This ACCV 2024 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2024 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

Dessie: Disentanglement for Articulated 3D Horse Shape and Pose Estimation from Images

Ci Li¹^(b), Yi Yang^{1,2}^(c), Zehang Weng¹^(b), Elin Hernlund³^(c), Silvia Zuffi⁴^(c), and Hedvig Kjellström^{1,3}^(c)

KTH, Sweden {cil, yiya, zehang, hedvig}@kth.se
 ² Scania, Sweden carol-yi.yang@scania.com
 ³ SLU, Sweden Elin.Hernlund@slu.se
 ⁴ IMATI-CNR, Italy silvia@mi.imati.cnr.it



Fig. 1: We estimate 3D shape and pose of horses from monocular images. The figure shows pictures of Dessie, a famous racehorse, together with our 3D reconstruction.

Abstract. In recent years, 3D parametric animal models have been developed to aid in estimating 3D shape and pose from images and video. While progress has been made for humans, it's more challenging for animals due to limited annotated data. To address this, we introduce the first method using synthetic data generation and disentanglement to learn to regress 3D shape and pose. Focusing on horses, we use textbased texture generation and a synthetic data pipeline to create varied shapes, poses, and appearances, learning disentangled spaces. Our method, Dessie, surpasses existing 3D horse reconstruction methods and generalizes to other large animals like zebras, cows, and deer. See the project website at: https://celiali.github.io/Dessie/.

Keywords: Animal 3D reconstruction, disentanglement

1 Introduction

Capturing 3D articulated motion of humans and animals is essential for understanding motion dynamics and enhancing various applications such as virtual reality [49]. In addition, motion can be often associated with diseases and

pain, making motion capture an effective diagnosis tool for animal subjects that do not communicate discomfort verbally. Recent advances in computer vision have promoted 3D markerless motion capture [69], allowing motion inference from visual data without invasive equipment. Substantial efforts in these markerless approaches have focused on using parameterized models like SMPL for humans [47] and SMAL for animals [84], which define local (shape and pose) and global (rotation and translation) properties. Their precise capture makes them important analysis tools across various fields, particularly in biomedical downstream tasks [28, 29, 50].

In this work, we focus on horses. Horses are particularly relevant for their economic value [12, 23]. Maintaining horses requires substantial resources [61]. Moreover, despite their large bodies, horses are delicate, owning an anatomical fragility that often leads to severe, irrecoverable injuries, in particular for the most precious racing subjects. Consequently, horses have been extensively studied from both behavioral and biomechanics perspectives. In computer vision, recent efforts have developed 3D articulated parameterized shape models specific for horses [38, 85], confirming our view that horses, for their particular wide presence in several human activities, require ad-hoc, specific models. What the horses have in common with other animals is the scarcity of 3D annotated datasets that would facilitate learning to regress 3D poses from images with endto-end regressors. In addition, horses might share similarities in how they move with other species of the same Equine family. Motivated by these observations, we explore disentanglement in learning image-to-3D model parameter regressors.

Disentanglement is a powerful machine learning technique, offering advancements in different fields like image generation [21] and style transfer [34]. By separating underlying factors of variation in the data, it facilitates more interpretable models and enhances generalization and robustness. In image recognition, disentangled representations have achieved state-of-the-art (SOTA) results [18,59], showcasing their ability to isolate and manipulate salient features in complex data. Motivated by these successes, we investigate the potential of disentanglement to address challenges in monocular 3D horse mesh reconstruction.

Recent advancements in 3D markerless motion capture domain have shown that regression-based approaches, utilizing neural networks to directly map image pixels to model parameters, achieve SOTA results for human [17, 25, 26, 31] and animal [3,83] modeling. These methods employ encoders [13,20], for feature extraction, with decoders predicting model parameters from the extracted features. However, these methods embed all image attributes—ranging from subject shape and pose to the background—into a single feature space, entangling diverse properties, and complicating the decoder's task of isolating relevant information. In this work, we employ a specialized encoder to separate features relevant to the subject's local (pose, shape) and global information (rotation, translation, camera), coupled with distinct decoders to regress separate model parameters. To the best of our knowledge, this is the first work in which disentangled feature spaces are used for animal 3D reconstruction. To enable disentanglement learning and regression-based 3D reconstruction of animals, a large-scale dataset is essential. Collecting real-world data is challenging, primarily due to the scarcity of real-world animal data, and the high costs associated with data collection, unlike the relative abundance of available human data [17]. To address this challenge, we develop a synthetic pipeline, DessiePIPE, capable of generating realistic animal images on the fly, to support regression-based 3D animal reconstruction while encouraging disentanglement learning. DessiePIPE utilizes a newly published horse motion capture dataset [39] and the hSMAL model [38], rendering horse images, in various poses (standing, walking, trotting, cantering, eating, bending neck, sitting, rearing), shapes, and species-related textures, with randomized orientations and backgrounds. Additionally, we introduce a novel technique, to create image tuples with controllable variations, allowing the study of disentangled features by producing pairs of images that vary one feature while maintaining others (*e.g.* different appearances with the same pose and global rotation).

Leveraging DessiePIPE, we present DinoHMR, a variant of the popular HMR framework [25], and introduce Dessie⁵, a multi-stream framework combining disentanglement learning. Both methods are based on the vision foundation model DINO [8]. Unlike HMR, Dessie disentangles image features into different spaces with independent processing paths, such that information that pertains to disentangled features can propagate through the layers without spurious influences.

Through extensive experiments, we validate our methods, demonstrating their effectiveness by training exclusively on synthetic data from DessiePIPE. Our methods generalize well across unseen real-world datasets and deliver both qualitative and quantitative improvements when fine-tuned on real image datasets of varying sizes, from limited to large. Dessie outperforms existing 3D animal reconstruction methods even with limited real data, offering a solution to data scarcity, and generalizes well across out-of-distribution image domains and to other horse-like species.

