

Who Left the Dogs Out?

3D Animal Reconstruction with Expectation Maximization in the Loop

Benjamin Biggs¹, Oliver Boyne¹, James Charles¹,
Andrew Fitzgibbon², and Roberto Cipolla¹

¹ Department of Engineering, University of Cambridge, Cambridge, UK
{bjb56,ob312,jjc75,rc10001}@cam.ac.uk

² Microsoft, Cambridge, UK awf@microsoft.com

Abstract. We introduce an automatic, end-to-end method for recovering the 3D pose and shape of dogs from monocular internet images. The large variation in shape between dog breeds, significant occlusion and low quality of internet images makes this a challenging problem. We learn a richer prior over shapes than previous work, which helps regularize parameter estimation. We demonstrate results on the Stanford Dog Dataset, an ‘in the wild’ dataset of 20,580 dog images for which we have collected 2D joint and silhouette annotations to split for training and evaluation. In order to capture the large shape variety of dogs, we show that the natural variation in the 2D dataset is enough to learn a detailed 3D prior through expectation maximization (EM). As a by-product of training, we generate a new parameterized model (including limb scaling) SMBLD which we release alongside our new annotation dataset *StanfordExtra* to the research community.

1 Introduction

Animals contribute greatly to our society, in numerous ways both economic and otherwise (there are more than 63 million pet dogs in the US alone [3]). Consequently, there has been considerable attention in the computer vision research community to the interpretation of animal imagery. Although these techniques share similarities to those used for understanding images of humans, a key difference is that obtaining labelled training data for animals is more difficult than for humans. This is due to the wide range of shapes and species of animals, and the difficulty of educating manual labellers in animal physiology.

Dogs are a particular species of interest, however it is noticeable that existing work has not yet demonstrated effective 3D reconstruction of dogs over large test sets. We postulate that this is partially because dog breeds are remarkably dissimilar in shape and texture, presenting a challenge to the current state of the art. The methods we propose extend the state of the art in several ways. While each of these qualities exist in some existing works, we believe ours is the first to exhibit this combination, leading to a new state of the art in terms of scale and object diversity.

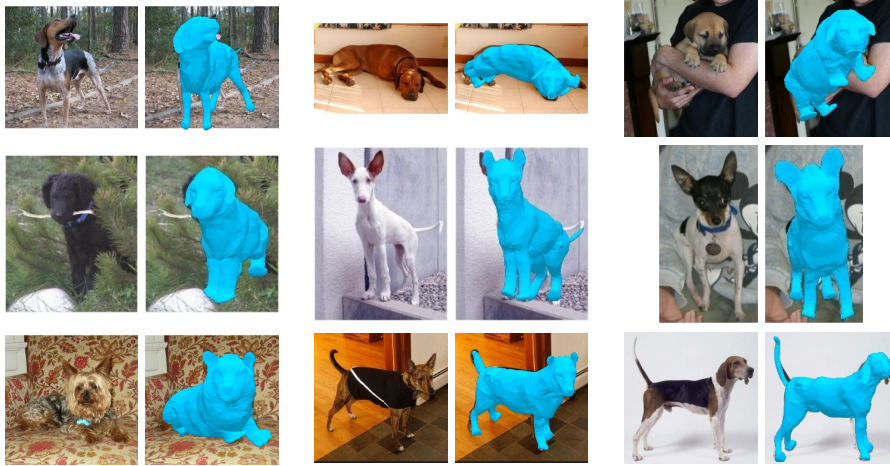


Fig. 1. End-to-end 3D dog reconstruction from monocular images. We propose a novel method that, given a monocular input image of a dog, directly predicts a set of SMBLD parameters to generate an accurate 3D dog model consistent in terms of shape and pose with the input. We regularize learning using a multi-modal shape prior, which is tuned during training with an expectation maximization scheme.

1. We reconstruct pose and shape on a test set of 1703 low-quality internet images of a complex 3D object class (dogs).
2. We directly regress to object pose and shape from a single image without a model fitting stage.
3. We use easily obtained 2D annotations in training, and none at test time.
4. We incorporate fitting of a new multi-modal prior into the training phase (via EM update steps), rather than fitting it to 3D data as in previous work.
5. We introduce new degrees of freedom to the SMAL model, allowing explicit scaling of subparts.

1.1 Related work

The closest work in terms of scale is the category-specific mesh reconstruction of Kanazawa et al. [15], where 2850 images of birds were reconstructed. However, doing so for the complex pose and shape variations of dogs required the advances described in this paper.

