

Faster Than Real-time Facial Alignment: A 3D Spatial Transformer Network Approach in Unconstrained Poses

Chandrasekhar Bhagavatula, Chenchen Zhu, Khoa Luu, and Marios Savvides
Carnegie Mellon University
Pittsburgh, PA

cbhagava@andrew.cmu.edu, zcckernel@cmu.edu, kluu@andrew.cmu.edu, msavvid@ri.cmu.edu

Abstract

Facial alignment involves finding a set of landmark points on an image with a known semantic meaning. However, this semantic meaning of landmark points is often lost in 2D approaches where landmarks are either moved to visible boundaries or ignored as the pose of the face changes. In order to extract consistent alignment points across large poses, the 3D structure of the face must be considered in the alignment step. However, extracting a 3D structure from a single 2D image usually requires alignment in the first place. We present our novel approach to simultaneously extract the 3D shape of the face and the semantically consistent 2D alignment through a 3D Spatial Transformer Network (3DSTN) to model both the camera projection matrix and the warping parameters of a 3D model. By utilizing a generic 3D model and a Thin Plate Spline (TPS) warping function, we are able to generate subject specific 3D shapes without the need for a large 3D shape basis. In addition, our proposed network can be trained in an end-to-end framework on entirely synthetic data from the 300W-LP dataset. Unlike other 3D methods, our approach only requires one pass through the network resulting in a faster than real-time alignment. Evaluations of our model on the Annotated Facial Landmarks in the Wild (AFLW) and AFLW2000-3D datasets show our method achieves state-of-the-art performance over other 3D approaches to alignment.

1. Introduction

Robust face recognition and analysis are contingent upon accurate localization of facial features. When modeling faces, the landmark points of interest consist of points that lie along the shape boundaries of facial features, e.g. eyes, lips, mouth, etc. When dealing with face images collected in the wild conditions, facial occlusion of landmarks becomes a common problem for off-angle faces. Predicting the occlusion state of each landmarking points is one of the

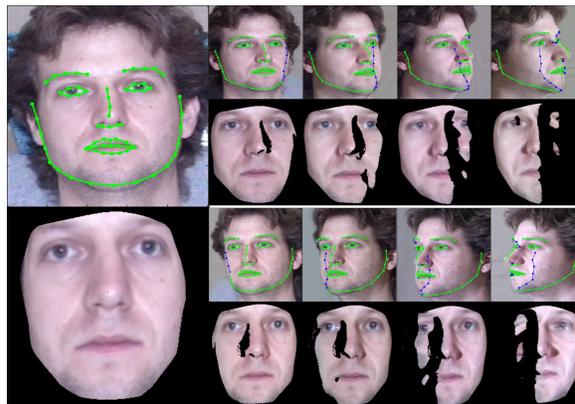


Figure 1. A subject from the CMU Multi-PIE dataset [18, 19] landmarked and frontalized by our method at various poses. Landmarks found by our model are overlaid in green if they are determined to be a visible landmark and blue if self-occluded. The non-visible regions of the face are determined by the estimated camera center and the estimated 3D shape. **Best viewed in color.**

challenges due to variations of objects in faces, e.g. beards and mustaches, sunglasses and other noisy objects. Additionally, face images of interest nowadays usually contain off-angle poses, illumination variations, low resolutions, and partial occlusions.

Many complex factors could affect the appearance of a face image in real-world scenarios and providing tolerance to these factors is the main challenge for researchers. Among these factors, pose is often the most important factor to be dealt with. It is known that as facial pose deviates from a frontal view, most face recognition systems have difficulty in performing robustly. In order to handle a wide range of pose changes, it becomes necessary to utilize 3D structural information of faces. However, many of the existing 3D face modeling schemes [1, 4, 42] have many drawbacks, such as computation time and complexity. Though these can be mitigated by using depth sensors [23] or by tracking results from frame to frame in video [39], this can cause difficulty when they have to be applied in real-world large

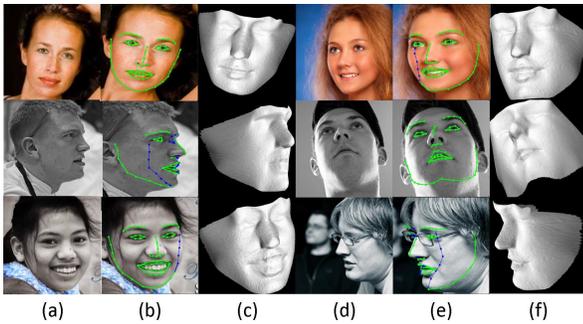


Figure 2. (a & d): Images in the wild from the AFLW dataset. (b & d): 3D landmarks (green: visible, blue: occluded) estimated from input image. (e & f): 3D model generated from input image. **Best viewed in color.**

scale unconstrained face recognition scenarios where video and depth information is not available. The 3D generic elastic model (3D-GEM) approach was proposed as an efficient and reliable 3D modeling method from a single 2D image. Heo *et al.* [22, 35] claim that the depth information of a face is not extremely discriminative when factoring out the 2D spatial location of facial features. In our method, we follow this idea and observe that fairly accurate 3D models can be generated by using a simple mean shape deformed to the input image at a relatively low computational cost compared to other approaches.

1.1. Our Contributions in this Work

(1) We take the approach of using a simple mean shape and using a parametric, non-linear warping of that shape through alignment on the image to be able to model any unseen example. A key flaw in many approaches that rely on a 3D Morphable Model (3DMM) is that it needs enough examples of the data to be able to model unseen samples. However, in the case of 3D faces, most datasets are very small.

(2) Our approach is efficiently implemented in an end-to-end deep learning framework allowing for the alignment and 3D modeling tasks to be codependent. This ensures that alignment points are semantically consistent across changing poses of the object which also allows for more consistent 3D model generation and frontalization on images in the wild as shown in Figs. 1 and 2.

(3) Our method only requires a single pass through the network allowing us to achieve faster than real-time processing of images with state-of-the-art performance over other 2D and 3D approaches to alignment.