To summarize, our contributions are:

- We propose DessiePIPE, a method to generate diverse and high-quality synthetic images for 3D horse markerless reconstruction learning.
- We present Dessie, a framework for 3D pose and shape estimation of horses from monocular images based on the foundation model DINO; on disentangled latent spaces; and on disentangled learning. For comparison, we also propose DinoHMR (Dessie without disentanglement).
- We show that disentangled learning results in more efficient training, and achieves SOTA performance on various benchmarks.

2 Related Work

Model-based methods for 3D reconstruction of humans and animals. The markerless motion capture of articulated deformable subjects, like humans and many

⁵ Desert Orchid (1979–2006), known as Dessie, was one of the most beloved and iconic racehorses in British history. Source: Wikipedia.

animals, requires prior models of body shape and pose when applied in monocular settings. In recent years, the development of 3D parametric shape models of the human body has seen explosive development in computer vision and graphics, with the SMPL model [47] being the de-facto standard. SMPL and its successors [47, 55, 56, 58] are data-driven models learned from thousands of 3D body scans of real people. They define statistical parametric linear shape models, for generating bodies of different shapes, sizes, and genders. A huge number of methods [5, 25, 26, 30–33, 36, 40, 41, 43, 57, 80] focus on recovering 3D shape and pose from monocular images or videos using SMPL models.

Given the challenges of learning animal models from 3D captured data, initial research on animal modeling focused on creating 3D models from images. A pioneering effort by Cashman and Fitzgibbon demonstrates the reconstruction of dolphins and bears from images [9]. Later, Zuffi et al. introduce SMAL [84], the animal equivalent of SMPL, learned from 3D scans of a limited set of toys representing different quadrupeds. Based on the SMAL model, Zuffi et al. [83] introduce an end-to-end regression network to reconstruct zebras, marking an important step towards more detailed and accurate animal 3D model-based monocular reconstruction. The 3D capture of dogs has seen innovation, with Kearney et al. [27] utilizing RGBD images for reconstruction, while Biggs et al. [3, 4] achieve pose and shape estimation from solely image data. Rüegg et al. [64, 65]enhance accuracy by incorporating breed-specific information and ground contact data. Bird modeling has also progressed, employing images and a 3D synthetic model [2,72] to learn shape variation of birds. Li et al. [38] introduce the hSMAL model, a horse-specific SMAL model that facilitates the recovery of 3D poses from images and aids in lameness detection in horses [45], showing the potential of these technologies for biological and medical applications. VAREN [85] builds upon hSMAL, leveraging real horse scan data.

Many existing regression methods [3, 25, 31, 33, 41, 43, 83] encode images as a single feature vector to directly predict camera, model shape and pose parameters. Notably, some methods take a different approach. PARE [30] employs part-guided attention from one module alongside 3D features from another to better understand occluded body parts. HybrIK [40] leverages separate modules to generate 3D heatmaps and predict shape and twist angles for inverse kinematics. Pavlakos et al. [57] estimate silhouette and heatmap from images to infer model shape and pose, showing that disentangling shape and pose enhances 3D prediction stability and accuracy. Recent works on dogs, Barc [65] and Bite [64], employ separate pose and shape modules. Unlike these approaches, our work proposes to learn three separate feature spaces for shape, pose, and global information (rotation, translation and camera) prediction in a disentangled fashion.

Feature Extraction Networks. The backbone networks for extracting image features are diverse, with ResNet [20], Stack Hourglass [53], and HRNet [68] being commonly used. Recently, DINO [8], a self-supervised self-distillation largescale foundation model, has proven effective in extracting image representation [71] and performs well on downstream tasks [1, 70]. Pretrained DINO features have been used in non-template based 3D animal reconstruction tasks, like, in Lassie [77], Hi-Lassie [78], and Artic3D [79] to optimize 3D animal shape from a small scale of images (less than 30 images); in Magicpony [74], Farm3D [24], 3D Fauna [44] to learn 3D priors from a large-scale dataset (over 10k images). Similar to our work, Digidogs [66] fine-tunes DINOv2 [54] for estimating 3D dog skeletons of real-world images using simulation data by extracting token features from DINOv2. Ours more focus on the extracted key features from DINO.

Disentanglement. Disentanglement aims to figure out the underlying latent representations that control different data properties. It has be applied to both 2D images [11, 21, 48] and 3D meshes [82]. Lorenz et al. [48] disentangle shape and appearance from 2D images by learning parts consistency over category instances, leveraging the equivariance and invariance constraints on the image representation to extract part-based information. Zhou et al. [82] disentangle shape and pose features from registered meshes by using self-consistency and cross-consistency constraints in an unsupervised setting. Hu et al. [22] extract disentangled 3D attributes only from 2D images to manipulate the 3D output, based on PiFu, which decodes image features from the backbone network into 3D implicit representations of humans. The image features encode three attributes. which are disentangled by the designed frameworks. Zhao et al. [81] disentangle pose and view factors from 2D human poses by maximizing mutual information of the same pose from different views in a contrastive learning manner. While these works focus on 2D pose estimation, 3D mesh generation or 3D mesh manipulation tasks, none address monocular model-based 3D reconstruction.

To encourage disentanglement, contrastive learning is an effective strategy. Traditional contrastive methods, such as SimCLR [10] and MoCo [19], typically involve pre-training an image encoder on large datasets in an unsupervised manner. These methods gain their effectiveness by generating positive and negative image pairs, which encourages the model to closely align the features of the positive pairs. Our method simplifies the process by leveraging the labels from the generated image pairs and ensuring feature similarity for each matching pair.