Table 1 summarizes previous work on animal reconstruction. It is interesting to note that while several papers demonstrate reconstruction across species, which *prima facie* is a richer class than just dogs, the test-time requirements (e.g. manually-clicked keypoints/silhouette segmentations, input image quality etc.) are considerably higher for those systems. Thus we claim that the achievement of reconstructing a full range of dog breeds, with variable fur length, varying shape and pose of ears, and with considerable occlusion, is a significant contribution.

| Paper | Animal Class | Training requirements | Template Model | Video required | Test Time Annotation | Model Fitting | Test Size |
|---|-------------------|-----------------------------------|-----------------------|----------------------|-----------------------------|---------------|-----------|
| This paper | Dogs | J2, S2, T3, P3 | SMAL | No | None | No | 1703 |
| 3D-Safari [32] | Zebras, horses | M3 (albeit synthetic), J2, S2, P3 | SMAL | 3-7 frames / animal | None | Yes | 200 |
| Lions, Tigers and Bears (SMALR) [33] | MLQ | Not trained | SMAL | 3-7 frames / animal | J2, S2 | Yes | 14 |
| 3D Menagerie (SMAL) [34] | MLQ | Not trained | SMAL | No | J2, S2 | Yes | 48 |
| Creatures Great and SMAL [5] | MLQ | Not trained | SMAL | Yes | S2 (for best results shown) | Yes | 9 |
| Category Specific Mesh Reconstructions [15] | Birds | J2, S2 | Bird convex hull | No | None | No | 2850 |
| What Shape are Dolphins [7] | Dolphins, Pigeons | Not trained | Dolphin Template | 25 frames / category | J2, S2 | Yes | 25 |
| Animated 3D Creatures [29] | MLQ | Not trained | Generalized Cylinders | Yes | J2, S2 | Yes | 15 |

Table 1. Literature summary: Our paper extends large-scale ‘in the wild’ reconstruction to the difficult class of diverse breeds of dogs. MLQ: Medium-to-large quadrupeds. J2: 2D Joints. S2: 2D Silhouettes. T3: 3D Template. P3: 3D Priors. M3: 3D Model.

Monocular 3D reconstruction of human bodies The majority of recent work in 3D pose and shape recovery from monocular images tackles the special case of 3D *human* reconstruction. As a result, the research community has collected a multitude of open source human datasets which provide strong supervisory signals for training deep neural networks. These include accurate 3D deformable template models [23] generated from real human scans, 3D motion capture datasets [11, 24] and large 2D datasets [22, 12, 4] which provide keypoint and silhouette annotations.

The abundance of available human data has supported the development of successful monocular 3D reconstruction pipelines [21, 13]. Such approaches rely on accurate 3D data to build detailed priors over the distribution of human shapes and poses, and use large 2D keypoint datasets to promote generalization to ‘in the wild’ scenarios. Silhouette data has also been shown to assist in accurate reconstruction of clothes, hair and other appearance detail [30, 2]. While the dominant paradigm in human reconstruction is now end-to-end deep learning methods, SPIN [20] shows impressive improvement by incorporating an energy minimization process within their training loop to further minimize a 2D reprojection loss subject to fixed pose & shape priors. Inspired by this innovation, we learn an iteratively-improving shape prior by applying expectation maximization during the training process.

Monocular 3D reconstruction of animal categories. While animals are often featured in computer vision literature, there are still relatively few works that focus on accurate 3D animal reconstruction.

A primary reason for this is absence of large scale 3D datasets³ stemming from the practical challenges associated with 3D motion capture, as well as a lack of 2D data which captures a wide variety of animals. The recent Animal Pose dataset [6] is one such 2D alternative, but contains significantly fewer labelled images than in our new StanfordDogs dataset (4,000 compared to 20,580). On the other hand, animal silhouette data is plentiful [22, 9, 18].

Zuffi et al. [34] made a significant contribution to 3D animal reconstruction research by releasing SMAL, a deformable 3D quadruped model (analogous to SMPL [23] for human reconstruction) from 41 scans of artist-designed toy figurines. The authors also released shape and pose priors generated from artist data. In this work we develop *SMBLD*, an extension of SMAL that better represents the diverse dog category by adding scale parameters and refining the shape prior using our large image dataset.

While there have been various ‘model-free’ approaches which do not rely on an initial template model to generate the 3D animal reconstruction, these techniques often do not produce a mesh [1, 26] or rely heavily on input 2D keypoints or video at test-time [31, 28]. An exception is the end-to-end network of Kanazawa et al. [15], although we argue that the bird category exhibits more limited articulation than our dog category.

We instead focus on model-based approaches. The SMAL authors [34] demonstrate fitting their deformable 3D model to quadruped species using user-provided keypoint and silhouette dataset. SMALR [33] then demonstrated fitting to broader animal categories by incorporating multi-view constraints from video sequences. Biggs et al. [5] overcame the need for hand-clicked keypoints by training a joint predictor on synthetic data. 3D-Safari [32] further improve by training a deep network on synthetic data (built using SMALR [33]) to recover detailed zebra shapes ‘in the wild’.

A drawback of these approaches is their reliance on a test-time energy-based optimization procedure, which is susceptible to failure with poor quality keypoint/silhouette predictions and increases the computational burden. By contrast our method requires no additional energy-based refinement, and is trained purely from single ‘in the wild’ images. The experimental section of this paper contains a robust comparison between our end-to-end method and relevant optimization-based approaches.

A major impediment to research in 3D animal reconstruction has been the lack of a strong evaluation benchmark, with most of the above methods showing only qualitative evaluations or providing quantitative results on fewer than 50 examples. To remedy this, we introduce *StanfordExtra*, a new large-scale dataset which we hope will drive further progress in the field.

