

3D Morphable Models as Spatial Transformer Networks

Anil Bas^{*}, Patrik Huber[†], William A. P. Smith^{*}, Muhammad Awais[†], Josef Kittler[†]

^{*}Department of Computer Science, University of York, UK

[†]Centre for Vision, Speech and Signal Processing, University of Surrey, UK

{ab1792,william.smith}@york.ac.uk, {p.huber,m.a.rana,j.kittler}@surrey.ac.uk

Abstract

In this paper, we show how a 3D Morphable Model (i.e. a statistical model of the 3D shape of a class of objects such as faces) can be used to spatially transform input data as a module (a 3DMM-STN) within a convolutional neural network. This is an extension of the original spatial transformer network in that we are able to interpret and normalise 3D pose changes and self-occlusions. The trained localisation part of the network is independently useful since it learns to fit a 3D morphable model to a single image. We show that the localiser can be trained using only simple geometric loss functions on a relatively small dataset yet is able to perform robust normalisation on highly uncontrolled images including occlusion, self-occlusion and large pose changes.

1. Introduction

Convolutional neural networks (CNNs) are usually trained with such large amounts of data that they can learn invariance to scale, translation, in-plane rotation and, to a certain degree, out-of-plane rotations, without using any explicit geometric transformation model. However, most networks do require a rough bounding box estimate as input and don't work for larger variations. Recently, Jaderberg et al. [14] proposed the Spatial Transformer Network (STN) - a module that can be incorporated into a neural network architecture, giving the network the ability to explicitly account for the effects of pose and nonrigid deformations (which we refer to simply as "pose"). An STN explicitly estimates pose and then resamples a specific part of the input image to a fixed-size output image. It is thus able to work on inputs with larger translation and pose variation in general, since it can explicitly compensate for it, and feed a transformed region of interest to the subsequent neural network layers. By exploiting and "hard-coding" knowledge of geometric transformation, the amount of training data and the required complexity of the network can be vastly reduced.

In this paper, we show how to use a 3D morphable model as a spatial transformer network (we refer to this as a 3DMM-STN). In this setting, the locations in the input image that are resampled are determined by the 2D projection of a 3D deformable mesh. Hence, our 3DMM-STN estimates both 3D shape and pose. This allows us to explicitly estimate and account for 3D rotations as well as self occlusions. The output of our 3DMM-STN is a resampled image in a flattened 2D texture space in which the images are in dense, pixel-wise correspondence. Hence, this output can be fed to subsequent CNN layers for further processing. We focus on face images and use a 3D morphable face model [4, 19], though our idea is general and could be applied to any object for which a statistical 3D shape model is available (though note that the loss functions proposed in Sections 3.1 and 3.2 do assume that the object is bilaterally symmetric). We release source code for our 3DMM-STN in the form of new layers for the MatConvNet toolbox [30]¹.

1.1. Related work

In a lot of applications, the process of pose normalisation and object recognition are disjoint. For example, in the breakthrough deep learning face recognition paper DeepFace, Taigman et al. [27] use a 3D mean face as preprocessing, before feeding the pose-normalised image to a CNN.

Spatial transformers The original STN [14] aimed to combine these two processes into a single network that is trainable end to end. The localiser network estimated a 2D affine transformation that was applied to the regular output grid meaning the network could only learn a fairly restricted space of transformations. Jaderberg et al. [14] also proposed the concept of a 3D transformer, which takes 3D voxel data as input, applies 3D rotation and translation, and outputs a 2D projection of the transformed data. Working with 3D (volumetric data) removes the need to model occlusion or camera projection parameters. In contrast, we work with

¹The source code is available at <https://github.com/anilbas/3DMMasSTN>.

regular 2D input and output images but transform them via a 3D model.

A number of subsequent works were inspired by the original STN. Yan et al. [31] use an encoder-decoder architecture in which the encoder estimates a 3D volumetric shape from an image and is trained by combining with a decoder which uses a perspective transformer network to compute a 2D silhouette loss. Handa et al. [11] present the gynn (Geometric Vision with Neural Networks) toolbox that, like in this paper, has layers that explicitly implement 3D geometric transformations. However, their goal is very different to ours. Rather than learning to fit a statistical shape model, they seek to use 3D transformations in low level vision tasks such as relative pose estimation. Chen et al. [6] use a spatial transformer that applies a 2D similarity transform as part of an end to end network for face detection. Henriques and Vedaldi [12] apply a spatial warp prior to convolutions such that the convolution result is invariant to a class of two-parameter spatial transformations. Like us, Yu et al. [32] incorporate a parametric shape model, though their basis is 2D (and trainable), models only sparse shape and combines pose and shape into a single basis. They use a second network to locally refine position estimates and train end to end to perform landmark localisation. Bhagavatula et al. [3] fit a generic 3D face model and estimate face landmarks, before warping the projected face model to better fit the landmarks. They estimate 2D landmarks in a 3D-aware fashion, though they require known landmarks for training.

Analysis-by-synthesis Our localiser learns to fit a 3DMM to a single image. This task has traditionally been posed as a problem of analysis-by-synthesis and solved by optimisation. The original method [4] used stochastic gradient descent to minimise an appearance error, regularised by statistical priors. Subsequent work used a more complex feature-based objective function [23] and the state-of-the-art method uses Markov Chain Monte Carlo for probabilistic image interpretation [24].