3 Method

3.1 The hSMAL model

Analogous to the SMPL model, the hSMAL model [38], a horse-specific version of the SMAL model, is a parameterized 3D model of horses. It factorizes surface variations into shape, and pose. Given 37 horse toy scans as training data, the average shape over the training set is calculated as V_{mean} , and PCA is performed on the variations of each scan to the mean shape. The shape, noted as $\beta \in \mathbb{R}^9$, is the set of PCA coefficients representing vertex-based deformations from the mean shape. The pose, noted as $\theta = (\theta_G, \theta_J)$, represents the global rotation $\theta_G \in \mathbb{R}^3$ and the relative rotation of each joint with respect to its parents in the kinematic tree $\theta_J \in \mathbb{R}^{35\times 3}$ in axis angle representation. The hSMAL model defines a mapping function, $\Theta : (\beta, \theta, \xi) \mapsto \mathbf{v}$, which outputs a 3D mesh based on linear blend skinning. ξ is the global translation.



Fig. 2: DessiePIPE: Synthetic generation pipeline. Top: the training image generation process; bottom: generated data samples with controlled variations.

3.2 DessiePIPE

As noted, to advance disentanglement learning and regression-based 3D reconstruction techniques, we require a large-scale animal dataset. However, compiling this dataset through real-world data collection presents significant challenges. Particularly, it's highly impractical to obtain e.g. a pair of animal images in identical poses with different appearances, which is crucial for promoting feature disentanglement by having image pairs that share similar features with variations in another feature. To overcome this challenge, we develop DessiePIPE. DessiePIPE allows for the creation of horse images with diverse combinations of camera viewpoints, horse appearances, and poses, utilizing Pytorch3D [60] and the hSMAL model as shown in Fig 2. It differs from previous offline synthetic data generation, using limited textures and poses learned from 2D images [75], or fixed shape [37,52]. A concurrent method is FoundationPose [73], which generates synthetic data using language and diffusion models for RGBD 3D rendering.

To generate realistic horse appearances, we create a finite set containing 80 highly realistic UV texture maps following TEXTure [62], which generates texture maps for a given 3D shape with a text prompt by leveraging pretrained depth-to-image stable diffusion model [63]. Specifically, the texture set is constructed with text prompts based on eight distinct horse species, including Bay Thoroughbred, Palomino Quarter Horse, Chestnut Morgan, Buckskin Tennessee Walker, White Arabian, Black Friesian, Dapple Gray Andalusian, Pinto Paint Horse. We refer the readers to *Supplementary Materials* for more details. For the horse's shape, we obtain an infinite set by randomly sampling within the hSMAL



Fig. 3: Network architectures of Dessie. Dessie extracts latent features $(\gamma_A, \gamma_P, \gamma_G)$ with DINO and predicts the hSMAL parameters with the corresponding decoders. The model is trained with keypoint loss, silhouette loss and a contrastive loss to encourage the disentanglement.

shape space with a Gaussian distribution. Regarding the pose, we create a finite set composed of realistic horse poses extracted from PFERD [39], including daily motion poses (standing, walking, trotting, eating, bending neck) and advanced poses (sitting, rearing). Note that these poses also encapsulate the global rotations. Random image backgrounds come from the COCO dataset [46]. With the described components, DessiePIPE randomly selects an item from each of the appearance, pose, and background sets to compose a horse image. This process is further exploited by DessiePIPE to generate pairs of images (I1, I2), where I1 and I2 are images that differ in one aspect—either model shape, pose, or global rotation. Fig 2 provides examples of how individual changes impact the rendered outcomes. Images are generated using Pytorch3D, each with a resolution of (256, 256). Ground-truth annotations for each image include: global rotation $\theta_{G_{qt}}$, pose $\theta_{J_{qt}}$, shape β_{gt} , global translation ξ_{gt} , silhouette S_{gt} , and landmark locations K_{at} . We define 17 landmarks on the 3D model template, corresponding to the 2D animal keypoint labels in ViTPose+ [76], and apply the Pytorch3D rasterizer to filter out and retain landmarks visible in the camera views.

3.3 Network Architecture

Inspired by HMR [25] and DINO-related feature studies [1, 8, 71, 74], we propose two extractor-decoder frameworks that incorporate the DINO model: 1) DinoHMR, an HMR-like model, and 2) Dessie, a multi-stream structure with a modified contrastive loss L_{DFL} , for disentangled shape and pose estimation (Fig 3). Unlike HMR, which encodes image features with ResNet, both DinoHMR and Dessie exploit the capabilities of DINO's key features from the last layer.

In DinoHMR, we employ an *extractor* with a Conv2D network to extract a single feature representation from the keys, and a *decoder* for predicting shape β , pose θ_J , global rotation θ_G and translation ξ parameters. In contrast, the Dessie *extractor* $\mathcal{F}(.)$ employs three Conv2D networks on the keys to isolate features $\gamma_A \in \Gamma_A, \gamma_P \in \Gamma_P$, and $\gamma_G \in \Gamma_G$, corresponding to the feature spaces for subject appearance, subject pose, and global information, respectively, with $\Gamma_A, \Gamma_P, \Gamma_G \subseteq \mathbb{R}^{640}$. The Dessie *decoder* is composed of three distinct submodules $\mathcal{M}_A, \mathcal{M}_P$, and \mathcal{M}_G , each processing a specific set of features extracted by the

Dessie extractor. As shown in Fig. 3, the submodule \mathcal{M}_A is responsible for inferring the shape parameters β from the appearance features γ_A . The submodule \mathcal{M}_P deduces the pose parameters, θ_J leveraging the pose features γ_P . Lastly, \mathcal{M}_G estimates both the global rotation θ_G and the global translation ξ , utilizing the global information features γ_G . The entire framework is described by:

Dessie Extractor:
$$(\gamma_A, \gamma_P, \gamma_C) = \mathcal{F}(x),$$

Dessie Decoder: $\beta = \mathcal{M}_A(\gamma_A),$
 $\theta_J = \mathcal{M}_P(\gamma_P),$
 $(\theta_G, \xi) = \mathcal{M}_G(\gamma_G)$
(1)

We employ a weak-perspective camera model $\Omega(s, t_x, t_y)$, with learnable scale $s \in R$ and translation $(t_x, t_y) \in R^2$, and pre-defined focal length f = 5000. Given the relativity between the camera and the 3D model, the global translation of the model is defined as $\xi = [t_x, t_y, \frac{2 \cdot f}{r \cdot s}]$, where r = 256 is the image resolution. We use a differentiable render, Pytorch3d, to render the silhouette \tilde{S} and 2D keypoints \tilde{K} by projecting the model back to the image.