³ Released after the submission of this paper, RGBD-Dog dataset [17] is the first open-source 3D motion capture dataset for dogs.

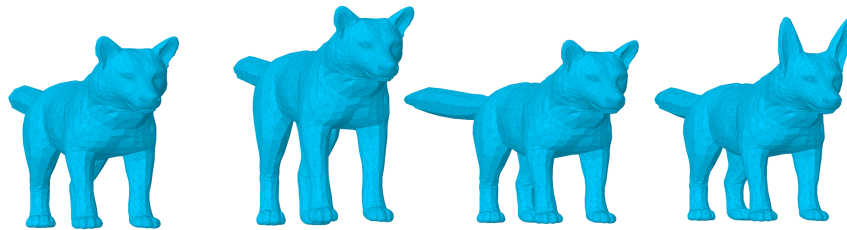


Fig. 3. Effect of varying SMBLD scale parameters. *From left to right: Mean SMBLD model, 25% leg elongation, 50% tail elongation, 50% ear elongation.*

scale parts of the mesh. For each model joint, we define parameters $\kappa_x, \kappa_y, \kappa_z$ which apply a local scaling of the mesh along the local coordinate x, y, z axes, before pose is applied. Allowing each joint to scale entirely independently can however lead to unrealistic deformations, so we share scale parameters between multiple joints, e.g. leg lengths. The new Skinned Multi-Breed Linear Model for Dogs (SMBLD) is therefore adapted from SMAL by adding 6 scale parameters to the existing set of shape parameters. Figure 3 shows how introducing scale parameters increases the flexibility of the SMAL model. We also extend the provided SMAL shape prior (which later initializes our EM procedure) to cover the new scale parameters by fitting SMBLD to a set of 13 artist-designed 3D dog meshes. Further details left to the supplementary.

3 End-to-end dog reconstruction from monocular images

We now consider the task of reconstructing a 3D dog mesh from a monocular image. We achieve this by training an end-to-end convolutional network that predicts a set of SMBLD model and perspective camera parameters. In particular, we train our network to predict pose θ and shape β SMBLD parameters together with translation t and focal length f for a perspective camera. A complete overview of the proposed system is shown in Figure 2.

3.1 Model architecture

Our network architecture is inspired by the model of 3D-Safari [32]. Given an input image cropped to (224, 224), we apply a Resnet-50 [10] backbone network to encode a 1024-dimensional feature map. These features are passed through various linear prediction heads to produce the required parameters. The pose, translation and camera prediction modules follow the design of 3D-Safari, but we describe the differences in our shape module.

Pose, translation and camera prediction. These modules are independent multi-layer perceptrons which map the above features to the various parameter types. As with 3D-Safari we use two linear layers to map to a set of 35×3 3D pose parameters (three parameters for each joint in the SMBLD kinematic tree) given in Rodrigues form. We use independent heads to predict camera frame

translation $t_{x,y}$ and depth t_z independently. We also predict the focal length of the perspective camera similarly to 3D-Safari.

Shape and scale prediction. Unlike 3D-Safari, we design our network to predict the set of shape parameters (including scale) rather than vertex offsets. We observe improvement by handling the standard 20 blend-shape parameters and our new scale parameters in separate linear prediction heads. We retrieve the scale parameters by $\kappa = \exp x$ where x are the network predictions, as we find predicting log scale helps stabilise early training.

3.2 Training losses

A common approach for training such an end-to-end system would be to supervise the prediction of (θ, β, t, f) with 3D ground truth annotations [20, 14, 27]. However, building a suitable 3D annotation dataset would require an experienced graphics artist to design an accurate ground truth mesh for each of 20,520 StanfordExtra dog images, a prohibitive expense.

We instead develop a method that instead relies on *weak 2D supervision* to guide network training. In particular, we rely on only 2D keypoints and silhouette segmentations, which are significantly cheaper to obtain.

The rest of this section describes the set of losses used to supervise the network at train time.

Joint reprojection. The most important loss to promote accurate limb positioning is the joint reprojection loss L_{joints} , which compares the projected model joints $\pi(F_J(\theta, \beta), t, f)$ to the ground truth annotations \hat{X} . Given the parameters predicted by the network, we apply the SMLD model to transform the pose and shape parameters into a set of 3D joint positions $J \in \mathbb{R}^{35 \times 3}$, and project them to the image plane using translation and camera parameters. The joint loss L_{joints} is given by the ℓ_2 error between the ground truth and projected joints:

$$L_{\text{joints}}(\theta, \beta, t, f; \hat{X}) = \|\hat{X} - \pi(F_J(\theta, \beta), t, f)\|_2 \quad (1)$$

Note that many of our training images exhibit significant occlusion, so \hat{X} contains many invisible joints. We handle this by masking L_{joints} to prevent invisible joints contributing to the loss.