Supervised CNN regression Analysis-by-synthesis approaches are computationally expensive, prone to convergence on local minima and fragile when applied to in-the-wild images. For this reason, there has been considerable recent interest in using CNNs to directly regress 3DMM parameters from images. The majority of such work is based on supervised learning. Jourabloo and Liu [15] fit a 3DMM to detected landmarks and then train a CNN to directly regress the fitted pose and shape parameters. Tr an et al. [29] use a recent multi-image 3DMM fitting algorithm [20] to obtain pooled 3DMM shape and texture parameters (i.e. the same parameters for all images of the same subject). They then train a CNN to directly regress these parameters from a single image. They do not estimate pose and hence do

not compute an explicit correspondence between the model and image. Kim et al. [16] go further by also regressing illumination parameters (effectively performing inverse rendering) though they train on synthetic, rendered images (using a breeding process to increase diversity). They estimate a 3D rotation but rely on precisely cropped input images such that scale and translation is implicit. Richardson et al. [21] also train on synthetic data though they use an iteratively applied network architecture and a shape-from-shading refinement step to improve the geometry. Jackson et al. [13] regress shape directly using a volumetric representation.

The DenseReg [10] approach uses fully convolutional networks to directly compute dense correspondence between a 3D model and a 2D image. The network does not explicitly estimate or model 3D pose or shape (though these are implied by the correspondence) and is trained by using manually annotated 2D landmarks to warp a 3D template onto the training images (providing the supervision). Sela et al. [25] also use a fully convolutional network to predict correspondence and also depth. They then merge the model-based and data-driven geometries for improved quality.

The weakness of all of these supervised approaches is that they require labelled training data (i.e. images with fitted morphable model parameters). If the images are real world images then the parameters must come from an existing fitting algorithm in which case the best the CNN can do is learn to replicate the performance of an existing algorithm. If the images are synthetic with known ground truth parameters then the performance of the CNN on real world input is limited by the realism and variability present in the synthetic images. Alternatively, we must rely on 3D supervision provided by multiview or RGBD images, in which case the available training data is vastly reduced.

Unsupervised CNN regression Richardson et al. [22] take a step towards removing the need for labels by presenting a semi-supervised approach. They still rely on supervised training for learning 3DMM parameter regression but then refine the coarse 3DMM geometry using a second network that is trained in an unsupervised manner. Very recently, Tewari et al. [28] presented MoFA, a completely unsupervised approach for training a CNN to regress 3DMM parameters, pose and illumination using an autoencoder architecture. The regression is done by the encoder CNN. The decoder then uses a hand-crafted differentiable renderer to synthesise an image. The unsupervised loss is the error between the rendered image and the input, with convergence aided by losses for priors and landmarks. Note that the decoder is exactly equivalent to the differentiable cost function used in classical analysis-by-synthesis approaches. Presumably, the issues caused by the non-convexity of this cost function are reduced in a CNN setting since the gradient

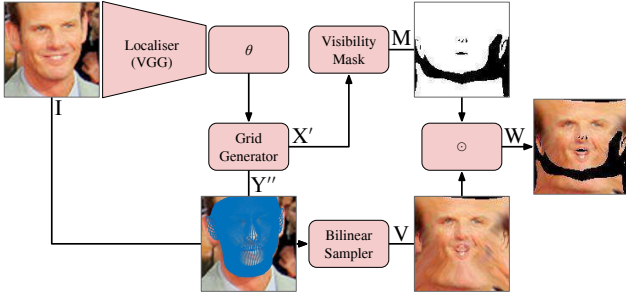


Figure 1. Overview of the 3DMM-STN. The localiser predicts 3DMM shape parameters and pose. The grid generator projects the 3D geometry to 2D. The bilinear sampler resamples the input image to a regular output grid which is then masked by an occlusion mask computed from the estimated 3D geometry.

is averaged over many images.

While the ability of [28] to learn from unlabelled data is impressive, there are a number of limitations. The complexity required to enable the hand-crafted decoder to produce photorealistic images of any face under arbitrary real world illumination, captured by a camera with arbitrary geometric and photometric properties, is huge. Arguably, this has not yet been achieved in computer graphics. Moreover, the 3DMM texture should only capture intrinsic appearance parameters such as diffuse and specular albedo (or even spectral quantities to ensure independence from the camera and lighting). Such a model is not currently available.

1.2. Contributions

In this paper we propose a purely geometric approach in which only the shape component of a 3DMM is used to geometrically normalise an image. Unlike [10, 13, 15, 16, 21, 25, 29], our method can be trained in an unsupervised fashion, and thus does not depend on synthetic training data or the fitting results of an existing algorithm. In contrast to [28], we avoid the complexity and potential fragility of having to model illumination and reflectance parameters. Moreover, our 3DMM-STN can form part of a larger network that performs a face processing task and is trained end to end. Finally, in contrast to all previous 3DMM fitting networks, the output of our 3DMM-STN is a 2D resampling of the original image which contains all of the high frequency, discriminating detail in a face rather than a model-based reconstruction which only captures the gross, low frequency aspects of appearance that can be explained by a 3DMM.