3.4 Loss

Dessie is trained in an end-to-end manner, with loss defined as:

$$L = L_{KP} + L_{SIL} + L_{HSMAL} + L_{DFL}.$$
(2)

 L_{KP} is a 2D keypoint loss defined as :

$$L_{KP} = \omega_K \frac{\sum_{j=1}^J \lambda_j^2 \rho \left\| \tilde{K}^j - K_{gt}^j \right\|_2}{\sum_{j=1}^J \lambda_j^2}, \ \tilde{K}^j = \Pi(\kappa^j, \Omega), \tag{3}$$

where κ are the body keypoints on the model, projected by the camera Ω . K_{gt} are the ground-truth 2D keypoints with visibility flag λ , and ρ is the Geman-McClure robustifier [15]. Π is the differentiable rendering function. L_{SIL} is a silhouette loss defined as:

$$L_{SIL} = \omega_S smooth L1(\tilde{S}, S_{gt}), \ \tilde{S} = \Pi(\mathbf{v}, \Omega), \tag{4}$$

where S_{gt} is the groundtruth silhouette, \tilde{S} is the projected model silhouette, and smoothL1(.) is the SmoothL1 loss [16]. L_{HSMAL} is the weighted sum of the shape and the pose priors of the hSMAL model, defined in [38], with corresponding weights ω_{Prior}^{β} and ω_{Prior}^{θ} . L_{DFL} is a variant of contrastive loss designed to encourage the disentanglement between different features and minimize intrapair differences with known input pair variations, defined as:

$$L_{DFL} = \omega_D \begin{cases} MSE(\gamma_{P_1}, \gamma_{P_2}) + MSE(\theta_{G_1}, \theta_{G_2}), & \text{if appearance difference.} \\ MSE(\gamma_{A_1}, \gamma_{A_2}) + MSE(\theta_{G_1}, \theta_{G_2}), & \text{if pose difference.} \\ MSE(\gamma_{A_1}, \gamma_{A_2}) + MSE(\gamma_{P_1}, \gamma_{P_2}), & \text{if global rotation difference.} \end{cases}$$
(5)

Given the ambiguity between the camera and model, the loss on global translation is not applied. It is important to note that ground-truth shape and pose labels from DessiePIPE are not necessary for the learning process. For DinoHMR, the contrastive loss term L_{DFL} is omitted.

4 Experiments

4.1 Datasets

Synthetic Dataset. Through DessiePIPE, we construct synthetic validation dataset, using 10% of the texture maps and around 12% of the poses, and the validation set from the COCO dataset for the background. For each training epoch, DessiePIPE creates 6.4k and 640 images for training and validation.

Real-world Dataset. Our evaluation utilizes several public datasets: 1) The horse category from PASCAL [14] and AnimalPose [7]; 2) The horse category in MagicPony [74] dataset, which compiles images from Weizmann Horse Database [6], PASCAL [14], and Horse-10 Dataset [51]. Each real-world dataset includes annotated segmentation masks. Given the keypoints labels are not available in the MagicPony dataset, we generate keypoint labels with ViTPose+ [76].

4.2 Network Implementation Details

DinoHMR and Dessie both use DINO-ViTs8 as the backbone for the extractor \mathcal{F} , with the last three layers unfrozen for adaptation to the 3D animal reconstruction task. In the decoder \mathcal{D} , each module comprises two fully connected layers with dropout applied in between, followed by a fully connected layer to predict the residual of each parameter. The associated features and the corresponding parameters are concatenated as input to each module, which outputs residual to update the parameters for three iterations. Training spans up to 800 epochs, with a learning rate of 5×10^{-5} . Model selection is based on achieving the lowest loss on the validation dataset. For the weights in the loss function, from Eq. 2 to Eq. 5, we set $\omega_K = 0.001$, $\omega_S = 0.0001$, $\omega_{\text{Prior}}^{\beta} = 0.01$, $\omega_{\text{Prior}}^{\theta} = 0.01$, and $\omega_D = 0.02$. We apply color jittering to the images. The training process is conducted on two Nvidia A100 GPUs for approximately two days.

4.3 Experiments

We conduct two main experiments to demonstrate the effectiveness of Dessie. First, we train both DinoHMR and Dessie solely based on synthetic data generated by DessiePIPE with the same sampling strategy, and validate their performance on unseen real-world datasets. Second, we fine-tune DinoHMR and Dessie using either a limited amount of real data from Staths [67] or a comparatively larger dataset from MagicPony [74], to showcase the model's capacity for generalization across different scales of data.

DinoHMR vs. Dessie with DessiePIPE. Given that DessiePIPE is fully controllable and observable, we can incorporate ground-truth shape and pose labels from DessiePIPE for supervised learning. We introduce a ground-truth loss term $L_{\rm gt}$, for both shape β and pose θ , using the mean squared error (MSE) between the predicted values and the ground truth labels from DessiePIPE.