Silhouette loss. The silhouette loss L_{sil} is used to promote shape alignment between the SMLD dog mesh and the input dog. In order to compute the silhouette loss, we define a rendering function $R : (\nu, t, f) \mapsto S$ which projects the SMLD mesh to produce a binary segmentation mask. In order to allow derivatives to be propagated through R , we implement R using the differentiable Neural Mesh Renderer [16]. The loss is computed as the ℓ_2 difference between a projected silhouette and the ground truth mask \hat{S} :

$$L_{\text{sil}}(\theta, \beta, t, f; \hat{S}) = \|\hat{S} - R(F_V(\theta, \beta), t, f)\|_2 \quad (2)$$

Priors. In the absence of 3D ground truth training data, we rely on priors obtained from artist graphics models to encourage realism in the network predictions. We model both pose and shape using a multivariate Gaussian prior, consisting of means μ_θ, μ_β and covariance matrices $\Sigma_\theta, \Sigma_\beta$. The loss is given as the log likelihood of a given shape or pose vector under these distributions, which corresponds to the Mahalanobis distance between the predicted parameters and their corresponding means:

$$L_{\text{pose}}(\theta; \mu_\theta, \Sigma_\theta) = (\theta - \mu_\theta)^T \Sigma_\theta^{-1} (\theta - \mu_\theta) \quad (3)$$

$$L_{\text{shape}}(\beta; \mu_\beta, \Sigma_\beta) = (\beta - \mu_\beta)^T \Sigma_\beta^{-1} (\beta - \mu_\beta) \quad (4)$$

Unlike previous work, we find there is no need to use a loss to penalize pose parameters if they exceed manually specified joint angle limits. We suspect our network learns this regularization naturally because of our large dataset.

3.3 Learning a multi-modal shape prior.

The previous section introduced a unimodal, multivariate Gaussian shape prior, based on mean μ_β and covariance matrix Σ_β . However, we find enforcing this prior throughout training tends to result in predictions which appear similar in 3D shape, even when tested on dog images of different breeds. We propose to improve diversity among predicted 3D dog shapes by extending the above formulation to a mixture of M Gaussians prior. The mixture shape loss is then given as:

$$L_{\text{mixture}}(\beta_i; \mu_\beta, \Sigma_\beta, \Pi_\beta) = \sum_{m=1}^M \Pi_\beta^m (\beta_i - \mu_\beta^m)^T \Sigma_\beta^{m-1} (\beta_i - \mu_\beta^m) \quad (5)$$

$$= \sum_{m=1}^M \Pi_\beta^m L_{\text{shape}}(\beta_i; \mu_\beta^m, \Sigma_\beta^m) \quad (6)$$

Where μ_β^m , Σ_β^m and Π_β^m are the mean, covariance and mixture weight respectively for Gaussian component m . For each component the mean is sampled from our existing unimodal prior, and the covariance is set equal to the unimodal prior i.e. $\Sigma_\beta^m := \Sigma_\beta$. All mixture weights are initially set to $\frac{1}{M}$.

Each training image i is assigned a set of latent variables $\{w_i^1, \dots, w_i^M\}$ encoding the probability of the dog shape in image i being generated by component m .

3.4 Expectation Maximization in the loop

As previously discussed, our initial shape prior is obtained from artist data which we find is unrepresentative of the diverse shapes present in our real dog dataset. We address this by proposing to recover the latent variables w_i^m and parameters $(\mu_\beta^m, \Sigma_\beta^m \text{ and } \Pi_\beta^m)$ of our 3D shape prior by learning from monocular images of ‘in the wild’ dogs and their 2D training labels in our training dataset.

We achieve this using Expectation Maximization (EM), which regularly updates the means and variances for each mixture component and per-image mixture weights based on the observed shapes in the training set. While training our 3D reconstruction network, we progressively update our shape mixture model with an alternating ‘E’ step and ‘M’ step described below:

The ‘E’ Step. The ‘E’ step computes the expected value of the latent variables w_i^m assuming fixed $(\mu_\beta^m, \Sigma_\beta^m, \Pi_\beta^m)$ for all $i \in \{1, \dots, N\}, m \in \{1, \dots, M\}$.

The update equation for an image i with latest shape prediction β_i and cluster m with parameters $(\mu_\beta^m, \Sigma_\beta^m, \Pi_\beta^m)$ is given as:

$$w_i^m := \frac{\mathcal{N}(\beta_i | \mu_\beta^m, \Sigma_\beta^m) \Pi_\beta^m}{\sum_{m'}^M \mathcal{N}(\beta_i | \mu_\beta^{m'}, \Sigma_\beta^{m'}) \Pi_\beta^{m'}} \quad (7)$$

The ‘M’ Step. The ‘M’ step computes new values for $(\mu_\beta^m, \Sigma_\beta^m, \Pi_\beta^m)$, assuming fixed w_i^m for all $i \in \{1, \dots, N\}, m \in \{1, \dots, M\}$.

The update equations are given as follows:

$$\mu_\beta^m := \frac{\sum_i w_i^m \beta_i}{\sum_i w_i^m} \quad \Sigma_\beta^m := \frac{\sum_i w_i^m (\beta_i - \Sigma_\beta^m)(\beta_i - \Sigma_\beta^m)^T}{\sum_i w_i^m} \quad \Pi_\beta^m := \frac{1}{N} \sum_i w_i^m \quad (8)$$

4 Experiments

In this section we compare our method to competitive baselines. We begin by describing our new large-scale dataset of annotated dog images, followed by a quantitative and qualitative evaluation.