2. 3DMM-STN

Our proposed 3DMM-STN has the same components as a conventional STN, however each component must be modified to incorporate the statistical shape model, 3D transformations and projection and self-occlusion. In this

section we describe each component of a 3DMM-STN and the layers that are required to construct it. We show an overview of our architecture in Figure 1.

2.1. Localiser network

The localiser network is a CNN that takes an image as input and regresses the pose and shape parameters, θ , of the face in the image. Specifically, we predict the following vector of parameters:

$$\theta = (\underbrace{\mathbf{r}, \mathbf{t}, \log s}_{\text{pose}}, \underbrace{\boldsymbol{\alpha}}_{\text{shape}}). \quad (1)$$

Here, $\mathbf{t} \in \mathbb{R}^2$ is a 2D translation, $\mathbf{r} \in \mathbb{R}^3$ is an axis-angle representation of a 3D rotation with rotation angle $\|\mathbf{r}\|$ and axis $\mathbf{r}/\|\mathbf{r}\|$. Since scale must be positive, we estimate log scale and later pass this through an exponentiation layer, ensuring that the estimated scale is positive. The shape parameters $\boldsymbol{\alpha} \in \mathbb{R}^D$ are the principal component weights used to reconstruct the shape.

For our localiser network, we use the pretrained VGG-Faces [18] architecture, delete the classification layer and add a new fully connected layer with $6 + D$ outputs. The weights for the new layer are randomly initialised but scaled so that the elements of the axis-angle vector are in the range $[-\pi, \pi]$ for typical inputs. The whole localiser is then fine-tuned as part of the subsequent training.

2.2. Grid generator network

In contrast to a conventional STN, the warped sampling grid is not obtained by applying a global transformation to the regular output grid. Instead, we apply a 3D transformation and projection to a 3D mesh that comes from the morphable model. The intensities sampled from the source image are then assigned to the corresponding points in a flattened 2D grid.

For this reason, the grid generator network in a 3DMM-STN is more complex than in a conventional STN, although we emphasise that it remains differentiable and hence suitable for use in end to end training. The sample points in our grid generator are determined by the transformation parameters θ estimated by the localiser network. Our grid generator combines a linear statistical model with a scaled orthographic projection as shown in Figure 2. Note that we could alternatively use a perspective projection (modifying the localiser to predict a 3D translation as well as camera parameters such as focal length). However, recent results show that interpreting face shape under perspective is ambiguous [2, 26] and so we use the more restrictive orthographic model here.

We now describe the transformation applied by each layer in the grid generator and provide derivatives.

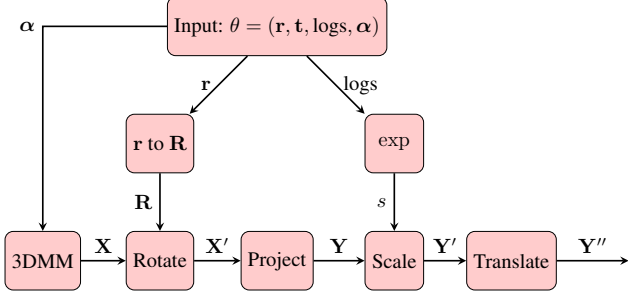


Figure 2. The grid generator network within a 3DMM-STN.

3D morphable model layer The 3D morphable model layer generates a shape $\mathbf{X} \in \mathbb{R}^{3 \times N}$ comprising N 3D vertices by taking a linear combination of D basis shapes (principal components) stored in the matrix $\mathbf{P} \in \mathbb{R}^{3N \times D}$ and the mean shape $\boldsymbol{\mu} \in \mathbb{R}^{3N}$ according to shape parameters $\boldsymbol{\alpha} \in \mathbb{R}^D$:

$$\mathbf{X}(\boldsymbol{\alpha})_{i,j} = \mathbf{x}(\boldsymbol{\alpha})_{3(j-1)+i}, \quad i \in [1, 3], j \in [1, N],$$

where

$$\mathbf{x}(\boldsymbol{\alpha}) = \mathbf{P}\boldsymbol{\alpha} + \boldsymbol{\mu}$$

and the derivatives are given by:

$$\frac{\partial \mathbf{x}}{\partial \boldsymbol{\alpha}} = \mathbf{P}, \quad \frac{\partial X_{i,j}}{\partial \alpha_k} = P_{3(j-1)+i,k}.$$

Note that such a linear model is exactly equivalent to a fully connected layer (and hence a special case of a convolutional layer) with fixed weights and biases. This is not at all surprising since a linear model is exactly what is implemented by a single layer linear decoder. In this interpretation, the shape parameters play the role of the input map, the principal components the role of weights and the mean shape the role of biases. This means that this layer can be implemented using an existing implementation of a convolution layer and also, following our later suggestion for future work, that the model could itself be made trainable simply by having non-zero learning rate for the convolution layer.

In our network, we use some of the principal components to represent shape variation due to identity and the remainder to represent deformation due to expression. We assume that expressions are additive and we can thus combine the two into a single linear model. Note that the shape parameters relating to identity may contain information that is useful for recognition, so these could be incorporated into a descriptor in a recognition network after the STN.