Network	L_{gt}	$ L_{DFL} $	Magicpony	Dataset	AnimalPose		Pascal	
			$\mathrm{PCK}_{@0.1}(\uparrow)$	$IoU(\uparrow)$	$\mathrm{PCK}_{@0.1}(\uparrow)$	$\mathrm{IoU}(\uparrow)$	$PCK_{@0.1}(\uparrow)$	$\mathrm{IoU}(\uparrow)$
DingUMP	-	-	67.70	57.39	68.42	56.87	47.65	47.63
DIHOITIWIK	\checkmark	-	61.02	54.68	68.70	54.87	45.23	45.85
	-	-	59.94	54.61	67.13	55.52	43.11	44.58
Dessie	\checkmark	-	69.42	58.02	73.63	58.24	47.73	46.93
Dessie	\checkmark	\checkmark	65.66	56.39	70.26	56.15	45.15	47.69
	-	\checkmark	70.97	59.99	75.43	59.69	52.27	52.51

Table 1: DinoHMR vs. Dessie with DessiePIPE on unseen image datasets.

Note: \checkmark indicates inclusion of the loss component in training; - indicates exclusion.

We examine the performance of DinoHMR and Dessie when utilizing the contrastive loss L_{DFL} , in comparison to their supervised-learning variants employing $L_{\rm gt}$, as presented in Table 1. Notably, DinoHMR and Dessie are trained without any real images. Our evaluation metrics include the Intersection Over Union (IoU) and the Percentage of Correct Keypoint (PCK) with a threshold of 0.10, applied to manually defined keypoints on the model across all unseen real-world datasets. For MagicPony dataset, we report PCK among the 17 predicted keypoints defined by ViTPose+. For Pascal and AnimalPose datasets, we follow the evaluation in [67], employing 16 keypoints as specified for Pascal. Dessie (w/o L_{gt} , w/ L_{DFL}) consistently outperforms other methods, supporting our hypothesis that a model promoting feature disentanglement achieves better generalization in zero-shot real image reconstructions.

DINO Key Feature. Given that we train Dessie with DINO backbone unfrozen, we are interested in one question: How does the 3D reconstruction learning affect the DINO feature space? To answer this question, we utilize PCA to visualize the key features extracted from the DINO model's last layer, and examine the self-similarity of these features, as in previous studies [71]. Brighter areas represent higher component values, indicating regions of intense focus for reconstruction. This analysis uses two synthetic test images and two real images from the AnimalPose dataset, as shown in Fig 4. The comparison reveals that while the vanilla DINO features capture the spatial layout and structure of both the subjects and environments, the features in Dessie shift to focus more toward the targeted reconstructed subjects, effectively minimizing background distractions.

Dessie vs. SOTA with limited real shots. We assess the performance of DinoHMR and Dessie against A-CSM [35] and Staths [67], with a constrained training dataset of 150 labeled images [67]. Notably, Staths is an augmented A-CSM framework with additional web-sourced images for training. During the training, we construct batches to maintain an equal proportion of synthetic and real images, selecting model based on the performance of the synthetic validation set. Real images are excluded from computing L_{DFL} and the silhouette loss L_{SIL} to maintain consistency with [67] during training. The effectiveness is quantified

Dessie 11



Fig. 4: Visualization of the leading PCA components of key features for two synthetic and two real images from top to down. For each image, we visualize the original DINO (row 1) and the Dessie key features (row 2).

on AnimalPose and Pascal datasets as shown in Table 2, using the Area Under the Curve (AUC) metric to capture PCK performance across varying thresholds from 0.06 to 0.1. The results show that both DinoHMR and Dessie, even when fine-tuned with only the 150 images, outperform Staths, highlighting their effectiveness in utilizing limited real-world data resources.

AUC (\uparrow)	A-CSM [35]†	Staths $[67]^{\dagger}$	DinoHMR	Dessie	${\rm DinoHMR}\circledast$	$\mathrm{Dessie}\circledast$
AnimalPose	51.0	75.1	55.6	62.2	81.8	82.9
Pascal	37.4	55.9	36.4	40.7	60.5	61.2
Note: (a): Model finatured on detection [67]						

Table 2: KeyPoint Evaluation. †: number taken from [67].

Note: \circledast : Model finetuned on dataset in [67].

Dessie vs. SOTA with larger real dataset. We extend our evaluation of DinoHMR and Dessie by training them with a larger dataset from MagicPony [74], denoted as DinoHMR \star and Dessie \star , and conducting assessments using the PASCAL horse subset [14]. For real images, we apply both silhouette and keypoints loss.

We evaluate the quality of keypoint transfer [35], by generating source-target image pairs from the PASCAL validation list. For each source image, we project the visible vertices of the predicted mesh onto it and align each annotated 2D groundtruth keypoint with its closest vertex. This vertex is projected onto the target image to compute the error with the ground-truth target points, quantified by PCK@0.1. Table 3a shows that Dessie outperforms the latest learning-based SO-TAs. Remarkably, 3D Fauna is trained on a dataset larger than MagicPony.

We further conduct a quantitative 3D evaluation on a subset of PFERD [39] by computing the Chamfer Distance (CD) after Procrustes alignment between the 3D GT and predicted results. As shown in Table 3b, Dessie outperforms other SOTAs.

Table 3: Dessie v.s. SOTA methods

(a) Keypoint Transfer Tasks on Pascal.†: number taken from [24, 44].

 	[].
Network	PCK _{@0.1} ↑
UMR [42]†	28.4
A-CSM [35]†	32.9
MagicPony [74] [†]	42.9
Farm3D [24]†	42.5
3D Fauna [44]†	53.9
DinoHMR	33.8
Dessie	37.1
$DinoHMR \star$	54.3
$Dessie \star$	55.8

(b) 3D evaluation on PFERD [39].

Network	Dessie*	MagicPony	3D Fauna
CD (mm) \downarrow	11.85	32.28	30.10
		0	

Note: *: Model fine-tuned on the whole dataset in [74].