4.1 StanfordExtra: A new large-scale dog dataset with 2D keypoint and silhouette annotations

In order to evaluate our method, we introduce StanfordExtra: a new large-scale dataset with annotated 2D keypoints and binary segmentation masks for dogs. We opted to take source images from the existing Stanford Dog Dataset [19], which consists of 20,580 dog images taken ‘in the wild’ and covers 120 dog breeds. The dataset contains vast shape and pose variation between dogs, as well as nuisance factors such as self/environmental occlusion, interaction with humans/other animals and partial views. Figure 4 (left) shows samples from the new dataset.

We used Amazon Mechanical Turk to collect a binary silhouette mask and 20 keypoints per image: 3 per leg (knee, ankle, toe), 2 per ear (base, tip), 2 per tail (base, tip), 2 per face (nose and jaw). We can approximate the difficulty of the dataset by analysing the variance between 3 annotators at both the joint labelling and silhouette task. Figure 4 (right) illustrates typical per-joint variance in joint labelling. Further details of the data curation procedure are left to the supplementary materials.



Fig. 4. StanfordExtra example images. *Left:* outlined segmentations and labelled keypoints for a selection of StanfordExtra images. *Right:* heatmap showing annotator uncertainty in per-keypoint clicking, computed over the entire dataset.

4.2 Evaluation protocol

Our evaluation is based on our new StanfordExtra dataset. In line with other methods which tackle ‘in the wild’ 3D reconstruction of articulated subjects [20, 21], we filter images from the original dataset of 20,580 for which the majority of dog keypoints are invisible. We consider these images unsuitable for our full-body dog reconstruction task. We also remove images for which the consistency in keypoint/silhouette segmentations between the 3 annotators is below a set threshold. This leaves us with 8,476 images which we divide per-breed into an 80%/20% train and test split.

We consider two primary evaluation metrics. IoU is the intersection-over-union of the projected model silhouette compared to the ground truth annotation and indicates the quality of the reconstructed 3D shape. Percentage of Correct Keypoints (PCK) computes the percentage of joints which are within a normalized distance (based on square root of 2D silhouette area) to the ground truth locations, and evaluates the quality of reconstructed 3D pose. We also produce PCK results on various joint groups (legs, tail, ears, face) to compare the reconstruction accuracy for different parts of the dog model.

4.3 Training procedure

We train our model in two stages. The first omits the silhouette loss which we find can lead the network to unsatisfactory local minima if applied too early. With the silhouette loss turned off, we find it satisfactory to use the simple unimodal prior (and without EM) for this preliminary stage since there is no loss to specifically encourage a strong shape alignment. After this, we introduce the silhouette loss, the mixture prior and begin applying the expectation maximization updates over $M = 10$ clusters. We train the first stage for 250 epochs, the second stage for 150 and apply the EM step every 50 epochs. All losses are weighted, as described in the supplementary. The entire training procedure takes 96 hours on a single P100 GPU.

4.4 Comparison to baselines

We first compare our method to various baseline methods. SMAL [34] is an approach which fits the 3D SMAL model using per-image energy minimization. Creatures Great and SMAL (CGAS) [5] is a three-stage method, which employs a joint predictor on silhouette renderings from synthetic 3D dogs, applies a genetic algorithm to clean predictions, and finally applies the SMAL optimizer to produce the 3D mesh.

At test-time both SMAL and CGAS rely on manually-provided segmentation masks, and SMAL also relies on hand-clicked keypoints. In order to produce a fair comparison, we produce a set of *predicted* keypoints for StanfordExtra by training the Stacked Hourglass Network [25] with 8 stacks and 1 block, and *predicted* segmentation masks using DeepLab v3+ [8]. The Stacked Hourglass Network achieves 71.4% PCK score, DeepLab v3+ achieves 83.4% IoU score and the CGAS joint predictor achieves 41.8% PCK score.

Table 2 and Figure 5 show the comparison against competitive methods. For full examination, we additionally provide results for SMAL and CGAS in the scenario that ground-truth keypoints and/or segmentations are available at test time. The results show our end-to-end method outperforms the competitors when they are provided with predicted keypoints/segmentations (white rows). Our method therefore achieves a new state-of-the-art on this 3D reconstruction task. In addition, we show our method achieves improved average IoU/PCK scores than competitive methods, even when they are provided ground truth annotations at test time (grey rows). We also demonstrate wider applicability of two contributions from our work (scale parameters and improved prior) by showing improved performance of the SMAL method when these are incorporated. Finally, our model’s test-time speed is significantly faster than the competitors as it does not require an optimizer.

4.5 Generalization to unseen dataset

Table 3 shows an experiment to compare how well our model generalizes to a new data domain. We test our model against the SMAL [34] method (using predicted keypoints and segmentations as above for fairness) on the recent Animal Pose dataset [6]. The data preparation process is the same as for StanfordExtra and no fine-tuning was used for either method. We achieve strong results in this unseen domain and still improve over the SMAL optimizer.