Axis-angle to rotation matrix layer This layer converts an axis-angle representation of a rotation, $\mathbf{r} \in \mathbb{R}^3$, into a rotation matrix:

$$\mathbf{R}(\mathbf{r}) = \cos \theta \mathbf{I} + \sin \theta [\bar{\mathbf{r}}]_{\times} + (1 - \cos \theta) \bar{\mathbf{r}}\bar{\mathbf{r}}^T,$$

where $\theta = \|\mathbf{r}\|$ and $\bar{\mathbf{r}} = \mathbf{r}/\|\mathbf{r}\|$ and

$$[\mathbf{a}]_{\times} = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix}$$

is the cross product matrix. The derivatives are given by [8]:

$$\frac{\partial \mathbf{R}}{\partial r_i} = \begin{cases} [\mathbf{e}_i]_{\times} & \text{if } \mathbf{r} = \mathbf{0} \\ \frac{r_i [\mathbf{r}]_{\times} + [\mathbf{r} \times (\mathbf{I} - \mathbf{R}(\mathbf{r})) \mathbf{e}_i]_{\times}}{\|\mathbf{r}\|^2} \mathbf{R} & \text{otherwise} \end{cases}$$

where \mathbf{e}_i is the i th vector of the standard basis in \mathbb{R}^3 .

3D rotation layer The rotation layer takes as input a rotation matrix \mathbf{R} and N 3D points $\mathbf{X} \in \mathbb{R}^{3 \times N}$ and applies the rotation:

$$\mathbf{X}'(\mathbf{R}, \mathbf{X}) = \mathbf{R}\mathbf{X}$$

$$\frac{\partial X'_{i,j}}{\partial R_{i,k}} = X_{k,j}, \quad \frac{\partial X'_{i,j}}{\partial X_{k,j}} = R_{i,k}, \quad i, k \in [1, 3], j \in [1, N].$$

Orthographic projection layer The orthographic projection layer takes as input a set of N 3D points $\mathbf{X}' \in \mathbb{R}^{3 \times N}$ and outputs N 2D points $\mathbf{Y} \in \mathbb{R}^{2 \times N}$ by applying an orthographic projection along the z axis:

$$\mathbf{Y}(\mathbf{X}') = \mathbf{P}\mathbf{X}', \quad \mathbf{P} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix},$$

$$\frac{\partial Y_{i,j}}{\partial X'_{i,j}} = 1, \quad i \in [1, 2], j \in [1, N].$$

Scaling The log scale estimated by the localiser is first transformed to scale by an exponentiation layer:

$$s(\text{logs}) = \exp(\text{logs}), \quad \frac{\partial s}{\partial \text{logs}} = \exp(\text{logs}).$$

Then, the 2D points $\mathbf{Y} \in \mathbb{R}^{2 \times N}$ are scaled:

$$\mathbf{Y}'(s, \mathbf{Y}) = s\mathbf{Y}, \quad \frac{\partial Y'_{i,j}}{\partial s} = Y_{i,j}, \quad \frac{\partial Y'_{i,j}}{\partial Y_{i,j}} = s$$

Translation Finally, the 2D sample points are generated by adding a 2D translation $\mathbf{t} \in \mathbb{R}^2$ to each of the scaled points:

$$\mathbf{Y}''(\mathbf{t}, \mathbf{Y}') = \mathbf{Y}' + \mathbf{1}_N \otimes \mathbf{t}, \quad \frac{\partial Y''_{i,j}}{\partial t_i} = 1, \quad \frac{\partial Y''_{i,j}}{\partial Y'_{i,j}} = 1,$$

where $\mathbf{1}_N$ is the row vector of length N containing ones and \otimes is the Kronecker product.




Figure 3. The output grid of our 3DMM-STN: a Tutte embedding of the mean shape of the Basel Face Model. On the left we show a visualisation using the mean texture (though note that our 3DMM-STN does not use a texture model). On the right we show the mean shape as a geometry image [9].

2.3. Sampling

In the original STN, the sampler component used bilinear sampling to sample values from the input image and transform them to an output grid. We make a number of modifications. First, the output grid is a texture space flattening of the 3DMM mesh. Second, the bilinear sampler layer will incorrectly sample parts of the face onto vertices that are self-occluded so we introduce additional layers that calculate which vertices are occluded and mask the sampled image appropriately.

Output grid The purpose of an STN is to transform an input image into a canonical, pose-normalised view. In the context of a 3D model, one could imagine a number of analogous ways that an input image could be normalised. For example, the output of the STN could be a rendering of the mean face shape in a frontal pose with the sampled texture on the mesh. Instead, we choose to output sampled textures in a 2D embedding obtained by flattening the mean shape of the 3DMM. This ensures that the output image is approximately area uniform with respect to the mean shape and also that the whole output image contains face information.

Specifically, we compute a Tutte embedding [7] using conformal Laplacian weights and with the mesh boundary mapped to a square. To ensure a symmetric embedding we map the symmetry line to the symmetry line of the square, flatten only one side of the mesh and obtain the flattening of the other half by reflection. We show a visualisation of our embedding using the mean texture in Figure 3. In order that the output warped image produces a regularly sampled image, we regularly re-sample (i.e. re-mesh) the 3DMM (mean and principal components) over a uniform grid of size $H' \times W'$ in this flattened space. This effectively makes our 3DMM a deformable geometry image [9]. The re-sampled 3DMM that we use in our STN therefore has $N = H'W'$ vertices and each vertex i has an associated UV coordinate (x_i^t, y_i^t) . The corresponding sample coordinate produced by the grid generator is given by $(x_i^s, y_i^s) = (Y_{1,i}'', Y_{2,i}'')$.