Qualitative and Generalization Results: DinoHMR vs. Dessie. We qualitatively compare DinoHMR and Dessie, trained with and without real-world data, as depicted in Fig 5. Incorporating additional real-world data in training enhances both models' performance. Both methods exhibit robust capabilities across different image domains, even when applied to out-of-distribution data (the last three rows in Fig 5), highlighting their adaptability and effectiveness.

Qualitative Results: Dessie vs. SOTA. Fig. 6 showcases Dessie versus SOTA methods, highlighting better mesh reconstruction results. In addition to shape and pose reconstruction, it has better global orientation prediction in challenging scenarios where MagicPony and Farm3D struggle as shown in the last rows.

Qualitative Results: Dessie with other species. Fig. 7(a,b) demonstrates Dessie reconstructing the 3D mesh better for horse-like and other distinct species. Despite being fine-tuned solely on limited horse data, Dessie captures better poses (Fig. 7(c)) compared to MagicPony, which is also trained on horse data.

5 Conclusion

In this work, we introduce DinoHMR and Dessie, two simple and effective frameworks for horse mesh recovery based on the vision foundational model DINO, utilizing the synthetic data generation pipeline DessiePIPE. The performances further increase with finetuning using real images. Our experiments show that Dessie generalizes well to unseen real images, both horses with new skin patterns

Dessie 13



Fig. 5: DinoHMR and Dessie before and after* real-world data fine-tuning.

and shapes, but also similar species like zebra, cow and deer. Dessie is a step towards a viable solution, particularly, for rare species where the scarcity of data poses a big challenge to training comprehensive models from scratch. Our results indicate that even with a minimal dataset for fine-tuning, Dessie maintains its efficacy, underscoring its applicability for scarce-data animal applications.

Limitations. Currently, Dessie operates on cropped images with a single subject, which poses challenges if this setting is violated (Fig. 8, right). To overcome the limited shape expressiveness of the hSMAL model, which generalizes poorly to shapes very different from those in the training set (Fig. 8, left), the latest horse shape model, VAREN [85] can be used. Finally, while model-based approaches can resolve ambiguities deriving from monocular settings, they require learning a shape and pose prior, which can be challenging for the animal species of interest.

Acknowledgements. We thank Qingwen Zhang, Alfredo Reichlin, and Giovanni Luca Marchetti for the fruitful discussions. Silvia Zuffi is supported by PNRR FAIR Future AI Research (PE00000013), Spoke 8 Pervasive AI (CUP H97G22000210007) under the NRRP MUR program by NextGenerationEU. Elin Hernlund is supported by an SLU career grant no. 57122058, and the Beijer foundation. The computations were enabled by the supercomputing resource Berzelius provided by the National Supercomputer Centre at Linköping University and the Knut and Alice Wallenberg Foundation, Sweden.



Fig. 6: Qualitative comparison between SOTA and Dessie. Non-Dessie figures are taken from the original publications [24, 44, 74, 78], except for the third row.



Fig. 7: Reconstructions of other species not present in the training set.



Fig. 8: Failure cases. Left: Novel shapes. Right: Two subjects.

References

- Amir, S., Gandelsman, Y., Bagon, S., Dekel, T.: Deep vit features as dense visual descriptors. ECCVW What is Motion For? (2022)
- Badger, M., Wang, Y., Modh, A., Perkes, A., Kolotouros, N., Pfrommer, B.G., Schmidt, M.F., Daniilidis, K.: 3d bird reconstruction: a dataset, model, and shape recovery from a single view. In: European Conference on Computer Vision. pp. 1–17 (2020)
- Biggs, B., Boyne, O., Charles, J., Fitzgibbon, A., Cipolla, R.: Who left the dogs out? 3D animal reconstruction with expectation maximization in the loop. In: European Conference on Computer Vision (2020)
- Biggs, B., Roddick, T., Fitzgibbon, A., Cipolla, R.: Creatures great and smal: Recovering the shape and motion of animals from video. In: Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part V 14. pp. 3–19. Springer (2019)
- Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14. pp. 561–578. Springer (2016)
- Borenstein, E., Ullman, S.: Class-specific, top-down segmentation. In: Computer Vision—ECCV 2002: 7th European Conference on Computer Vision Copenhagen, Denmark, May 28–31, 2002 Proceedings, Part II 7. pp. 109–122. Springer (2002)
- Cao, J., Tang, H., Fang, H.S., Shen, X., Lu, C., Tai, Y.W.: Cross-domain adaptation for animal pose estimation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9498–9507 (2019)
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021)
- Cashman, T.J., Fitzgibbon, A.W.: What shape are dolphins? building 3d morphable models from 2d images. IEEE transactions on pattern analysis and machine intelligence 35(1), 232–244 (2012)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
- Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., Abbeel, P.: Infogan: Interpretable representation learning by information maximizing generative adversarial nets. Advances in neural information processing systems 29 (2016)
- 12. Council, A.H.: Economic impact study (2023), https://horsecouncil.org/ economic-impact-study/
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
- Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. International journal of computer vision 111, 98–136 (2015)
- Geman, S.: Statistical methods for tomographic image reconstruction. Bull. Int. Stat. Inst 4, 5–21 (1987)