2. Related Work

There have been numerous studies related to face alignment since the first work of Active Shape Models (ASM) [14] in 1995. A comprehensive literature review in face

alignment is beyond the scope of this work. In this paper, we mainly focus on recent Convolutional Neural Network (CNN) approaches to solve the face alignment problem. Especially those methods aimed at using 3D approaches to achieve robust alignment results.

2.1. Face Alignment Methods

While Principal Component Analysis (PCA) and its variants [14, 13, 15] were successfully used to model the facial shapes and appearances, there have since been many advances in facial alignment. Landmark locations can be directly predicted by a regression from a learned feature space [9, 16, 44]. Xiong *et al.* [45] presented the Global Supervised Descent Method (GSDM) method to solve the problem of 2D face alignment. The objective function in GSDM is divided into multiple regions of similar gradient directions. It then constructs a separate cascaded shape regressor for each region. Yu *et al.* [46] incorporated 3D pose landmarking models with group sparsity to indicate the best landmarks. These kind of methods shows an increase of performance on landmark localization. However, these methods all rely on hand-crafted features. Recently, CNN-based methods have achieved good results in facial alignment [51, 48]. 3DDFA [51] fits a dense 3D face model to the image via CNN and DDN [48] proposes a novel cascaded framework incorporating geometric constraints for localizing landmarks in faces and other non-rigid objects. Recently, shape regression has been used in numerous facial landmarking methods [41, 36, 43].

There are several recent works studying the human head rotations [12, 53], nonlinear statistical models ([17]) and 3D shape models [8, 20]. Nonlinear statistical model approaches are impractical in real-time applications. View-based methods employ a separate model for each viewpoint mode. Traditionally, the modes are specified as part of the algorithm design, and problems can arise at midpoints between models.

2.2. CNNs for 3D Object Modeling

While estimating a 3D model from images is not a new problem, the challenging task of modeling objects from a single image has always posed a challenge. This is, of course, due to the ambiguous nature of images where depth information is removed. With the recent success of deep learning and especially CNNs in extracting salient information from images, there have been many explorations into how to best use CNNs for modeling objects in 3 dimensions. Many of these approaches are aimed creating a depth estimation for natural images [33, 2, 37, 32, 31]. While the results on uncontrolled images are impressive, the fact that these models are very general means they tend to suffer when applied to specific objects, such as faces. In fact, many times, the depth estimate for faces in the scene tend

to be fairly flat. By limiting the scope of the method, the resulting estimated 3D model can be made much more accurate. Hassner et al. [21] use a 3D model of the face to be able to frontalize faces in unseen images with the end goal of improving face recognition by limiting the variations the matcher has to learn. However, this approach requires landmarks on the input face in the same fashion as other methods [22, 35, 21, 34, 52].

A 2D approach to landmarking inevitably suffers from the problem of visibility and self-occlusion. As Zhu et al. [52] show, the problem of landmark marching, where landmarks tend to move to the visible boundary, can cause issues when estimating 3D models from purely 2D alignment. However, this problem can be alleviated by using a 3D model of the face in the alignment step itself as done in [27, 51]. Both of these methods make use of an underlying 3D Morphable Model (3DMM) and try to fit the model to the input image in order to find the 2D landmarks. This of course requires a basis to use and the Basel Face Model (BFM) [24] is a very popular model to use. However, the BFM is only created from a set of 100 male and 100 female scans. As any basis can only recreate combinations of the underlying samples, this can severely limit the capability of these models to fit outlier faces or expressions not seen before. Although there has been recent efforts to generate more accurate 3DMMs [6], neither the data nor the model is available to researchers in the field of biometrics. Therefore, we propose to use a smooth warping function, Thin Plate Splines (TPS) [5], to warp mean shapes to fit the input image and generate new 3D shapes. In this fashion, any new face can be modeled, even if its shape cannot be reconstructed by the BFM.

3. 3D Spatial Transformer Networks

In order to model how a face truly changes from viewpoint to viewpoint, it is necessary to have both the true 3D model of the subject in the image and the properties of the camera used to capture the image, usually in the form of the camera projection matrix. However, knowledge of the true 3D model and the camera projection matrix are almost always not available. Jaderberg *et al.* [25], in their work on Spatial Transformer Networks, use a deep network to estimate the parameters of either an affine transformation or a 2D Thin Plate Spline (TPS) transformation. These parameters are then used to generate a new sampling grid which can then be used to generate the transformed image.

We approach finding the unknown camera projection matrix parameters and the parameters needed to generate the 3D model of the head in a similar fashion. Both the camera projection parameters and the warping parameters, a TPS warp in this case, can be estimated from deep features generated from the image using any architecture. The TPS parameters can be used to warp a model of the face

to match what the network estimates the true 3D shape is and the camera projection parameters can be used to texture the 3D coordinates from the 2D image. Additionally, the pose of the face can be determined from the camera parameters allowing for a visibility map to be generated for the 3D model. This allows us to only texture vertexes that are visible in the image as opposed to vertexes that are occluded by the face itself. The architecture of our model is shown in Figure 3. Sections 3.1, 3.2, and 3.3 detail how to create differentiable modules to utilize the camera projection and TPS parameters that are estimated by the deep network to warp and project a 3D model to a 2D image plane for texture sampling.

3.1. Camera Projection Transformers

In order to be able to perform end-to-end training of a network designed to model 3D transformations of the face, a differentiable module that performs a camera projection must be created. This will be part of the grid generator portion of the Spatial Transformer. Modeling how a 3D point will map to the camera coordinates is expressed by the well known camera projection equation

$$\mathbf{p}_c \approx \mathbf{M}\mathbf{p}_w \quad (1)$$

where \mathbf{p}_c is the homogeneous 2D point in the camera coordinate system, \mathbf{p}_w is the homogeneous 3D point in the world coordinate system, and \mathbf{M} is the 3x4 camera projection matrix. This relationship is only defined up to scale due to the ambiguity of scale present in projective geometry, hence the \approx instead of a hard equality. The camera projection matrix has only 11 degrees of freedom since it is only defined up to scale as well. Therefore, this module takes in the 11 parameters estimated by a previous layer as the input in the form of a length 11 vector, \mathbf{a} . In order to perform backpropagation on the new grid generator, the derivative of the generated grid with respect to \mathbf{a} must be computed.

