

DSFNet: Dual Space Fusion Network for Occlusion-Robust 3D Dense Face Alignment

Heyuan Li¹, Bo Wang², Yu Cheng¹, Mohan Kankanhalli¹, Robby T. Tan¹

¹National University of Singapore ²CtrsVision

liheyuan@u.nus.edu, hawk.rsrch@gmail.com, e0321276@u.nus.edu,
 mohan@comp.nus.edu.sg, robbly.tan@nus.edu.sg

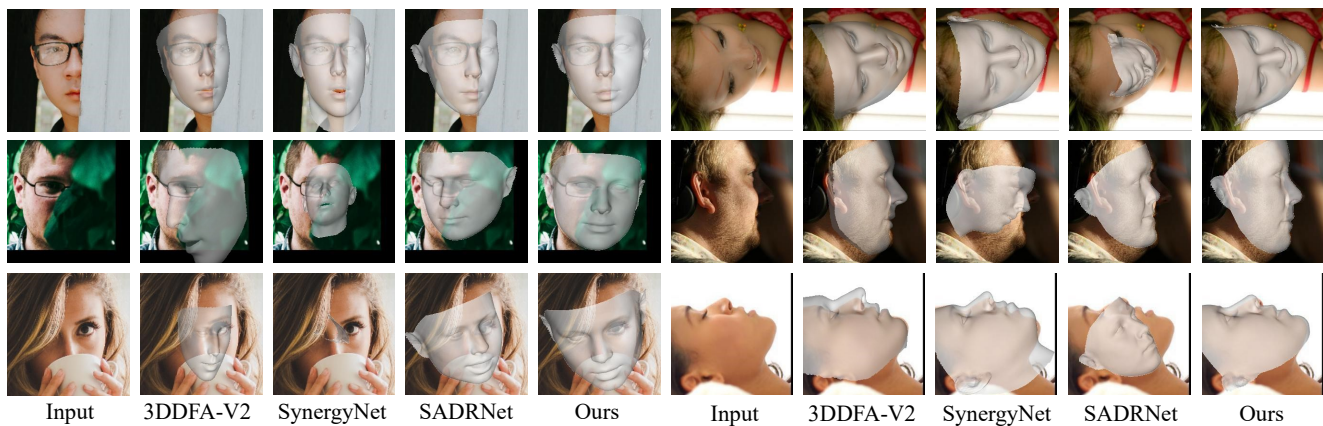


Figure 1. The qualitative results compared to state-of-the-art methods on severely occluded or large view angle cases from datasets [38,47]. For each case, from left to right are input, results of 3DDFA-V2 [12], SynergyNet [41], SADRNet [26], and our method.

Abstract

Sensitivity to severe occlusion and large view angles limits the usage scenarios of the existing monocular 3D dense face alignment methods. The state-of-the-art 3DMM-based method, directly regresses the model’s coefficients, underutilizing the low-level 2D spatial and semantic information, which can actually offer cues for face shape and orientation. In this work, we demonstrate how modeling 3D facial geometry in image and model space jointly can solve the occlusion and view angle problems. Instead of predicting the whole face directly, we regress image space features in the visible facial region by dense prediction first. Subsequently, we predict our model’s coefficients based on the regressed feature of the visible regions, leveraging the prior knowledge of whole face geometry from the morphable models to complete the invisible regions. We further propose a fusion network that combines the advantages of both the image and model space predictions to achieve high robustness and accuracy in unconstrained scenarios. Thanks to the proposed fusion module, our method is robust not only to occlusion and large pitch and roll view angles, which is the bene-

fit of our image space approach, but also to noise and large yaw angles, which is the benefit of our model space method. Comprehensive evaluations demonstrate the superior performance of our method compared with the state-of-the-art methods. On the 3D dense face alignment task, we achieve 3.80% NME on the AFLW2000-3D dataset, which outperforms the state-of-the-art method by 5.5%. Code is available at <https://github.com/lhyfst/DSFNet>.

1. Introduction

3D dense face alignment is an important problem with many applications, e.g. video conferencing, AR/VR/metaverse, games, facial analysis, etc. Many methods have been proposed [8–12, 16, 21, 26, 27, 29, 33, 39, 41, 47, 49]. However, these methods are sensitive to severe occlusion and large view angles [19, 26, 31, 32], limiting their applicability of 3D dense face alignment on wild images where occlusion and view angles often occur.

