distance, $Q(x)$ denotes the 3D surface features of point x , and $K(p)$ is the 2D image features at pixel p . In effect, \mathcal{L}_{sem} optimizes the 3D surface coordinates such that the aggregated 3D point features would project closer to the similar pixel features in the image ensemble. In LASSIE, the 3D surface features

are maintained and updated for only a sparse set of surface points, which limits the capabilities of \mathcal{L}_{sem} to supervise part localization. Instead, we represent the 3D surface feature of each part with an MLP (Fig 4), which is similar to the part shape MLPs shared by all instances. Consequently, we can obtain the semantic features $Q(x)$ given an arbitrary surface coordinate x . We update the part surface and feature MLPs alternatively in an EM-style optimization. That is, we update the feature MLPs by sampling 3D surface points and projecting them onto 2D images in the E-step. In the M-step, we use the updated features to optimize 3D surface MLPs via minimizing the semantic consistency loss.

Pose and shape regularizations. To constrain the output articulations and part shapes, we apply the following pose and shape regularizations. First, we impose a part rotation loss \mathcal{L}_{rot} to limit the angle offsets from resting pose as: $\mathcal{L}_{rot} = \sum_j \|R^j - \bar{R}\|^2$, where R^j is the part rotations of instance j and \bar{R} denotes the part rotations of shared resting pose. Moreover, we remove the side-way rotation constraint on animal legs in LASSIE since the leg parts are not specified in our self-discovered skeletons. Instead, we propose a more general regularization based on 3D symmetry prior. We define the symmetry loss \mathcal{L}_{sym} on the 3D joints to prevent overlapping parts or irregular poses as: $\mathcal{L}_{sym} = \sum_i \|J_i - \Psi(J_i^*)\|^2$ where J_i^* is the symmetric joint of J_i and Ψ is the reflection operation w.r.t. the symmetry plane. Finally, to encourage smooth part surfaces, we apply common mesh regularizations like Laplacian loss \mathcal{L}_{lap} and surface normal loss \mathcal{L}_{norm} . \mathcal{L}_{lap} encourages smooth 3D surfaces by pulling each vertex towards the center of its neighbors, and \mathcal{L}_{norm} enforces neighboring mesh faces to have similar normal vectors. Note that all the pose and shape regularizations are generic and applicable to a wide range of articulated shapes.

Optimization and texture sampling. The overall optimization objective is given by the weighted sum of aforementioned losses (\mathcal{L}_{sil} , \mathcal{L}_{part} , \mathcal{L}_{sem} , \mathcal{L}_{rot} , \mathcal{L}_{sym} , \mathcal{L}_{lap} , \mathcal{L}_{norm}). We optimize the camera, pose, and shape parameters in a multi-stage manner. Superficially, we first estimate the camera viewpoints and fix the rest. Then, we optimize the part transformations and shared part MLPs along with cameras until convergence. Finally, we freeze shared part MLPs (shallow layers) and fine-tune the instance-specific part deformations (deep layers) on each instance individually. Note that the surface feature MLPs are also updated during all optimization stages in an EM-style. The final textured outputs are generated by densely sampling the colors of visible surface points from individual images. The invisible (self-occluded) surfaces, on the other hand, are textured by their symmetric surfaces or nearest visible neighbors. More details can be found in the supplemental material.