Since Eqn. 1 is only defined up to scale, the final output of this module will have to divide out the scale factor. By first rewriting the camera projection matrix as

$$\mathbf{M} = \begin{bmatrix} a_1 & a_2 & a_3 & a_4 \\ a_5 & a_6 & a_7 & a_8 \\ a_9 & a_{10} & a_{11} & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{m}_1^T \\ \mathbf{m}_2^T \\ \mathbf{m}_3^T \end{bmatrix} \quad (2)$$

where a_i is the i^{th} element of \mathbf{a} , the final output of the camera projection module can be written as

$$\mathbf{O} = \begin{bmatrix} x_c \\ y_c \end{bmatrix} = \begin{bmatrix} \frac{\mathbf{m}_1^T \mathbf{p}_w}{\mathbf{m}_3^T \mathbf{p}_w} \\ \frac{\mathbf{m}_2^T \mathbf{p}_w}{\mathbf{m}_3^T \mathbf{p}_w} \end{bmatrix} \quad (3)$$

The gradient with respect to each of the rows of \mathbf{M} can be

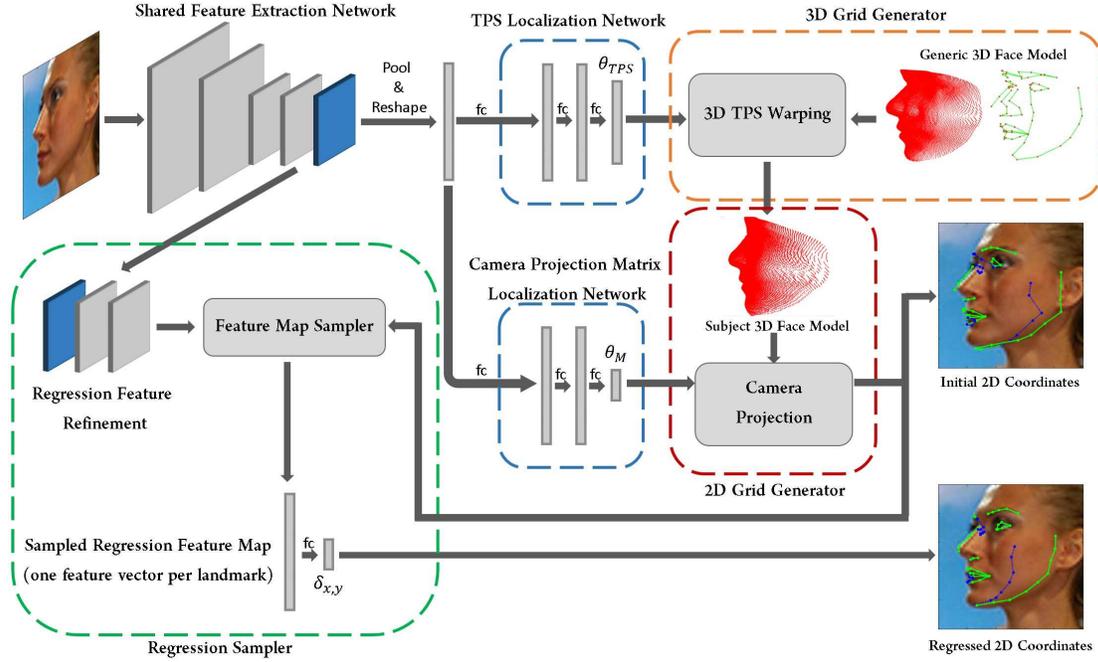


Figure 3. Network design of the 3D TPS Spatial Transformer for facial alignment. Because a 3D model and an estimate of the camera position are found in the output of the network, visibility of landmarks can also be determined. Visible landmarks are shown in green while non-visible landmarks are shown in blue.

shown to be

$$\begin{aligned} \frac{\delta \mathbf{O}}{\delta \mathbf{m}_1^T} &= \begin{bmatrix} \mathbf{p}_w^T \\ \mathbf{m}_3^T \mathbf{p}_w \\ \mathbf{0} \end{bmatrix} & \frac{\delta \mathbf{O}}{\delta \mathbf{m}_2^T} &= \begin{bmatrix} \mathbf{0} \\ \mathbf{p}_w^T \\ \mathbf{m}_3^T \mathbf{p}_w \end{bmatrix} \\ \frac{\delta \mathbf{O}}{\delta \mathbf{m}_3^T} &= \begin{bmatrix} -\mathbf{p}_w^T (\mathbf{m}_1^T \mathbf{p}_w) \\ (\mathbf{m}_3^T \mathbf{p}_w)^2 \\ -\mathbf{p}_w^T (\mathbf{m}_2^T \mathbf{p}_w) \\ (\mathbf{m}_3^T \mathbf{p}_w)^2 \end{bmatrix} \end{aligned} \quad (4)$$

Using the chain rule, the gradient of the loss of the network with respect to the input can be found as

$$\frac{\delta L}{\delta \mathbf{a}} = \begin{bmatrix} \left(\frac{\delta L}{\delta \mathbf{O}} \frac{\delta \mathbf{O}}{\delta \mathbf{m}_1^T} \right)^T \\ \left(\frac{\delta L}{\delta \mathbf{O}} \frac{\delta \mathbf{O}}{\delta \mathbf{m}_2^T} \right)^T \\ \left(\frac{\delta L}{\delta \mathbf{O}} \frac{\delta \mathbf{O}}{\delta \mathbf{m}_3^T} \right)^T \end{bmatrix} \quad (5)$$

Since \mathbf{M} is only defined up to scale, the last element of \mathbf{M} can be defined to be a constant which means that only the first 11 elements of this gradient are used to actually perform the backpropagation on \mathbf{a} . Since \mathbf{M} relates many pairs of 2D and 3D points, the gradient is computed for every pair and added together to give the final gradient that is used for updating \mathbf{a} .