Bilinear sampling We use bilinear sampling, exactly as in the original STN such that the re-sampled image V_i^c at location (x_i^t, y_i^t) in colour channel c is given by:

$$V_i^c = \sum_{j=1}^H \sum_{k=1}^W I_{jk}^c \max(0, 1 - |x_i^s - k|) \max(0, 1 - |y_i^s - j|)$$

where I_{jk}^c is the value in the input image at pixel (j, k) in colour channel c . I has height H and width W . This bilinear sampling is differentiable (see [14] for derivatives) and so the loss can be backpropagated through the sampler and back into the grid generator.

Self-occlusions Since the 3DMM produces a 3D mesh, parts of the mesh may be self-occluded. The occluded vertices can be computed exactly using ray-tracing or z-buffering or they can be precomputed and stored in a lookup table. For efficiency, we approximate occlusion by only computing which vertices have backward facing normals. This approximation would be exact for any object that is globally convex. For objects with concavities, the approximation will underestimate the set of occluded vertices. Faces are typically concave around the eyes, the nose boundary and the mouth interior but we find that typically only around 5% of vertices are mislabelled and the accuracy is sufficient for our purposes.

This layer takes as input the rotation matrix \mathbf{R} and the shape parameters α and outputs a binary occlusion mask $\mathbf{M} \in \{0, 1\}^{H' \times W'}$. The occlusion function is binary and hence not differentiable at points where the visibility of a vertex changes, everywhere else the gradient is zero. Hence, we simply pass back zero gradients:

$$\frac{\partial \mathbf{M}}{\partial \alpha} = 0, \quad \frac{\partial \mathbf{M}}{\partial \mathbf{R}} = 0.$$

Note that this means that the network is not able to learn how changes in occlusion help to reduce the loss. Occlusions are applied in a forward pass but changes in occlusion do not backpropagate.

Masking layer The final layer in the sampler combines the sampled image and the visibility map via pixel-wise products:

$$W_i^c = V_i^c M_{x_i^t, y_i^t}, \quad \frac{\partial W_i^c}{\partial V_i^c} = M_{x_i^t, y_i^t}, \quad \frac{\partial W_i^c}{\partial M_{x_i^t, y_i^t}} = V_i^c.$$

3. Geometric losses for localiser training

An STN is usually inserted into a network as a pre-processor of input images and its output is then passed to a classification or regression CNN. Hence, the pose normalisation that is learnt by the STN is the one that produces optimal performance on the subsequent task. In the context of




Figure 4. Siamese multiview loss. An image and its horizontal reflection yield two sampled images. We penalise differences in these two images.

a 3D morphable face model, an obvious task would be face recognition. While this is certainly worth pursuing, we have observed that the optimal normalisation for recognition may not correspond to the correct model-image correspondence one would expect. For example, if context provided by hair and clothing helps with recognition, then the 3DMM-STN may learn to sample this.

Instead, we show that it is possible to train an STN to perform accurate localisation using only some simple geometric priors without even requiring identity labels for the images. We describe these geometric loss functions in the following sections.

3.1. Bilateral symmetry loss

Faces are approximately bilaterally symmetric. Ignoring the effects of illumination, this means that we expect sampled face textures to be approximately bilaterally symmetric. We can define a loss that measures asymmetry of the sampled texture over visible pixels:

$$\ell_{\text{sym}} = \sum_{i=1}^N \sum_{c=1}^3 M_{x_i^t, y_i^t} M_{x_{\text{sym}(i)}^t, y_{\text{sym}(i)}^t} (V_i^c - V_{\text{sym}(i)}^c)^2, \quad (2)$$

where $V_{\text{sym}(i)}^c$ is the value in the resampled image at location $(W' + 1 - x_i^s, y_i^s)$.




Figure 5. Landmark loss. Left: The diagram shows the implementation of the regression layer that computes the Euclidean distance between selected 2D points and ground truth positions. Right: Predicted positions are in red and landmark positions are in green.



Figure 6. Overview of the 3DMM-STN. From left to right: input image; rendering of estimated shape in estimated pose; sampled image; occlusion mask; final output of 3DMM-STN.

3.2. Siamese multi-view fitting loss

If we have multiple images of the same face in different poses (or equivalently from different viewpoints), then we expect that the sampled textures will be equal (again, neglecting the effects of illumination). If we had such multi-view images, this would allow us to perform Siamese training where a pair of images in different poses were sampled into images V_i^c and W_i^c with visibility masks M and N giving a loss:

$$\ell_{\text{multiview}} = \sum_{i=1}^N \sum_{c=1}^3 M_{x_i^t, y_i^t} N_{x_i^t, y_i^t} (V_i^c - W_i^c)^2. \quad (3)$$

Ideally, this loss would be used with a multiview face database or even a face recognition database where images of the same person in different in-the-wild conditions are present. We use an even simpler variant which does not require multiview images; again based on the bilateral symmetry assumption. A horizontal reflection of a face image approximates what that face would look like in a reflected pose. Hence, we perform Siamese training on an input image and its horizontal reflection. This is different to the bilateral symmetry loss and is effectively encouraging the localiser to behave symmetrically.