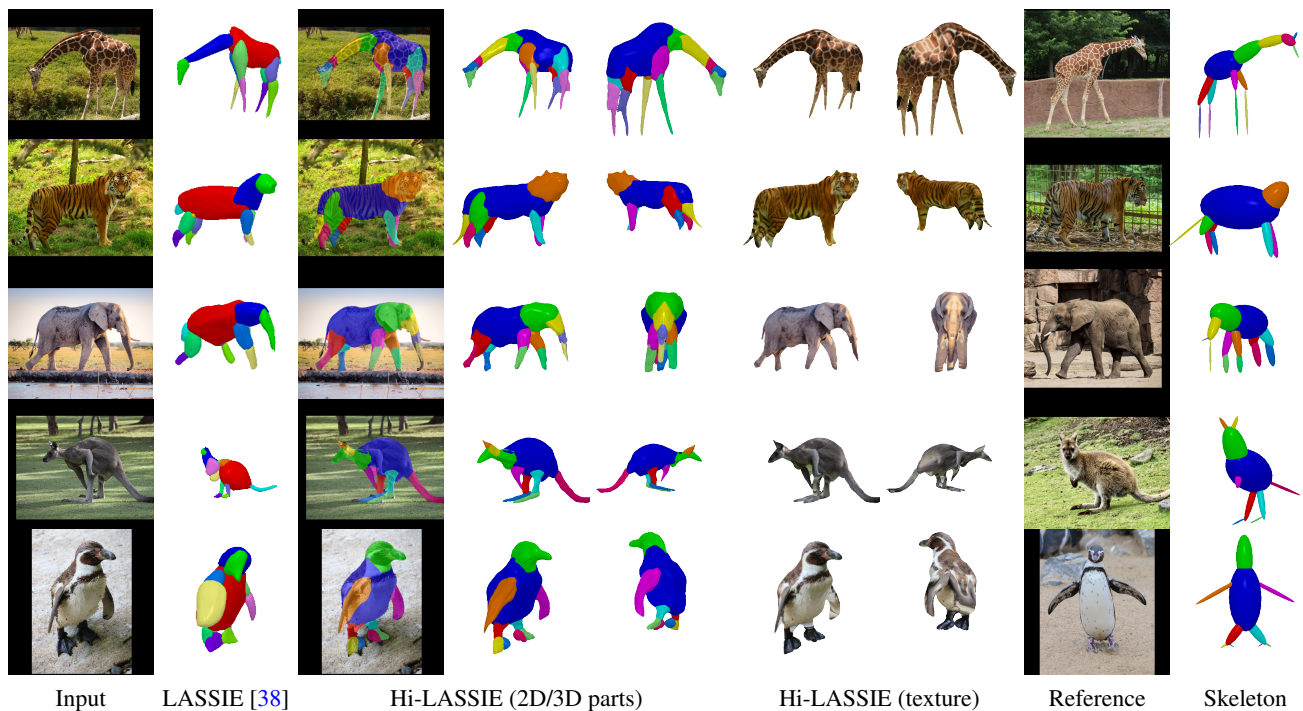


Figure 7. **Qualitative results on in-the-wild images.** We show example results of prior arts and Hi-LASSIE on the LASSIE [38] animal image ensembles as well as the reference image and 3D skeleton discovered by Hi-LASSIE. The results demonstrate the effectiveness of our 3D skeleton discovery and the high-fidelity shape/texture reconstruction across diverse animal classes.

4. Experiments

Datasets and baselines. We follow the same evaluation protocols as in LASSIE [38]. That is, we optimize and evaluate Hi-LASSIE on individual image ensembles in the Pascal-Part [8] and LASSIE [38] datasets. From the Pascal-Part dataset, we select the same set of images of horse, cow, and sheep as in LASSIE. These image are annotated with 2D part segmentation masks, which we use to automatically find 2D keypoints for evaluation. The LASSIE dataset contains sparse image ensembles (with CC-licensed web images) of some diverse animals (zebra, tiger, giraffe, elephant, kangaroo, and penguin) with 2D keypoints annotations for evaluation. Considering the novelty of this problem setting (sparse image optimization for articulated animal shapes), we mainly compare Hi-LASSIE with LASSIE as well as some learning-based mesh reconstruction methods. Most recent mesh reconstruction methods either cannot handle articulations [13, 17, 20, 30, 39] or leverage different inputs [19, 34, 36]. 3D Safari [43] and A-CSM [16], on the other hand, are more comparable to Hi-LASSIE since they explicitly model articulations for animal classes of our interest. Since both 3D Safari and A-CSM require large-scale image sets for training (not available in our setting), we use their released models of the closest animal classes to evaluate on our datasets. For instance, we use the zebra model of 3D Safari and the horse model of A-CSM to evaluate on similar quadrupeds.

Visual comparisons. Fig. 7 shows the qualitative results of Hi-LASSIE and prior methods on LASSIE dataset images. LASSIE results can fit the object silhouettes quite well but lack shape details. Our results demonstrate that Hi-LASSIE can effectively discover a good 3D skeleton that well explains other instance in the image ensemble. Moreover, the output articulated parts are detailed, high-fidelity, and faithful to input images. More qualitative results of skeleton and shape discovery are shown in the supplemental material.