3.2. 3D Thin Plate Spline Transformers

When modeling the 3D structure of a face, a generic model cannot represent the variety of shapes that might be

seen in an image. Therefore, some method of warping a model must be used to allow the method to handle unseen shapes. Thin Plate Spline (TPS) warping has been used by many applications to great effect [5, 11]. TPS warps have the very desirable features of providing a closed form of a smooth, parameterized warping given a set of control points and desired destination points. Jaderberg *et al.* [25] showed how 2D TPS Spatial Transformers could lead to good normalization of nonlinearly transformed input images. Applying a TPS to a 3D set of points follows a very similar process. As in [25], the TPS parameters would be estimated from a deep network of some sort and passed as input to a 3D grid generator module.

A 3D TPS function is of the form

$$f_{\Delta_x}(x, y, z) = \begin{bmatrix} b_{1x} \\ b_{2x} \\ b_{3x} \\ b_{4x} \end{bmatrix}^T \begin{bmatrix} 1 \\ x \\ y \\ z \end{bmatrix} + \sum_{j=1}^n w_{jx} U(|\mathbf{c}_j - (x, y)|) \quad (6)$$

where b_{1x} , b_{2x} , b_{3x} , b_{4x} , and w_{jx} are the parameters of the function, \mathbf{c}_j is the j^{th} control point used in determining the function parameters, and $U(r) = r^2 \log r$. This function is normally learned by setting up a system of linear equations using the known control points, \mathbf{c}_j and the corresponding points in the warped 3D object. The function finds the change in a single coordinate, the change in the x -coordinate in the case of Eqn. 6. Similarly, one such

function is created for each dimension, i.e. $f_{\Delta_x}(x, y, z)$, $f_{\Delta_y}(x, y, z)$, and $f_{\Delta_z}(x, y, z)$. The 3D TPS module would then take in the parameters for all three of these functions as input and output the newly transformed points on a 3D structure as

$$\mathbf{O} = \begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \begin{bmatrix} f_{\Delta_x}(x, y, z) \\ f_{\Delta_y}(x, y, z) \\ f_{\Delta_z}(x, y, z) \end{bmatrix} + \begin{bmatrix} x \\ y \\ z \end{bmatrix} \quad (7)$$

This means that the 3D TPS module must have all of the 3D vertices of the generic model and the control points on the generic model as fixed parameters specified from the start. This will allow the module to warp the specified model by the warps specified by the TPS parameters.

As in 3.1, the gradient of the loss with respect to the input parameters must be computed in order to perform backpropagation on this module. As usual, the chain rule can be used to find this by computing the gradient of the output with respect to the input parameters. Since each 3D vertex in the generic model will give one 3D vertex as an output, it is easier to compute the gradient on one of these points, $\mathbf{p}_i = (x_i, y_i, z_i)$, first. This can be shown to be

$$\frac{\delta \mathbf{O}}{\delta \theta_{\Delta_x}} = \begin{bmatrix} 1 & 0 & 0 \\ x_i & 0 & 0 \\ y_i & 0 & 0 \\ z_i & 0 & 0 \\ U(|\mathbf{c}_1 - (x_i, y_i, z_i)|) & 0 & 0 \\ \vdots & \vdots & \vdots \\ U(|\mathbf{c}_n - (x_i, y_i, z_i)|) & 0 & 0 \end{bmatrix}^T \quad (8)$$

where θ_{Δ_x} are the parameters of f_{Δ_x} . Similarly, the gradients for θ_{Δ_y} and θ_{Δ_z} are the same with only the non-zero values in either the second or third row respectively. The final gradient of the loss with respect to the parameters can be computed as

$$\frac{\delta L}{\delta \theta_{\Delta_x}} = \frac{\delta L}{\delta \mathbf{O}} \frac{\delta \mathbf{O}}{\delta \theta_{\Delta_x}} \quad (9)$$

Since this is only for a single point, once again the gradient can be computed for every point and added for each set of parameters to get the final gradient for each set of parameters that can be used to update previous layers of the network.

3.3. Warped Camera Projection Transformers

In order to make use of the TPS warped 3D points in the camera projection module of the transformer network, the module must take in as input the warped coordinates. This means that such a module would also have to do backpropagation on the 3D coordinates as well as the camera projection parameters. Since 3.1 already specified how to

compute the gradient of the loss with respect to the camera projection parameters, all that is left to do is compute the gradient of the loss with respect to the 3D coordinates in this module. Taking the derivative of the output in Eqn. 3 with respect to the 3D point, \mathbf{p}_w results in

$$\frac{\delta \mathbf{O}}{\delta \mathbf{p}_w} = \begin{bmatrix} \frac{\mathbf{m}_1^T}{\mathbf{m}_3^T \mathbf{p}_w} - \frac{\mathbf{m}_1^T \mathbf{p}_w}{(\mathbf{m}_3^T \mathbf{p}_w)^2} \mathbf{m}_3^T \\ \frac{\mathbf{m}_2^T}{\mathbf{m}_3^T \mathbf{p}_w} - \frac{\mathbf{m}_2^T \mathbf{p}_w}{(\mathbf{m}_3^T \mathbf{p}_w)^2} \mathbf{m}_3^T \end{bmatrix} \quad (10)$$

However, since \mathbf{p}_w is in homogeneous coordinates and only the gradient with respect to the x , y , and z coordinates are needed, the actual gradient becomes

$$\frac{\delta \mathbf{O}}{\delta \mathbf{p}'_w} = \begin{bmatrix} \frac{\mathbf{m}'_1{}^T}{\mathbf{m}'_3{}^T \mathbf{p}'_w} - \frac{\mathbf{m}'_1{}^T \mathbf{p}'_w}{(\mathbf{m}'_3{}^T \mathbf{p}'_w)^2} \mathbf{m}'_3{}^T \\ \frac{\mathbf{m}'_2{}^T}{\mathbf{m}'_3{}^T \mathbf{p}'_w} - \frac{\mathbf{m}'_2{}^T \mathbf{p}'_w}{(\mathbf{m}'_3{}^T \mathbf{p}'_w)^2} \mathbf{m}'_3{}^T \end{bmatrix} \quad (11)$$

where

$$\mathbf{p}'_w = \begin{bmatrix} x_w \\ y_w \\ z_w \end{bmatrix} \quad \mathbf{m}'_i = \begin{bmatrix} m_{i_1} \\ m_{i_3} \\ m_{i_3} \end{bmatrix} \quad (12)$$

and m_{i_j} is the j^{th} element of \mathbf{m}_i . This gradient is computed for every 3D point independently and used in the chain rule to compute

$$\frac{\delta L}{\delta \mathbf{p}_w} = \frac{\delta L}{\delta \mathbf{O}} \frac{\delta \mathbf{O}}{\delta \mathbf{p}_w} \quad (13)$$

which can then be used to perform backpropagation on each \mathbf{p}_w .