- 16 C. Li et al.
- Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision (2015)
- Goel, S., Pavlakos, G., Rajasegaran, J., Kanazawa^{*}, A., Malik^{*}, J.: Humans in 4D: Reconstructing and tracking humans with transformers. In: International Conference on Computer Vision (ICCV) (2023)
- Hao, S., Han, K., Wong, K.Y.K.: Learning attention as disentangler for compositional zero-shot learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-vae: Learning basic visual concepts with a constrained variational framework. In: International conference on learning representations (2016)
- 22. Hu, X., Li, X., Busam, B., Zhou, Y., Leonardis, A., Yuan, S.: Disentangling 3d attributes from a single 2d image: Human pose, shape and garment. The British Machine Vision Conference (2022)
- 23. IBISWorld: Horse & other equine production in the us market size, industry analysis, trends and forecasts (2023), https://www.ibisworld.com/ united-states/market-research-reports/horse-other-equine-productionindustry/#ExternalEnvironment
- Jakab, T., Li, R., Wu, S., Rupprecht, C., Vedaldi, A.: Farm3D: Learning articulated 3d animals by distilling 2d diffusion. In: 2024 International Conference on 3D Vision (3DV). IEEE (2024)
- Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7122–7131 (2018)
- 26. Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3d human dynamics from video. In: Computer Vision and Pattern Regognition (CVPR) (2019)
- Kearney, S., Li, W., Parsons, M., Kim, K.I., Cosker, D.: Rgbd-dog: Predicting canine pose from rgbd sensors. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8336–8345 (2020)
- Keller, M., Arora, V., Dakri, A., Chandhok, S., Machann, J., Fritsche, A., Black, M.J., Pujades, S.: HIT: Estimating internal human implicit tissues from the body surface. In: Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR) (Jun 2024)
- Keller, M., Werling, K., Shin, S., Delp, S., Pujades, S., Liu, C.K., Black, M.J.: From skin to skeleton: Towards biomechanically accurate 3d digital humans. In: ACM ToG, Proc. SIGGRAPH Asia (Dec 2023)
- Kocabas, M., Huang, C.P., Hilliges, O., Black, M.J.: Pare: Part attention regressor for 3d human body estimation. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11107–11117. IEEE Computer Society, Los Alamitos, CA, USA (oct 2021)
- Kocabas, M., Athanasiou, N., Black, M.J.: Vibe: Video inference for human body pose and shape estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5253–5263 (2020)

- Kolotouros, N., Pavlakos, G., Jayaraman, D., Daniilidis, K.: Probabilistic modeling for human mesh recovery. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 11585–11594. IEEE Computer Society, Los Alamitos, CA, USA (oct 2021)
- 33. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2252–2261 (2019)
- Kotovenko, D., Sanakoyeu, A., Lang, S., Ommer, B.: Content and style disentanglement for artistic style transfer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4422–4431 (2019)
- Kulkarni, N., Gupta, A., Fouhey, D.F., Tulsiani, S.: Articulation-aware canonical surface mapping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 452–461 (2020)
- 36. Lassner, C., Romero, J., Kiefel, M., Bogo, F., Black, M.J., Gehler, P.V.: Unite the people: Closing the loop between 3d and 2d human representations. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4704–4713. IEEE Computer Society, Los Alamitos, CA, USA (jul 2017)
- 37. Li, C., Lee, G.H.: From synthetic to real: Unsupervised domain adaptation for animal pose estimation. CVPR (2021)
- Li, C., Ghorbani, N., Broomé, S., Rashid, M., Black, M.J., Hernlund, E., Kjellström, H., Zuffi, S.: hsmal: Detailed horse shape and pose reconstruction for motion pattern recognition. arXiv preprint arXiv:2106.10102 (2021)
- 39. Li, C., Mellbin, Y., Krogager, J., Polikovsky, S., Holmberg, M., Ghorbani, N., Black, M.J., Kjellström, H., Zuffi, S., Hernlund, E.: The poses for equine research dataset (pferd). Scientific Data 11(1), 497 (2024)
- 40. Li, J., Xu, C., Chen, Z., Bian, S., Yang, L., Lu, C.: Hybrik: A hybrid analyticalneural inverse kinematics solution for 3d human pose and shape estimation. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3382–3392. IEEE Computer Society, Los Alamitos, CA, USA (jun 2021)
- 41. Li, J., Bian, S., Liu, Q., Tang, J., Wang, F., Lu, C.: Niki: Neural inverse kinematics with invertible neural networks for 3d human pose and shape estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12933–12942 (2023)
- 42. Li, X., Liu, S., Kim, K., De Mello, S., Jampani, V., Yang, M.H., Kautz, J.: Selfsupervised single-view 3d reconstruction via semantic consistency. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16. pp. 677–693. Springer (2020)
- Li, Z., Liu, J., Zhang, Z., Xu, S., Yan, Y.: Cliff: Carrying location information in full frames into human pose and shape estimation. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) Computer Vision – ECCV 2022. pp. 590–606. Springer Nature Switzerland, Cham (2022)
- 44. Li, Z., Litvak, D., Li, R., Zhang, Y., Jakab, T., Rupprecht, C., Wu, S., Vedaldi, A., Wu, J.: Learning the 3d fauna of the web. In: Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR) (2024)
- 45. Liautard, A.F.A.: Lameness of horses and diseases of the locomotory apparatus. WR Jenkins (1888)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision– ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014)

- 18 C. Li et al.
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. ACM transactions on graphics (TOG) 34(6), 1–16 (2015)
- 48. Lorenz, D., Bereska, L., Milbich, T., Ommer, B.: Unsupervised part-based disentangling of object shape and appearance. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10955–10964 (2019)
- Luo, H., Xu, T., Jiang, Y., Zhou, C., Qiu, Q., Zhang, Y., Yang, W., Xu, L., Yu, J.: Artemis: Articulated neural pets with appearance and motion synthesis. ACM Trans. Graph. (2022)
- Luo, X., Hu, Y., Gao, Z., Guo, H., Su, Y.: Automated measurement of livestock body based on pose normalisation using statistical shape model. Biosystems Engineering (2023)
- 51. Mathis, A., Biasi, T., Schneider, S., Yuksekgonul, M., Rogers, B., Bethge, M., Mathis, M.W.: Pretraining boosts out-of-domain robustness for pose estimation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1859–1868 (2021)
- Mu, J., Qiu, W., Hager, G.D., Yuille, A.L.: Learning from synthetic animals. CVPR (2020)
- Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII 14. pp. 483–499. Springer (2016)
- 54. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Howes, R., Huang, P.Y., Xu, H., Sharma, V., Li, S.W., Galuba, W., Rabbat, M., Assran, M., Ballas, N., Synnaeve, G., Misra, I., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision (2023)
- Osman, A.A., Bolkart, T., Black, M.J.: Star: Sparse trained articulated human body regressor. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16. pp. 598–613. Springer (2020)
- Osman, A.A., Bolkart, T., Tzionas, D., Black, M.J.: Supr: A sparse unified partbased human body model. In: European Conference on Computer Vision (ECCV). vol. 2 (2022)
- 57. Pavlakos, G., Zhu, L., Zhou, X., Daniilidis, K.: Learning to estimate 3d human pose and shape from a single color image. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 459–468. IEEE Computer Society, Los Alamitos, CA, USA (jun 2018)
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10975–10985 (2019)
- Qing, Z., Zhang, S., Huang, Z., Zhang, Y., Gao, C., Zhao, D., Sang, N.: Disentangling spatial and temporal learning for efficient image-to-video transfer learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2023)
- Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv:2007.08501 (2020)
- 61. Research, T.M.: Equine healthcare market (2022), https://www.transparencymarketresearch.com/equine-healthcare-market.html