Keypoint transfer. Without ground-truth 3D annotations in our datasets, we follow a common practice [16, 43] to evaluate 3D reconstruction by transferring 2D keypoints from source to target images. That is, we map a set of 2D keypoints on a source image onto the canonical 3D part surfaces, and project them to a target image via the estimated camera, pose, and shape. Since the keypoints are transferred from 2D-to-3D and from 3D-to-2D, a successful transfer indicates accurate 3D reconstruction on both the source and target images. In Table 1, we report the percentage of correct keypoints (PCK) under a tight threshold $0.05 \times \max(h, w)$, where h and w are image height and width, respectively. The results show that Hi-LASSIE achieves higher PCK on most animal image ensembles compared to the baselines while requiring minimal user inputs. We also show ablative results of Hi-LASSIE without the instance part MLPs (frequency decomposition), feature MLPs, or zoomed-in part silhouette loss \mathcal{L}_{part} . All the pro-

Table 1. **Keypoint transfer evaluations on the Pascal-Part [8] and LASSIE [38] image ensembles.** For all pairs of images in each animal class, we report the average percentage of correct keypoints (PCK@0.05).

Method	Horse	Cow	Sheep	Zebra	Tiger	Giraffe	Elephant	Kangaroo	Penguin
3D Safari [43]	57.1	50.3	50.5	62.1	50.3	32.5	29.9	20.7	28.9
A-CSM [16]	55.3	60.5	54.7	60.3	55.7	52.2	39.5	26.9	33.0
LASSIE [38]	58.0	62.4	55.5	63.3	62.4	60.5	40.3	31.5	40.6
Hi-LASSIE w/o inst. MLPs	56.8	57.7	53.6	59.7	63.0	59.9	40.8	31.5	41.2
Hi-LASSIE w/o feat. MLPs	57.5	62.2	56.3	64.1	62.0	60.8	42.7	30.3	41.5
Hi-LASSIE w/o \mathcal{L}_{part}	58.8	62.8	55.9	63.8	62.8	61.1	41.5	33.3	42.5
Hi-LASSIE	59.6	63.1	56.2	64.2	63.1	61.6	42.7	35.0	44.4

Table 2. **Quantitative evaluations on the Pascal-part images.** We report the overall 2D IOU, part mask IOU, and percentage of correct pixels (PCP) under dense part mask transfer between all source-target image pairs.

Method	Overall IOU			Part IOU			Part transfer (PCP)		
	Horse	Cow	Sheep	Horse	Cow	Sheep	Horse	Cow	Sheep
SCOPS [15]	62.9	67.7	63.2	23.0	19.1	26.8	-	-	-
DINO clustering [1]	81.3	85.1	83.9	26.3	21.8	30.8	-	-	-
3D Safari [43]	72.2	71.3	70.8	-	-	-	71.7	69.0	69.3
A-CSM [16]	72.5	73.4	71.9	-	-	-	73.8	71.1	72.5
LASSIE	81.9	87.1	85.5	38.2	35.1	43.7	78.5	77.0	74.3
Hi-LASSIE w/o inst. MLPs	80.4	80.0	79.5	30.2	29.3	33.8	76.4	74.9	71.1
Hi-LASSIE w/o feat. MLPs	82.5	87.6	85.9	34.9	32.4	39.7	74.7	72.4	74.9
Hi-LASSIE w/o \mathcal{L}_{part}	81.6	84.7	83.8	38.6	35.2	43.6	79.2	77.4	72.0
Hi-LASSIE	83.4	88.1	86.3	39.0	35.3	43.4	79.9	77.8	75.5

posed modules can effectively increase the accuracy of keypoint transfer demonstrating their use.

2D overall/part IOU. In addition to keypoint transfer accuracy, we compare Hi-LASSIE with the baselines using different segmentation metrics (Overall/Part IOU) in Table 2. For Hi-LASSIE and prior 2D co-part segmentation methods like SCOPS [15] and DINO clustering [1], we manually assign each discovered part to the best matched part in the Pascal-part annotations. Hi-LASSIE can produce accurate overall and part masks in 2D by learning high-fidelity 3D shapes. Compared to prior methods, Hi-LASSIE outputs match the Pascal-part segmentation better and achieves consistently higher overall IOU.

Part transfer. Finally, we show the results of part transfer accuracy in Table 2 using the percentage of correct pixels (PCP) metric proposed in LASSIE [38]. The PCP metric is designed similarly as PCK for keypoint transfer, but it uses 2D part segmentations to more densely evaluate 3D reconstruction. In short, we densely transfer the part segmentation from source to target images through mapping 2D pixels and 3D canonical surfaces. A correct transfer is done when a pixel is mapped to the same 2D part in both the source and target images. The PCP results further demonstrate the favorable performance of Hi-LASSIE against prior arts.