3.4. 2D Landmark Regression

In order to further improve the landmark accuracy, we extend our network with a landmark refinement stage. This stage treats the projected 2D coordinates from the previous stage as initial points and estimates the offsets for each point. To extract the feature vector for each point, a 3×3 convolution layer is attached on top of the last convolution layer in the base model, followed by a 1×1 convolution layer for more nonlinearity, resulting in a feature map with D channels. Then each initial point is projected onto this feature map and its D -dimensional feature vector is extracted along the channel direction. Notice that the initial points are often not aligned with the grids on the feature map. Therefore, their feature vectors are sampled with bilinear interpolation.

Given the feature vector for each landmark, it goes through a fully-connected (FC) layer to output the offsets, i.e. δ_x and δ_y . Then the offsets are added to the coordinates of the initial location. For each landmark we use an independent FC layer. We don't share the FC layer for all landmarks because each landmark should have a unique behavior of offsets. For example, the center of the eye may move left after regression whereas the corner of the eye may

move right. Also, sometimes two initial landmarks may be projected to the same location due to a certain pose. We want them to move to different locations even when they have the same feature vector.

3.5. 3D Model Regression From 2D Landmarks

Once the 2D regression is performed, the mapping between the 3D model and the 2D landmarks is broken. While this is not necessarily a problem in the case of sparse facial alignment, if a denser scheme is needed, the entire model would have to be retrained. In order to avoid this, we create a new 3D model that does map to these 2D landmarks by finding a new set of 3D coordinates that project to the new 2D landmarks and warping the 3D model to fit these new points. To find the new 3D coordinates, we need to backproject rays through each of the 2D landmarks through 3D space using the camera projection matrix we have estimated. The equation for the ray of points associated with a given homogeneous 2D point, \mathbf{p}_{2D}^i , is defined as

$$\mathbf{p}_{3D}^{i'} = \begin{bmatrix} \mathbf{A}^{-1}\mathbf{b} \\ 1 \end{bmatrix} + \lambda \begin{bmatrix} \mathbf{A}^{-1}\mathbf{p}_{2D}^i \\ 0 \end{bmatrix} \quad (14)$$

where \mathbf{A} and \mathbf{b} are the first three and the last column of the estimated camera projection matrix respectively.

These rays represent all possible points in 3D that could project to the determined locations in the image. We then find the closest point, $\mathbf{p}_{3D}^{i'}$, on the ray to the original 3D coordinate, \mathbf{p}_{3D}^i , to use as the new 3D point as shown in Fig. 4. These new correspondences are used to perform a TPS warping of the model. After this warping, the landmark points on the model will project to exactly the regressed 2D landmarks, recovering the mapping between the 3D model and the 2D image. This new model can then be projected onto the image to generate a much more accurate texturing of the 3D model. This same style of warping can be used to move the 3D coordinates anywhere we choose. This means neutralizing out expressions, especially smiles, is very easy to do by using the texture from the regressed 3D shape. While the non-smiling shape will not be as accurate due to the fact that a non-smiling image was not seen, it still gives convincing qualitative results, as seen in Fig. 5, which indicate it may be a worthwhile avenue of exploration for future work, especially in face recognition.

4. Experiments

4.1. Datasets

300W-LP: The 300W-LP [51] dataset contains 122,450 synthetically generated views of faces from the AFW [54], LFPW [3], HELEN [30], and IBUG [38] datasets. These images not only contain rotated faces but also attempt to move the background in a convincing fashion, making it a

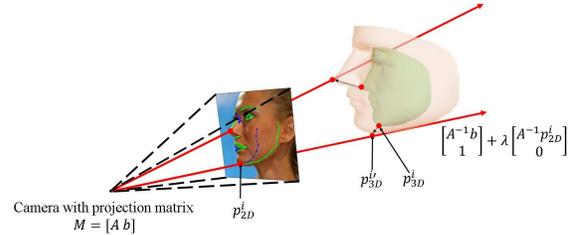


Figure 4. Backprojection or rays through image landmarks. The closest points are found for each ray-landmark pair to use as new 3D coordinates for the face model. The original model (green) is warped to fit the new landmarks with a 3D TPS warp resulting in a new face model (red).

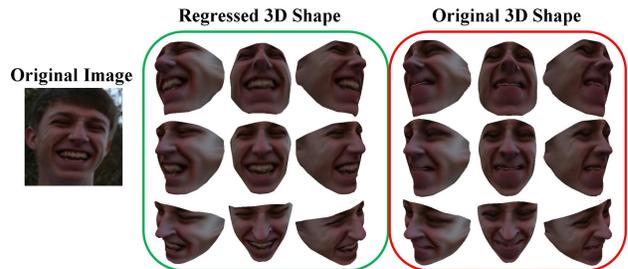


Figure 5. 3D renderings of input face with a smiling expression. The resulting regressed 3D model (green box) maintains the smile and is very similar to the input image while the same texture applied to the original shape (red box) suffers a small degradation in shape but allows for a non-smiling rendering of the input image.

very useful dataset for training 3D approaches to work on real world images.

AFLW: The Annotated Facial Landmarks in the Wild (AFLW) dataset [28] is a relatively large dataset for evaluating facial alignment on wild images. It contains approximately 25,000 faces annotated with 21 landmarks with visibility labels. The dataset provides pose estimates so results are grouped into three different pose ranges, $[0^\circ, 30^\circ]$, $(30^\circ, 60^\circ]$, and $(60^\circ, 90^\circ]$. Due to the inconsistency in the bounding boxes in the AFLW dataset, we adopt the use of a face detector first to normalize the scale of the faces. The Multiple Scale Faster Region-based CNN approach [49] has shown good results and at a fast speed. We use the recent extension to this work, the Contextual Multi-Scale Region-based CNN (CMS-RCNN) approach [50] to perform the face detection in any experiment where face detection is needed. The CMS-RCNN approach detects 98.8% (13,865), 95.9% (5,710), and 86.5% (3,830) of the faces in the $[0^\circ, 30^\circ]$, $(30^\circ, 60^\circ]$, and $(60^\circ, 90^\circ]$ pose ranges respectively.