4.6 Ablation study

We also produce a study in which we ablate individual components of our method and examine the effect on the PCK/IoU performance. We evaluate three variants: (1) **Ours w/o EM** that omits EM updates, (2) **Ours w/o MoG** which replaces our mixture shape prior with a unimodal prior, (3) **Ours w/o Scale** which removes the scale parameters.

The results in Table 4 indicate that each individual component has a positive impact on the overall method performance. In particular, it can be seen that the

| Method | Kps | Seg | IoU | PCK | | | | |
|----------------------------|------|------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | | | Avg | Legs | Tail | Ears | Face |
| SMAL [34] | Pred | Pred | 67.9 | 67.1 | 65.7 | 79.5 | 54.9 | 87.4 |
| SMAL | GT | GT | 69.2 | 72.6 | 69.9 | 92.0 | 58.6 | 96.9 |
| SMAL | GT | Pred | 68.6 | 72.6 | 70.2 | 91.5 | 58.1 | 96.9 |
| SMAL | Pred | GT | 68.5 | 67.4 | 66.0 | 79.9 | 55.0 | 88.2 |
| CGAS [5] | CGAS | Pred | 62.4 | 43.7 | 46.5 | 64.1 | 36.5 | 21.4 |
| CGAS | CGAS | GT | 63.1 | 43.6 | 46.3 | 64.2 | 36.3 | 21.6 |
| SMAL + scaling | Pred | Pred | 69.3 | 69.6 | 69.4 | 79.3 | 56.5 | 87.6 |
| SMAL + scaling + new prior | Pred | Pred | 70.7 | 71.6 | 71.5 | 80.7 | 59.3 | 88.0 |
| Ours | — | — | 73.6 | 75.7 | 75.0 | 77.6 | 69.9 | 90.0 |

Table 2. Quantitative comparison to baselines. PCK and silhouette IoU metrics are shown against competitive methods in various conditions. *White rows*: Experiments which use no test-time ground truth annotations, *Grey rows*: Experiments which either use test-time ground truth annotations (GT) or incorporate components introduced in this paper (SMBLD scale parameters or mixture shape prior).

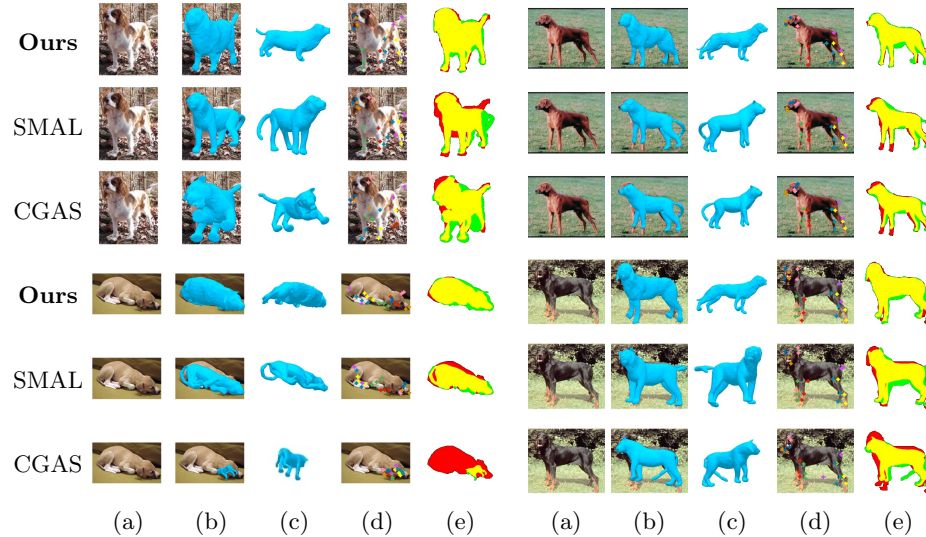


Fig. 5. Qualitative comparison to baselines. Comparison between our method, SMAL [34] and CGAS [5]. (a) Input image, (b) predicted 3D mesh, (c) canonical view 3D mesh, (d) reprojected model joints and (e) silhouette reprojection error.

inclusion of the EM and mixture of Gaussians prior leads to an improvement in IoU, suggesting that the shape prior refinements steps help the model accurately fit the exact dog shape. Interestingly, we notice that adding the mixture of Gaussians prior but omitting EM steps slightly hinders performance, perhaps due to a sub-optimal initialization for the M clusters. However, we find adding EM updates to the mixture of Gaussian model improves all metrics except the

| Method | IoU | PCK | | | | |
|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | Avg | Legs | Tail | Ears | Face |
| SMAL [34] | 63.6 | 69.1 | 60.9 | 83.5 | 75.0 | 93.0 |
| Ours | 66.9 | 73.8 | 65.1 | 85.6 | 84.0 | 93.6 |

Table 3. Animal Pose dataset [6]. Evaluation on recent Animal Pose dataset with no fine-tuning to our method nor joint/silhouette predictors used for SMAL.