Applications. Hi-LASSIE not only produces high-fidelity 3D shapes but also enables various part-based applications due to explicit skeleton and part-based representation. For instance, we can easily transfer/interpolate the 3D skeleton transformations to repose or animate the output 3D shapes. Likewise, we can transfer the surface texture or part deformation from one animal species/instance to another. Due to their explicit nature, Hi-LASSIE 3D shapes can also be used by graphics artists for downstream applications in AR/VR/games. We show some example results of these applications in the supplemental material.

5. Conclusion

We propose Hi-LASSIE, a technique for 3D articulated shape reconstruction from sparse image ensemble without using any 2D/3D annotations or templates. Hi-LASSIE automatically discovers 3D skeleton based on a single reference image from the input ensemble. We further design several optimization strategies to reconstruct high-resolution and instance-varying details of 3D part shapes across the given ensemble. Our results on Pascal-Part and LASSIE image ensembles demonstrate the favorable reconstructions of Hi-LASSIE against prior arts despite using minimal user annotations. In future, we hope to apply Hi-LASSIE on more general articulated objects in-the-wild.

References

- [1] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep ViT features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2021. 2, 3, 8
- [2] Xiang Bai and Longin Jan Latecki. Path similarity skeleton graph matching. *PAMI*, 30(7):1282–1292, 2008. 3, 4
- [3] Serge Belongie, Jitendra Malik, and Jan Puzicha. Shape matching and object recognition using shape contexts. *PAMI*, 24(4):509–522, 2002. 3
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. NeRD: Neural reflectance decomposition from image collections. In *CVPR*, pages 12684–12694, 2021. 3
- [5] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan T Barron, Hendrik Lensch, and Varun Jampani. SAMURAI: Shape and material from unconstrained real-world arbitrary image collections. *NeurIPS*, 2022. 3
- [6] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-PIL: Neural pre-integrated lighting for reflectance decomposition. *NeurIPS*, 34, 2021. 3
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 2, 4
- [8] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *CVPR*, pages 1971–1978, 2014. 2, 7, 8
- [9] Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Unsupervised part discovery from contrastive reconstruction. *NeurIPS*, 34:28104–28118, 2021. 3
- [10] Edo Collins, Radhakrishna Achanta, and Sabine Susstrunk. Deep feature factorization for concept discovery. In *ECCV*, pages 336–352, 2018. 3
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [12] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *ICLR*, 2020. 5
- [13] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *ECCV*, pages 88–104, 2020. 3, 7
- [14] Xingzhe He, Bastian Wandt, and Helge Rhodin. Autolink: Self-supervised learning of human skeletons and object outlines by linking keypoints. *arXiv preprint arXiv:2205.10636*, 2022. 3
- [15] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. SCOPS: Self-supervised co-part segmentation. In *CVPR*, pages 869–878, 2019. 3, 8
- [16] Nilesh Kulkarni, Abhinav Gupta, David F Fouhey, and Shubham Tulsiani. Articulation-aware canonical surface mapping. In *CVPR*, pages 452–461, 2020. 2, 3, 7, 8
- [17] Nilesh Kulkarni, Abhinav Gupta, and Shubham Tulsiani. Canonical surface mapping via geometric cycle consistency. In *ICCV*, pages 2202–2211, 2019. 2, 3, 7
- [18] Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, Nicu Sebe, et al. Motion-supervised co-part segmentation. *arXiv preprint arXiv:2004.03234*, 2020. 3
- [19] Xueting Li, Sifei Liu, Shalini De Mello, Kihwan Kim, Xiaolong Wang, Ming-Hsuan Yang, and Jan Kautz. Online adaptation for consistent mesh reconstruction in the wild. *NeurIPS*, 33:15009–15019, 2020. 7
- [20] 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. In *ECCV*, pages 677–693, 2020. 3, 7
- [21] David B Lindell, Dave Van Veen, Jeong Joon Park, and Gordon Wetzstein. Bacon: Band-limited coordinate networks for multiscale scene representation. In *CVPR*, pages 16252–16262, 2022. 5
- [22] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3D reasoning. In *CVPR*, pages 7708–7717, 2019. 6
- [23] Tiange Luo, Kaichun Mo, Zhiao Huang, Jiarui Xu, Siyu Hu, Liwei Wang, and Hao Su. Learning to group: A bottom-up framework for 3D part discovery in unseen categories. *arXiv preprint arXiv:2002.06478*, 2020. 3
- [24] Priyanka Mandikal, Navaneet KL, and R Venkatesh Babu. 3D-PSRNet: Part segmented 3D point cloud reconstruction from a single image. In *ECCV*, pages 0–0, 2018. 3
- [25] 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. In *ECCV*, pages 405–421, 2020. 3
- [26] Despoina Paschalidou, Luc Van Gool, and Andreas Geiger. Learning unsupervised hierarchical part decomposition of 3D objects from a single RGB image. In *CVPR*, pages 1060–1070, 2020. 3
- [27] Despoina Paschalidou, Angelos Katharopoulos, Andreas Geiger, and Sanja Fidler. Neural parts: Learning expressive 3D shape abstractions with invertible neural networks. In *CVPR*, pages 3204–3215, 2021. 3
- [28] Nadine Rueegg, Silvia Zuffi, Konrad Schindler, and Michael J Black. BARC: Learning to regress 3D dog shape from images by exploiting breed information. *arXiv preprint arXiv:2203.15536*, 2022. 3
- [29] Wei Shen, Yuan Jiang, Wenjing Gao, Dan Zeng, and Xinggang Wang. Shape recognition by bag of skeleton-associated contour parts. *Pattern Recognition Letters*, 83:321–329, 2016. 3
- [30] Shubham Tulsiani, Nilesh Kulkarni, and Abhinav Gupta. Implicit mesh reconstruction from unannotated image collections. *arXiv preprint arXiv:2007.08504*, 2020. 3, 7
- [31] Shubham Tulsiani, Hao Su, Leonidas J Guibas, Alexei A Efros, and Jitendra Malik. Learning shape abstractions by assembling volumetric primitives. In *CVPR*, pages 2635–2643, 2017. 3