- 62. Richardson, E., Metzer, G., Alaluf, Y., Giryes, R., Cohen-Or, D.: Texture: Textguided texturing of 3d shapes. In: ACM SIGGRAPH (2023)
- 63. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
- 64. Rüegg, N., Tripathi, S., Schindler, K., Black, M.J., Zuffi, S.: Bite: Beyond priors for improved three-d dog pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8867–8876 (2023)
- Rüegg, N., Zuffi, S., Schindler, K., Black, M.J.: Barc: Learning to regress 3d dog shape from images by exploiting breed information. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3876– 3884 (2022)
- 66. Shooter, M., Malleson, C., Hilton, A.: Digidogs: Single-view 3d pose estimation of dogs using synthetic training data. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops. pp. 80–89 (January 2024)
- 67. Stathopoulos, A., Pavlakos, G., Han, L., Metaxas, D.N.: Learning articulated shape with keypoint pseudo-labels from web images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13092–13101 (2023)
- Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5693–5703 (2019)
- Tian, Y., Zhang, H., Liu, Y., Wang, L.: Recovering 3d human mesh from monocular images: A survey. IEEE transactions on pattern analysis and machine intelligence (2023)
- Tschernezki, V., Laina, I., Larlus, D., Vedaldi, A.: Neural feature fusion fields: 3d distillation of self-supervised 2d image representations. In: 2022 International Conference on 3D Vision (3DV). pp. 443–453. IEEE (2022)
- Tumanyan, N., Bar-Tal, O., Bagon, S., Dekel, T.: Splicing vit features for semantic appearance transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10748–10757 (2022)
- 72. Wang, Y., Kolotouros, N., Daniilidis, K., Badger, M.: Birds of a feather: Capturing avian shape models from images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14739–14749 (2021)
- Wen, B., Yang, W., Kautz, J., Birchfield, S.: Foundationpose: Unified 6d pose estimation and tracking of novel objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17868–17879 (2024)
- Wu, S., Li, R., Jakab, T., Rupprecht, C., Vedaldi, A.: MagicPony: Learning articulated 3d animals in the wild. In: CVPR (2023)
- Xu, J., Zhang, Y., Peng, J., Ma, W., Jesslen, A., Ji, P., Hu, Q., Zhang, J., Liu, Q., Wang, J., et al.: Animal3d: A comprehensive dataset of 3d animal pose and shape. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9099–9109 (2023)
- Xu, Y., Zhang, J., Zhang, Q., Tao, D.: ViTPose+: Vision transformer foundation model for generic body pose estimation. arXiv preprint arXiv:2212.04246 (2022)
- 77. Yao, C.H., Hung, W.C., Li, Y., Rubinstein, M., Yang, M.H., Jampani, V.: Lassie: Learning articulated shapes from sparse image ensemble via 3d part discovery. Advances in Neural Information Processing Systems 35, 15296–15308 (2022)
- 78. Yao, C.H., Hung, W.C., Li, Y., Rubinstein, M., Yang, M.H., Jampani, V.: Hi-lassie: High-fidelity articulated shape and skeleton discovery from sparse image ensemble.

In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4853–4862 (2023)

- Yao, C.H., Raj, A., Hung, W.C., Rubinstein, M., Li, Y., Yang, M.H., Jampani, V.: Artic3d: Learning robust articulated 3d shapes from noisy web image collections. Advances in Neural Information Processing Systems 36 (2024)
- Yuan, Y., Iqbal, U., Molchanov, P., Kitani, K., Kautz, J.: Glamr: Global occlusionaware human mesh recovery with dynamic cameras. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11028–11039. IEEE Computer Society, Los Alamitos, CA, USA (jun 2022)
- Zhao, L., Wang, Y., Zhao, J., Yuan, L., Sun, J.J., Schroff, F., Adam, H., Peng, X., Metaxas, D., Liu, T.: Learning view-disentangled human pose representation by contrastive cross-view mutual information maximization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12793– 12802 (2021)
- Zhou, K., Bhatnagar, B.L., Pons-Moll, G.: Unsupervised shape and pose disentanglement for 3d meshes. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16. pp. 341–357. Springer (2020)
- Zuffi, S., Kanazawa, A., Berger-Wolf, T., Black, M.J.: Three-d safari: Learning to estimate zebra pose, shape, and texture from images" in the wild". In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5359–5368 (2019)
- Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3d menagerie: Modeling the 3d shape and pose of animals. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6365–6373 (2017)
- Zuffi, S., Mellbin, Y., Li, C., Hoeschle, M., Kjellström, H., Polikovsky, S., Hernlund, E., Black, M.J.: VAREN: Very accurate and realistic equine network. CVPR (2024)