| Method | IoU | PCK | | | | |
|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | Avg | Legs | Tail | Ears | Face |
| Ours | 73.6 | 75.7 | 75.0 | 77.6 | 69.9 | 90.0 |
| –EM | 67.7 | 74.6 | 72.9 | 75.2 | 72.5 | 88.3 |
| –MoG | 68.0 | 74.9 | 74.3 | 73.3 | 70.0 | 90.2 |
| –Scale | 67.3 | 72.6 | 72.9 | 75.3 | 62.3 | 89.1 |

Table 4. Ablation study. Evaluation with the following contributions removed: (a) EM updates, (b) Mixture Shape Prior, (c) SMBLD scale parameters.

ear keypoint accuracy. We observe the error here is caused by the our shape prior learning slightly imprecise shapes for dogs with extremely ‘floppy’ ears. Although there is good silhouette coverage for these regions, the fact our model has only a single articulation point per ear causes a lack of flexibility that results in occasionally misplaced ear tips for these instances. This could be improved in future work by adding additional model joints to the ear. Finally, our results show the increased model flexibility afforded by the SMBLD scale parameters has a positive effect on IoU/PCK scores.

4.7 Qualitative evaluation

Figure 5 shows a range of predicted 3D reconstructions from our method when tested on a wide range of StanfordExtra and Animal Pose [6] dogs. The examples include challenging poses, large diversity in dog shape and size, interaction with humans, various environmental conditions and handling occluded keypoints. Note that only StanfordExtra is used for training and we use no fine-tuning to produce results on the Animal Pose dataset.

5 Conclusions

This paper presents an end-to-end method for automatic, monocular 3D dog reconstruction. We achieve this using only weak 2D supervision, provided by our novel StanfordExtra dataset. Furthermore, we show we can learn a more detailed shape prior by tuning a Gaussian mixture during model training and this leads to improved reconstructions. We also show our method improves over competitive baselines, even when they are given access to ground truth data at test time.

Future work should involve tackling some failure cases of our system, for example handling multiple overlapping dogs or dealing with heavy motion blur. Other areas for research include extending our EM formulation to handle video input to take advantage of multi-view shape constraints, and transferring knowledge accumulated through training on StanfordExtra dogs to other species.

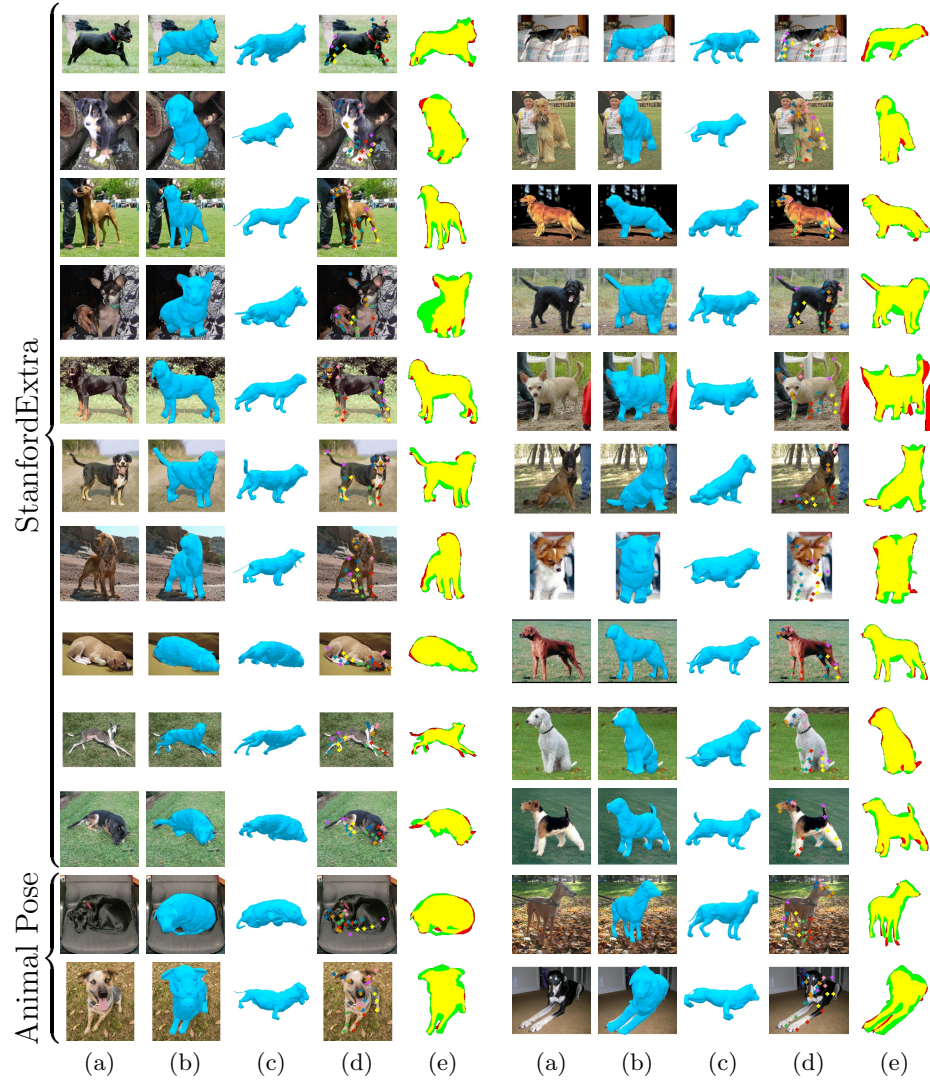


Fig. 6. Qualitative results on StanfordExtra and Animal Pose [6]. For each sample we show: (a) input image, (b) predicted 3D mesh, (c) canonical view 3D mesh, (d) reprojected model joints and (e) silhouette reprojection error.