- [32] Narek Tumanyan, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Splicing ViT features for semantic appearance transfer. *arXiv preprint arXiv:2201.00424*, 2022. [2](#)
- [33] Jun Xie, Pheng-Ann Heng, and Mubarak Shah. Shape matching and modeling using skeletal context. *Pattern Recognition*, 41(5):1756–1767, 2008. [3](#), [4](#)
- [34] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Huiwen Chang, Deva Ramanan, William T Freeman, and Ce Liu. LASR: Learning articulated shape reconstruction from a monocular video. In *CVPR*, pages 15980–15989, 2021. [3](#), [7](#)
- [35] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Ce Liu, and Deva Ramanan. ViSER: Video-specific surface embeddings for articulated 3d shape reconstruction. *NeurIPS*, 34, 2021. [3](#)
- [36] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. BANMo: Building animatable 3D neural models from many casual videos. *arXiv preprint arXiv:2112.12761*, 2021. [3](#), [7](#)
- [37] Chun-Han Yao, Wei-Chih Hung, Varun Jampani, and Ming-Hsuan Yang. Discovering 3d parts from image collections. In *ICCV*, pages 12981–12990, 2021. [3](#)
- [38] Chun-Han Yao, Wei-Chih Hung, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, and Varun Jampani. Lassie: Learning articulated shapes from sparse image ensemble via 3d part discovery. *arXiv preprint arXiv:2207.03434*, 2022. [2](#), [3](#), [7](#), [8](#)
- [39] Yufei Ye, Shubham Tulsiani, and Abhinav Gupta. Shelf-supervised mesh prediction in the wild. In *CVPR*, pages 8843–8852, 2021. [3](#), [7](#)
- [40] Jason Zhang, Gengshan Yang, Shubham Tulsiani, and Deva Ramanan. NeRS: Neural reflectance surfaces for sparse-view 3D reconstruction in the wild. *NeurIPS*, 34, 2021. [3](#)
- [41] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. PhySG: Inverse rendering with spherical Gaussians for physics-based material editing and relighting. In *CVPR*, pages 5453–5462, 2021. [3](#)
- [42] Tongjie Y Zhang and Ching Y. Suen. A fast parallel algorithm for thinning digital patterns. *CACM*, 27(3):236–239, 1984. [3](#)
- [43] Silvia Zuffi, Angjoo Kanazawa, Tanya Berger-Wolf, and Michael J Black. Three-D Safari: Learning to estimate zebra pose, shape, and texture from images ”in the wild”. In *ICCV*, pages 5359–5368, 2019. [2](#), [3](#), [7](#), [8](#)
- [44] Silvia Zuffi, Angjoo Kanazawa, and Michael J Black. Lions and tigers and bears: Capturing non-rigid, 3D, articulated shape from images. In *CVPR*, pages 3955–3963, 2018. [2](#), [3](#)
- [45] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *CVPR*, pages 6365–6373, 2017. [3](#)