AFLW2000-3D: Zhu et al. [51] accurately pointed out how merely evaluating an alignment scheme on the visible landmarks in a dataset can result in artificially low errors. Therefore, a true evaluation of any 3D alignment method

must also evaluate alignment on the non-visible landmarks as well. The AFLW2000-3D dataset contains the first 2000 images of the AFLW dataset but with all 68 points defined by the scheme in the CMU MPIE dataset [18, 19]. These points were found by aligning the Basel Face Model to the images. While this is a synthetic dataset, meaning the true location of the non-visible landmarks is not known, it is the best one can do when dealing with real images. As these images are from the AFLW dataset, they are also grouped into the same pose ranges.

4.2. Implementation Details

Our network is implemented in the Caffe [26] framework. A new layer is created consisting of the 3D TPS transformation module, the camera projection module and the bilinear sampler module. All modules are differentiable so that the whole network can be trained end-to-end.

We adopt two architectures, AlexNet [29] and VGG-16 [40], as the pre-trained models for our shared feature extraction networks in Fig. 3, i.e. we use the convolution layers from the pre-trained models to initialize ours. Since these networks already extract informative low-level features and we do not want to lose this information, we freeze some of the earlier convolution layers and finetune the rest. For the AlexNet architecture, we freeze the first layer while for the VGG-16 architecture, the first 4 layers are frozen.

The 2D landmark regression is implemented by attaching additional layers on top of the last convolution layer. With N landmarks to regress, we need N FC layers to compute the offsets for each individual landmark. While it's possible to setup N individual FC layers, here we implement this by adding one Scaling layer followed by a Reduction layer and Bias layer. During training only the new layers are updated and all previous layers are frozen.

4.3. Training on 300W-LP

When training our model, we train on the AFW, HELEN, and LFPW subsets of the 300W-LP dataset and use the IBUG portion as a validation set. All sets are normalized using the bounding boxes from the CMS-RCNN detector by reshaping the detected faces to 250 x 250 pixels. For the AlexNet architecture, we train for 100,000 iterations with a batch size of 50. The initial learning rate is set to 0.001 and drops by a factor of 2 after 50,000 iterations. When training the landmark regression, the initial learning rate is 0.01 and drops by a factor of 10 every 40,000 iterations. For the VGG-16 architecture, we train for 200,000 iterations with a batch size of 25. The initial learning rate is set to 0.001 and drops by a factor of 2 after 100,000 iterations. When training the landmark regression, the initial learning rate is 0.01 and drops by a factor of 10 every 70,000 iterations. The momentum for all experiments is set to 0.9. Euclidean loss is applied to 3D vertexes, 2D projected landmarks and 2D

Table 1. Alignment accuracy for both the AlexNet (AN) and VGG-16 (VGG) models. (LR: landmark regression)

	AFLW Dataset (21 pts)				
	[0, 30]	(30, 60]	(60, 90]	mean	std
AN	4.88	5.55	7.10	5.84	1.14
AN+LR	4.00	4.48	5.89	4.79	0.98
VGG	4.15	4.64	5.96	4.92	0.94
VGG+LR	3.46	3.78	4.77	4.00	0.69

regressed landmarks.

4.4. Ablation Experiments

To investigate the effect of each component in our network, we conduct two ablation studies. All the models in these experiments are trained on the same 300W_LP dataset and tested on the detected images in AFLW. We first test the effect of the different pre-trained models. We fine-tune our network from the AlexNet and VGG-16 models pre-trained on the ImageNet dataset and evaluate the landmark accuracy before the regression step. The VGG-16 model outperforms the AlexNet model in all three pose ranges on the AFLW detected set as shown in Table 1. This seems to indicate that a good base model is important for the parameter estimation portion of the network. Second, we evaluate the effect of landmark regression stage. We compare the errors between the regressed and projected landmarks. Table 1 shows that the landmark regression step greatly helps to improve the accuracy.

4.5. Comparison Experiments

AFLW: Since the CMS-RCNN approach may only detect the easier to landmark faces, we use the provided bounding box anytime the face is not detected by the detector. Due to the inconsistency between the two bounding box schemes, faces are not always normalized properly. However, we feel this is the only way to get a fair comparison to other methods without artificially making the dataset easier by only evaluating on detected faces. We compare against baseline methods used by [51] on the same dataset, namely Cascaded Deformable Shape Models (CDM) [47], Robust Cascaded Pose Regression (RCPR) [7], Explicit Shape Regression (ESR) [10], SDM [44] and 3DDFA [51]. All methods except for CDM were retrained on the 300W-LP dataset. The Normalized Mean Error (NME) is computed by averaging the error of the visible landmarks and normalizing it by the square root of the bounding box size ($h \times w$) provided in the dataset. Table 2 clearly shows that our model using the VGG-16 architecture has achieved better accuracy in all pose ranges, especially the $(60^\circ, 90^\circ]$ category, and has achieved a smaller standard deviation in the error. This means that not only are the landmarks more accurate, they are more consistent than the other methods.