3.3. Landmark loss

As has been observed elsewhere [28], convergence of the training can be speeded up by introducing surrogate loss functions that provide supervision in the form of landmark locations. It is straightforward to add a landmark loss to our network. First, we define a selection layer that selects $L < N$ landmarks from the N 2D points outputted by the grid generator:

$$\mathbf{L} = \mathbf{Y}'' \mathbf{S} \quad (4)$$

where $\mathbf{S} \in \{0, 1\}^{N \times L}$ is a selection matrix with $\mathbf{S}^T \mathbf{S} = \mathbf{I}_L$. Given L landmark locations $\mathbf{l}_1, \dots, \mathbf{l}_L$ and associated detection confidence values c_1, \dots, c_L , we computed a weighted Euclidean loss:

$$\ell_{\text{landmark}} = \sum_{i=1}^L c_i \|\mathbf{L}_i - \mathbf{l}_i\|^2. \quad (5)$$

Landmarks that are not visible (i.e. were not hand-labelled or detected) are simply assigned zero confidence.




Figure 7. 3DMM-STN output for multiple images of the same person in different poses.

3.4. Statistical prior loss

The statistical shape model provides a prior. We scale the shape basis vectors such that the shape parameters follow a standard multivariate normal distribution: $\alpha \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_D)$. Hence, the statistical prior can be encoded by the following loss function:

$$\ell_{\text{prior}} = \|\alpha\|^2. \quad (6)$$

4. Experiments

For our statistical shape model, we use $D = 10$ dimensions of which five are the first five (identity) principal components from the Basel Face Model [19]. The other five are expression components which come from FaceWarehouse [5] using the correspondence to the Basel Model provided by [33]. We re-mesh the Basel Model over a uniform grid of size 224×224 . We trained our 3DMM-STN with the four loss functions described in Section 3 using the AFLW database [17]. This provides up to 21 landmarks per subject for over 25k in-the-wild images. This is a relatively small dataset for training a deep network so we perform ‘fine-tuning’ by setting the learning rate on the last layer of the localiser to four times that of the rest of the network. Figure 6 shows the pipeline of an image passing through a 3DMM-STN. A by-product of the trained 3DMM-STN is that it can also act as a 2D landmark localiser. After training, the localiser achieves an average landmarking error of 2.35 pixels on the part of AFLW used as validation set, over the 21 landmarks, showing that overall, the training converges well.

We begin by demonstrating that our 3DMM-STN learns to predict consistent correspondence between model and image. In Figure 7 we show 3DMM-STN output for multiple images of the same person. Note that the features are consistently mapped to the same location in the transformed output. In Figure 8 we go further by applying the 3DMM-STN to multiple images of the same person and then average the resulting transformed images. We show results for 10 subjects from the UMDFaces [1] dataset. The number of images for each subject is shown in parentheses. The averages have well-defined features despite being computed from images with large pose variation.




Figure 8. A set of mean flattened images per subject. Real images are obtained from UMDFaces dataset. The number of images that are used for averaging is stated next to subject’s name.

In Figure 9 we provide a qualitative comparison to [29]. This is the only previous work on 3DMM fitting using a CNN for which the trained network is made publicly available. In columns one and five, we show input images from UMDFaces [1]. In columns two and six, we show the reconstruction provided by [29]. While the reconstruction captures the rough appearance of the input face, it lacks the discriminating detail of the original image. This method regresses shape and texture directly but not illumination or pose. Hence, we cannot directly compare the model-image correspondence provided by this method. To overcome this, we use the landmark detector used by [29] during training and compute the optimal pose to align their reconstruction to these landmarks. We replace their cropped model by the original BFM shape model and sample the image. This allows us to create the flattened images in columns three and seven. The output of our proposed 3DMM-STN is shown in columns four and eight. We note that our approach less frequently samples background and yields a more consistent correspondence of the resampled faces. In the bottom row of the figure we show challenging examples where [29] did not produce any output because the landmark detector failed. Despite occlusions and large out of plane rotations, the 3DMM-STN still does a good job of producing a normalised output image.

5. Conclusions

In this paper we have shown how to use a 3D morphable model as a spatial transformer within a CNN. Our proposed architecture has a number of interesting properties. First, the network (specifically, the localiser part of the network) learns to fit a 3D morphable model to a single 2D image without needing labelled examples of fitted models. Since the problem of fitting a morphable model to an image is an unsolved problem (and therefore no existing algorithm could be assumed to provide reliable ground truth fits), this kind of unsupervised learning is desirable. Second, the morphable model itself is fixed in our current architecture. However, there is no reason that this could not