6 Acknowledgements

The authors would like to thank the GSK AI team for providing access to their GPU cluster, Michael Sutcliffe, Matthew Allen, Thomas Roddick and Peter Fisher for useful technical discussions, and the GSK TDI team for project sponsorship.

References

1. Agudo, A., Pijoan, M., Moreno-Noguer, F.: Image collection pop-up: 3d reconstruction and clustering of rigid and non-rigid categories. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
2. Alldieck, T., Magnor, M., Bhatnagar, B.L., Theobalt, C., Pons-Moll, G.: Learning to reconstruct people in clothing from a single rgb camera. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1175–1186 (2019)
3. American Pet Products Association: 2019-2020 APPA National Pet Owners Survey (2020), <http://www.americanpetproducts.org>
4. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2014)
5. Biggs, B., Roddick, T., Fitzgibbon, A., Cipolla, R.: Creatures great and SMAL: Recovering the shape and motion of animals from video. In: ACCV (2018)
6. Cao, J., Tang, H., Fang, H., Shen, X., Tai, Y., Lu, C.: Cross-domain adaptation for animal pose estimation. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9497–9506 (2019)
7. 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 (2013)
8. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *CoRR abs/1606.00915* (2016)
9. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International journal of computer vision* **88**(2), 303–338 (2010)
10. 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)
11. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence* **36**(7), 1325–1339 (2013)
12. Johnson, S., Everingham, M.: Clustered pose and nonlinear appearance models for human pose estimation. In: Proceedings of the British Machine Vision Conference (2010), doi:10.5244/C.24.12
13. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Computer Vision and Pattern Recognition (CVPR) (2018)
14. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proc. CVPR (2018)
15. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: European Conference on Computer Vision. pp. 371–386 (2018)
16. Kato, H., Ushiku, Y., Harada, T.: Neural 3d mesh renderer. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
17. Kearney, S., Li, W., Parsons, M., Kim, K.I., Cosker, D.: Rgb-dog: Predicting canine pose from rgb-d sensors. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)

18. Khoreva, A., Benenson, R., Ilg, E., Brox, T., Schiele, B.: Lucid data dreaming for object tracking. The 2017 DAVIS Challenge on Video Object Segmentation - CVPR Workshops (2017)
19. Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L.: Novel dataset for fine-grained image categorization. In: First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition. Colorado Springs, CO (June 2011)
20. 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 International Conference on Computer Vision. pp. 2252–2261 (2019)
21. Kolotouros, N., Pavlakos, G., Daniilidis, K.: Convolutional mesh regression for single-image human shape reconstruction. In: Proc. CVPR (2019)
22. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollar, P., Zitnick, L.: Microsoft COCO: Common objects in context. In: ECCV. European Conference on Computer Vision (September 2014), <https://www.microsoft.com/en-us/research/publication/microsoft-coco-common-objects-in-context/>
23. 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), 248 (2015)
24. von Marcard, T., Henschel, R., Black, M., Rosenhahn, B., Pons-Moll, G.: Recovering accurate 3d human pose in the wild using imus and a moving camera. In: European Conference on Computer Vision (ECCV) (sep 2018)
25. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European Conference on Computer Vision. pp. 483–499. Springer (2016)
26. Novotny, D., Ravi, N., Graham, B., Neverova, N., Vedaldi, A.: C3DPO: Canonical 3d pose networks for non-rigid structure from motion. In: Proc. ICCV (2019)
27. Pavlakos, G., Zhu, L., Zhou, X., Daniilidis, K.: Learning to estimate 3D human pose and shape from a single color image. In: Proc. CVPR (2018)
28. Probst, T., Pani Paudel, D., Chhatkuli, A., Van Gool, L.: Incremental non-rigid structure-from-motion with unknown focal length. In: The European Conference on Computer Vision (ECCV) (2018)
29. Reinert B, Ritschel T, S.H.P.: Animated 3d creatures from single-view video by skeletal sketching. In: Proc. Graphics Interface (2016)
30. Saito, S., , Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. arXiv preprint arXiv:1905.05172 (2019)
31. Vicente, S., Agapito, L.: Balloon shapes: Reconstructing and deforming objects with volume from images. In: 2013 International Conference on 3D Vision - 3DV 2013. pp. 223–230 (2013)
32. 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: The IEEE International Conferene on Computer Vision (ICCV) (2019)
33. Zuffi, S., Kanazawa, A., Black, M.J.: Lions and tigers and bears: Capturing non-rigid, 3D, articulated shape from images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE Computer Society (2018)
34. Zuffi, S., Kanazawa, A., Jacobs, D., Black, M.J.: 3D menagerie: Modeling the 3D shape and pose of animals. In: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (Jul 2017)