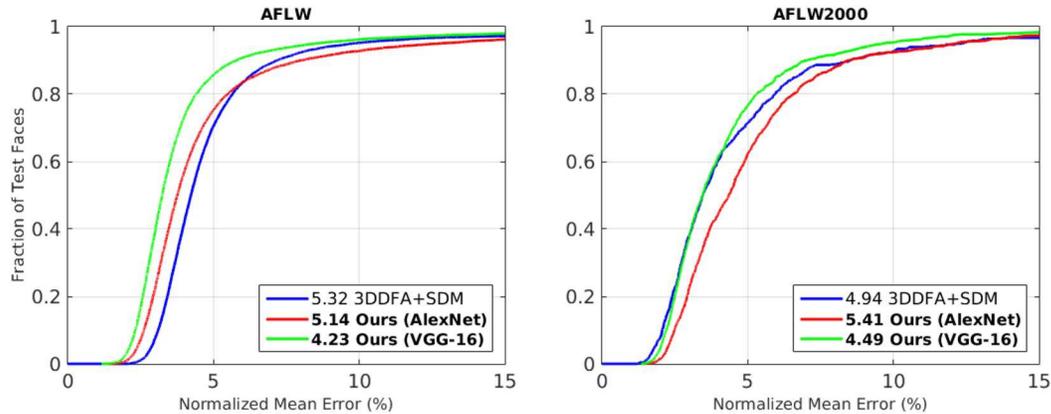


Figure 6. CED curves for both the AlexNet (red) and VGG-16 (green) architectures on both the AFLW (left) and AFLW2000-3D (right) dataset. To balance the distributions, we randomly sample 13,209 faces from AFLW and 915 faces from AFLW2000-3D, split evenly among the 3 categories, and compute the CED curve. This is done 10 times and the average of the resulting CED curves are reported. The mean NME% for each architecture from Table 2 is also reported in the legend.

Table 2. The NME(%) of face alignment results on AFLW and AFLW2000-3D. The best two numbers in each category are shown in bold.

Method	AFLW Dataset (21 pts)					AFLW 2000-3D Dataset (68 pts)				
	[0, 30]	(30, 60]	(60, 90]	mean	std	[0, 30]	(30, 60]	(60, 90]	mean	std
CDM	8.15	13.02	16.17	12.44	4.04	-	-	-	-	-
RCPR	5.43	6.58	11.53	7.85	3.24	4.26	5.96	13.18	7.80	4.74
ESR	5.66	7.12	11.94	8.24	3.29	4.60	6.70	12.67	7.99	4.19
SDM	4.75	5.55	9.34	6.55	2.45	3.67	4.94	9.76	6.12	3.21
3DDFA	5.00	5.06	6.74	5.60	0.99	3.78	4.54	7.93	5.42	2.21
3DDFA+SDM	4.75	4.83	6.38	5.32	0.92	3.43	4.24	7.17	4.94	1.97
Ours (AlexNet)	4.11	4.69	6.61	5.14	1.31	3.71	5.33	7.19	5.41	1.74
Ours (VGG-16)	3.55	3.92	5.21	4.23	0.87	3.15	4.33	5.98	4.49	1.42

AFLW2000-3D: The baseline methods were evaluated using the bounding box of the 68 landmarks so we retrained our models using the same bounding box on the training data. Generating these is trivial due to the 3D models. The NME is computed using the bounding box size. Here we see that though 3DDFA+SDM performs well, the VGG-16 architecture of our model still performs best in both the $[0^\circ, 30^\circ]$ and $(60^\circ, 90^\circ]$ ranges. While the VGG-16 model is only second best in the $(30^\circ, 60^\circ]$ range by a small amount, the improvement in $(60^\circ, 90^\circ]$ means that, once again, our method generates more accurate and more consistent landmarks, even in a 3D sense. Cumulative Error Distribution (CED) curves are reported for both architectures on both datasets in Fig. 6.

4.6. Running Speed

In order to evaluate the speed of our method, we evaluate the models on a random subset of 1200 faces from the AFLW subset split evenly into the $[0^\circ, 30^\circ]$, $(30^\circ, 60^\circ]$, and $(60^\circ, 90^\circ]$ pose ranges. The images are processed one at a time to avoid any benefit from batch processing. The models are evaluated on a 3.40 GHz Intel Core i7-6700 CPU and

an NVIDIA GeForce GTX TITAN X GPU. Our AlexNet trained model takes a total of 7.064 seconds to landmark the 1200 faces for an average of 0.0059 seconds per image or approximately 170 faces per second. The deeper and more accurate VGG-16 model landmarks the 1200 faces in 22.765 seconds for an average of 0.0190 seconds or approximately 52 faces per second. In comparison, the 3DDFA approach [51] takes 75.72 ms (3 iterations at 25.24 ms per iteration as specified in [51]) with 2/3 of the time being used to process data on the CPU.

5. Conclusions

In this paper we propose a method using 3D Spatial Transformer Networks with TPS warping to generate both a 3D model of the face and accurate 2D landmarks across large pose variation. The limited data used in the generation of a 3DMM can mean that unseen face shapes cannot be modeled. By using a TPS warp, any potential face can be modeled through a regression of 2D landmarks, of which there is much more data available. We have shown how this approach leads to more accurate and consistent landmarks over other 2D and 3D methods.