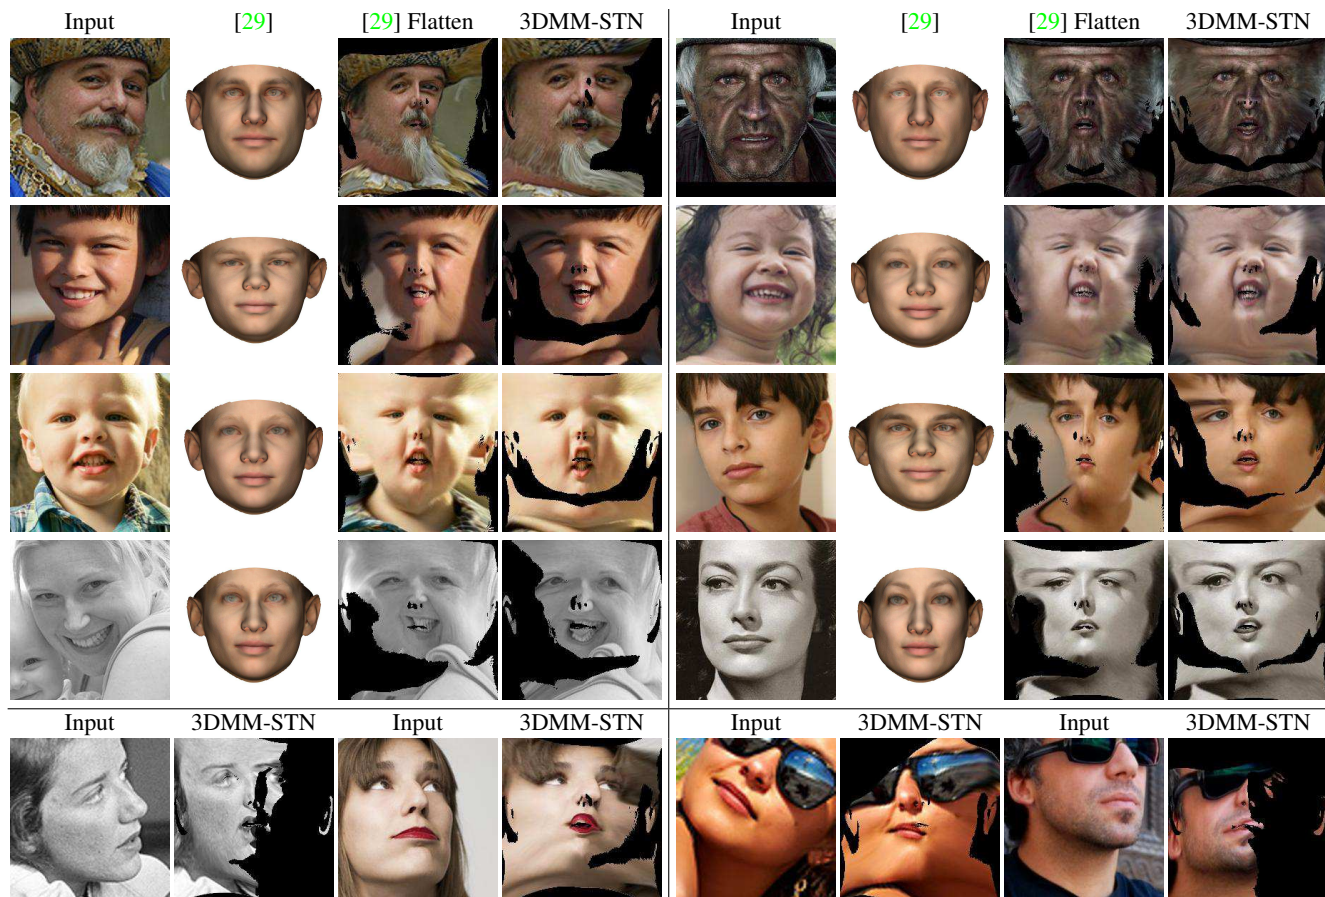


Figure 9. Qualitative comparison to [29]. The bottom row shows examples for which [29] failed to fit due to failure of the landmark detector.

also be learnt. In this way, it may be possible to learn a 3D deformable model for an object class simply from a collection of images that are labelled appropriately for the chosen proxy task.

There are many ways that this work can be extended. First, we would like to investigate training our 3DMM-STN in an end to end recognition network. We would hope that the normalisation means that a recognition network could be trained on less data and with less complexity than existing networks that must learn pose invariance implicitly. Second, the shape parameters estimated by the localiser may contain discriminative information and so these could be combined into subsequent descriptors for recognition. Third, we would like to further explore the multiview fitting loss. Using a multiview face database or video would provide a rich source of data for learning accurate localisation. Finally, the possibility of learning the shape model during training is exciting and we would like to explore other objects classes besides faces for which 3DMMs do not currently exist.

Acknowledgements

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research.

References

- [1] A. Bansal, A. Nanduri, C. D. Castillo, R. Ranjan, and R. Chellappa. UMDFaces: An annotated face dataset for training deep networks. *arXiv preprint arXiv:1611.01484v2*, 2016. 7
- [2] A. Bas and W. A. P. Smith. What does 2D geometric information really tell us about 3D face shape? *arXiv preprint arXiv:1708.06703*, 2017. 3
- [3] C. Bhagavatula, C. Zhu, K. Luu, and M. Savvides. Faster than real-time facial alignment: A 3d spatial transformer network approach in unconstrained poses. In *Proc. ICCV*, page to appear, 2017. 2
- [4] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *Proc. SIGGRAPH*, pages 187–194, 1999. 1, 2
- [5] C. Cao, Y. Weng, S. Zhou, Y. Tong, and K. Zhou. Face-warehouse: A 3D facial expression database for visual com-