References

- [1] J. J. Atick, P. A. Griffin, and A. N. Redlich. Statistical approach to shape from shading: Reconstruction of three-dimensional face surfaces from single two-dimensional images. *Neural Comput.*, 8(6):1321–1340, Aug. 1996.
- [2] A. Bansal, B. Russell, and A. Gupta. Marr revisited: 2d-3d alignment via surface normal prediction. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5965–5974, June 2016.
- [3] P. N. Belhumeur, D. W. Jacobs, D. J. Kriegman, and N. Kumar. Localizing parts of faces using a consensus of exemplars. In *CVPR 2011*, pages 545–552, June 2011.
- [4] V. Blanz, S. Romdhani, and T. Vetter. Face identification across different poses and illuminations with a 3d morphable model. In *Proceedings of Fifth IEEE International Conference on Automatic Face Gesture Recognition*, pages 192–197, May 2002.
- [5] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(6):567–585, June 1989.
- [6] J. Booth, A. Roussos, S. Zafeiriou, A. Ponniah, and D. Dunaway. A 3d morphable model learnt from 10,000 faces. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5543–5552, June 2016.
- [7] X. P. Burgos-Artizzu, P. Perona, and P. Dollár. Robust face landmark estimation under occlusion. In *Proceedings of the 2013 IEEE International Conference on Computer Vision, ICCV '13*, pages 1513–1520, Washington, DC, USA, 2013. IEEE Computer Society.
- [8] C. Cao, Y. Weng, S. Lin, and K. Zhou. 3d shape regression for real-time facial animation. In *ACM Transactions on Graphics*, 2013.
- [9] X. Cao, Y. Wei, F. Wen, and J. Sun. Face alignment by explicit shape regression. *International Journal of Computer Vision*, 107(2):177–190, 2014.
- [10] X. Cao, Y. Wei, F. Wen, and J. Sun. Face alignment by explicit shape regression. *Int. J. Comput. Vision*, 107(2):177–190, Apr. 2014.
- [11] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Underst.*, 89(2-3):114–141, Feb. 2003.
- [12] T. Cootes, K. Walker, , and C. Taylor. View-based active appearance models. In *IEEE Intl Conf. and Workshops on Automatic Face and Gesture Recognition*, 2015.
- [13] T. F. Cootes, G. J. Edwards, C. J. Taylor, et al. Active appearance models. *IEEE Transactions on pattern analysis and machine intelligence*, 23(6):681–685, 2001.
- [14] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995.
- [15] D. Cristinacce and T. Cootes. Automatic feature localisation with constrained local models. *Pattern Recognition*, 41(10):3054–3067, 2008.
- [16] M. Dantone, J. Gall, G. Fanelli, and L. Van Gool. Real-time facial feature detection using conditional regression forests. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2578–2585. IEEE, 2012.
- [17] C. N. Duong, K. Luu, K. G. Quach, and T. D. Bui. Beyond principal components: Deep boltzmann machines for face modeling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [18] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition*. IEEE Computer Society, September 2008.
- [19] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image Vision Comput.*, 28(5):807–813, May 2010.
- [20] L. Gu and T. Kanade. 3d alignment of face in a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [21] T. Hassner, S. Harel, E. Paz, and R. Enbar. Effective face frontalization in unconstrained images. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4295–4304, June 2015.
- [22] J. Heo and M. Savvides. Gender and ethnicity specific generic elastic models from a single 2d image for novel 2d pose face synthesis and recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(12):2341–2350, Dec. 2012.
- [23] P. L. Hsieh, C. Ma, J. Yu, and H. Li. Unconstrained real-time facial performance capture. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1675–1683, June 2015.
- [24] IEEE. *A 3D Face Model for Pose and Illumination Invariant Face Recognition*, Genova, Italy, 2009.
- [25] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks. *CoRR*, abs/1506.02025, 2015.
- [26] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014.
- [27] A. Jourabloo and X. Liu. Pose-invariant 3d face alignment. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 3694–3702, Dec 2015.
- [28] M. Koestinger, P. Wohlhart, P. M. Roth, and H. Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *First IEEE International Workshop on Benchmarking Facial Image Analysis Technologies*, 2011.
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [30] V. Le, J. Brandt, Z. Lin, L. Bourdev, and T. S. Huang. Interactive facial feature localization. In *Proceedings of the 12th European Conference on Computer Vision - Volume Part III, ECCV'12*, pages 679–692, Berlin, Heidelberg, 2012. Springer-Verlag.
- [31] B. Li, C. Shen, Y. Dai, A. van den Hengel, and M. He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1119–1127, June 2015.

- [32] F. Liu, C. Shen, and G. Lin. Deep convolutional neural fields for depth estimation from a single image. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5162–5170, June 2015.
- [33] F. Liu, C. Shen, G. Lin, and I. Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(10):2024–2039, Oct 2016.
- [34] I. Masi, S. Rawls, G. Medioni, and P. Natarajan. Pose-aware face recognition in the wild. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4838–4846, June 2016.
- [35] U. Prabhu, J. Heo, and M. Savvides. Unconstrained pose-invariant face recognition using 3d generic elastic models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(10):1952–1961, Oct 2011.
- [36] S. Ren, X. Cao, Y. Wei, and J. Sun. Face alignment at 3000 fps via regressing local binary features. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [37] A. Roy and S. Todorovic. Monocular depth estimation using neural regression forest. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [38] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 397–403, Dec 2013.
- [39] S. Saito, T. Li, and H. Li. *Real-Time Facial Segmentation and Performance Capture from RGB Input*, pages 244–261. Springer International Publishing, Cham, 2016.
- [40] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.
- [41] G. Tzimiropoulos. Project-out cascaded regression with an application to face alignment. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [42] S.-F. Wang and S.-H. Lai. *Efficient 3D Face Reconstruction from a Single 2D Image by Combining Statistical and Geometrical Information*, pages 427–436. 2006.
- [43] P. P. X. P. Burgos-Artizzu and P. Dollr. Robust face landmark estimation under occlusion. In *IEEE International Conference on Computer Vision*, June 2013.
- [44] X. Xiong and F. De la Torre. Supervised descent method and its applications to face alignment. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 532–539, 2013.
- [45] X. Xiong and F. De la Torre. Global supervised descent method. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2664–2673, 2015.
- [46] X. Yu, J. Huang, S. Zhang, W. Yan, and D. Metaxas. Pose-free facial landmark fitting via optimized part mixtures and cascaded deformable shape model. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1944–1951, 2013.
- [47] X. Yu, J. Huang, S. Zhang, W. Yan, and D. N. Metaxas. Pose-free facial landmark fitting via optimized part mixtures and cascaded deformable shape model. In *2013 IEEE International Conference on Computer Vision*, pages 1944–1951, Dec 2013.
- [48] X. Yu, F. Zhou, and M. Chandraker. Deep deformation network for object landmark localization. *arXiv preprint arXiv:1605.01014*, 2016.
- [49] Y. Zheng, C. Zhu, K. Luu, C. Bhagavatula, T. H. N. Le, and M. Savvides. Towards a deep learning framework for unconstrained face detection. In *2016 IEEE 8th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–8, Sept 2016.
- [50] C. Zhu, Y. Zheng, K. Luu, and M. Savvides. CMS-RCNN: contextual multi-scale region-based CNN for unconstrained face detection. *CoRR*, abs/1606.05413, 2016.
- [51] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li. Face alignment across large poses: A 3d solution. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 146–155, June 2016.
- [52] X. Zhu, Z. Lei, J. Yan, D. Yi, and S. Z. Li. High-fidelity pose and expression normalization for face recognition in the wild. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 787–796, June 2015.
- [53] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2012.
- [54] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2879–2886, June 2012.