- puting. *IEEE Trans. Vis. Comp. Gr.*, 20(3):413–425, 2014. 7
- [6] D. Chen, G. Hua, F. Wen, and J. Sun. Supervised transformer network for efficient face detection. In *Proc. ECCV*, pages 122–138, 2016. 2
- [7] M. S. Floater. Parametrization and smooth approximation of surface triangulations. *Comput. Aided Geom. Des.*, 14(3):231–250, 1997. 5
- [8] G. Gallego and A. Yezzi. A compact formula for the derivative of a 3-D rotation in exponential coordinates. *J. Math. Imaging Vis.*, 51(3):378–384, 2015. 4
- [9] X. Gu, S. J. Gortler, and H. Hoppe. Geometry images. *ACM Trans. Graphic.*, 21(3):355–361, 2002. 5
- [10] R. A. Güler, G. Trigeorgis, E. Antonakos, P. Snape, S. Zafeiriou, and I. Kokkinos. DenseReg: Fully convolutional dense shape regression in-the-wild. In *Proc. CVPR*, 2017. 2, 3
- [11] A. Handa, M. Bloesch, V. Pătrăucean, S. Stent, J. McCormac, and A. Davison. gynn: Neural network library for geometric computer vision. In *Proc. ECCV Workshop on Geometry Meets Deep Learning*, 2016. 2
- [12] J. F. Henriques and A. Vedaldi. Warped convolutions: Efficient invariance to spatial transformations. *CoRR*, abs/1609.04382, 2016. 2
- [13] A. S. Jackson, A. Bulat, V. Argyriou, and G. Tzimiropoulos. Large pose 3D face reconstruction from a single image via direct volumetric cnn regression. *arXiv preprint arXiv:1703.07834*, 2017. 2, 3
- [14] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks. In *Proc. NIPS*, pages 2017–2025, 2015. 1, 5
- [15] A. Jourabloo and X. Liu. Large-pose face alignment via CNN-based dense 3D model fitting. In *Proc. CVPR*, 2016. 2, 3
- [16] H. Kim, M. Zollhöfer, A. Tewari, J. Thies, C. Richardt, and C. Theobalt. Inversefacenet: Deep single-shot inverse face rendering from a single image. *arXiv preprint arXiv:1703.10956*, 2017. 2, 3
- [17] P. M. R. Martin Koestinger, Paul Wohlhart and H. Bischof. Annotated Facial Landmarks in the Wild: A Large-scale, Real-world Database for Facial Landmark Localization. In *Proc. First IEEE International Workshop on Benchmarking Facial Image Analysis Technologies*, 2011. 7
- [18] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proc. BMVC*, 2015. 3
- [19] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. A 3D face model for pose and illumination invariant face recognition. In S. Tubaro and J. Dugelay, editors, *Proc. AVSS*, pages 296–301, 2009. 1, 7
- [20] M. Pietraschke and V. Blanz. Automated 3D face reconstruction from multiple images using quality measures. In *Proc. CVPR*, pages 3418–3427, 2016. 2
- [21] E. Richardson, M. Sela, and R. Kimmel. 3D face reconstruction by learning from synthetic data. In *Proc. 3DV*, pages 460–469, 2016. 2, 3
- [22] E. Richardson, M. Sela, R. Or-El, and R. Kimmel. Learning detailed face reconstruction from a single image. In *Proc. CVPR*, 2017. 2
- [23] S. Romdhani and T. Vetter. Estimating 3d shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *Proc. CVPR*, volume 2, pages 986–993, 2005. 2
- [24] S. Schönborn, B. Egger, A. Morel-Forster, and T. Vetter. Markov chain monte carlo for automated face image analysis. *International Journal of Computer Vision*, 123(2):160–183, 2017. 2
- [25] M. Sela, E. Richardson, and R. Kimmel. Unrestricted facial geometry reconstruction using image-to-image translation. *arXiv preprint arXiv:1703.10131*, 2017. 2, 3
- [26] W. A. P. Smith. The perspective face shape ambiguity. In *Perspectives in Shape Analysis*, pages 299–319. Springer, 2016. 3
- [27] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proc. CVPR*, pages 1701–1708, 2014. 1
- [28] A. Tewari, M. Zollhöfer, H. Kim, P. Garrido, F. Bernard, P. Pérez, and C. Theobalt. MoFA: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. *arXiv preprint arXiv:1703.10580*, 2017. 2, 3, 6
- [29] A. T. Trăn, T. Hassner, I. Masi, and G. Medioni. Regressing robust and discriminative 3D morphable models with a very deep neural network. In *Proc. CVPR*, 2017. 2, 3, 7, 8
- [30] A. Vedaldi and K. Lenc. Matconvnet: Convolutional neural networks for MATLAB. In X. Zhou, A. F. Smeaton, Q. Tian, D. C. A. Bulterman, H. T. Shen, K. Mayer-Patel, and S. Yan, editors, *Proceedings of the 23rd Annual ACM Conference on Multimedia Conference, MM '15, Brisbane, Australia, October 26 - 30, 2015*, pages 689–692. ACM, 2015. 1
- [31] X. Yan, J. Yang, E. Yumer, Y. Guo, and H. Lee. Perspective transformer nets: Learning single-view 3D object reconstruction without 3D supervision. In *Advances in Neural Information Processing Systems*, pages 1696–1704, 2016. 2
- [32] X. Yu, F. Zhou, and M. Chandraker. Deep deformation network for object landmark localization. In *Proc. ECCV*, pages 52–70, 2016. 2
- [33] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li. Face alignment across large poses: A 3D solution. In *Proc. CVPR*, pages 146–155, 2016. 7