3D dense face alignment from a single image is an ill-posed problem, mainly because of the depth ambiguity. The

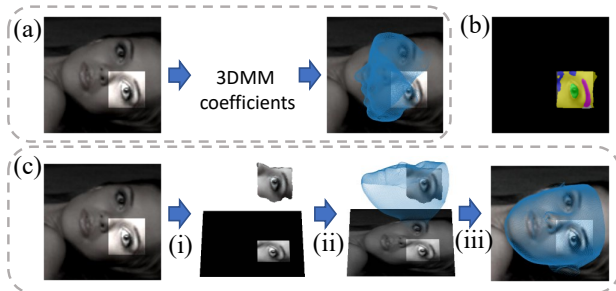


Figure 2. In this case, only one eye is visible. (a) The 3DMM-based method fails. (b) Face parsing algorithm [43] still works. (c) Our method first (i) predicts reliable geometry in visible region by dense prediction, then (ii) completes the whole face by facial geometry prior, producing a reasonable result. (iii) Viewed in image view.

existing methods [8, 10, 27, 29, 33] use a contractive CNN to predict the coefficients of 3DMM [4] directly. However, contractive CNNs are essentially ill-suited for this task [18] due to some reasons, including: mixed depth and 2D spatial information, the loss of low-level 2D spatial information as a result of invariance attribute of CNNs, and mixed facial region and occluder in the process of the contraction.

Severe occlusion and large view angles pose problems due to the complexity of the many-to-one mapping from 2D image to 3D shape. In contrast, low-level vision features are less variant according to geometry transform. Therefore, dense prediction is essentially more robust to the above problem in the visible region, because dense prediction relies more on local information, where an example is shown in Fig. 2 (b). Even if most of the face is masked out and only the left eye is visible, the face parsing algorithm is still able to deduce a reasonable parsing result.

Based on this observation, we decentralize the instance-level 3DMM coefficients regression (i.e., whole-face level) to pixel-level dense prediction in image space to improve the robustness against occlusion and large view angles, by proposing a 3D facial geometry’s 2D image space representation. To complete the invisible region due to extra- or self-occlusion, a novel post-process algorithm is proposed to convert the dense prediction for the visible face region into 3D facial geometry that includes the whole face area. Fig. 2 (c) shows that our image space prediction recovers reasonable results only seeing one eye, while the SOTA method fails to produce a reasonable result.

We further compare the robustness and accuracy between the image space prediction with the model space prediction that directly regresses 3DMM’s coefficients, and discover that there is a complementary relationship between these two spaces. Thus, we propose a dual space fusion network (DSFNet) that predicts using the image and model spaces using a two-branch architecture. With the fusion

module, our DSFNet effectively combines the advantages of both spaces. In summary, the main contributions of this paper are:

- We propose a novel 3D facial geometry’s 2D image space representation, followed by a novel post-processing algorithm. It achieves robust 3D dense face alignment to occlusion and large view angles.
- We introduce a fusion network, which combines the advantages of both the image and model space predictions to achieve high robustness and accuracy in unconstrained scenarios.
- On the 3D dense face alignment task, we achieve **3.80%** NME on AFLW2000-3D dataset, which outperforms the state-of-the-art method by **5.5%**.

2. Related Works

3D Dense Face Alignment 3D facial geometry is a finite deformable object, which has specific distributions according to gender, age, and ethnicity. Since 3DMM compactly represents such prior knowledge by Principal Component Analysis (PCA), 3DMM model-based methods have been dominating 3D facial geometry modeling since deep learning revived in the past decade. A batch of works [11, 12, 41, 47, 49] directly supervise the deep learning model by ground truth 3DMM coefficients in annotated datasets. Another batch of works [8–10, 27, 29, 39] supervise the deep learning model by photometric information in unlabelled images by taking advantage of differentiable rendering in a self-supervised way.

On the contrary, there are some model-free methods [11, 14, 26, 48] that do not involve 3DMM, but directly predict 3D coordinates of vertices of the face mesh. Therefore, they theoretically have more geometry representation flexibility. PRN [11] and SADRNet [26] predict a position map in UV space. VRN [14] outputs its prediction in 3D volumetric space, which leads to large parameter volume and heavy computing burden. Model-free methods always are able to give precise predictions in visible region, but cannot tackle invisible region well, because of lack of the prior knowledge of 3D facial geometry distribution.

Occlusion-Aware 3D Dense Face Alignment Recently, more and more works [19, 26, 34] realize the importance of robustness to occlusion and large view angles which are common in unconstrained images. Some methods enhanced their robustness by focusing on the visible facial region. SADRNet [26] proposes an attention mechanism that lets the model only focus on the visible facial region. [19] proposed a precise face skin segmentation network which masks out non-facial regions, so that it supervises only the model-based encoder by photometric signal in the visible

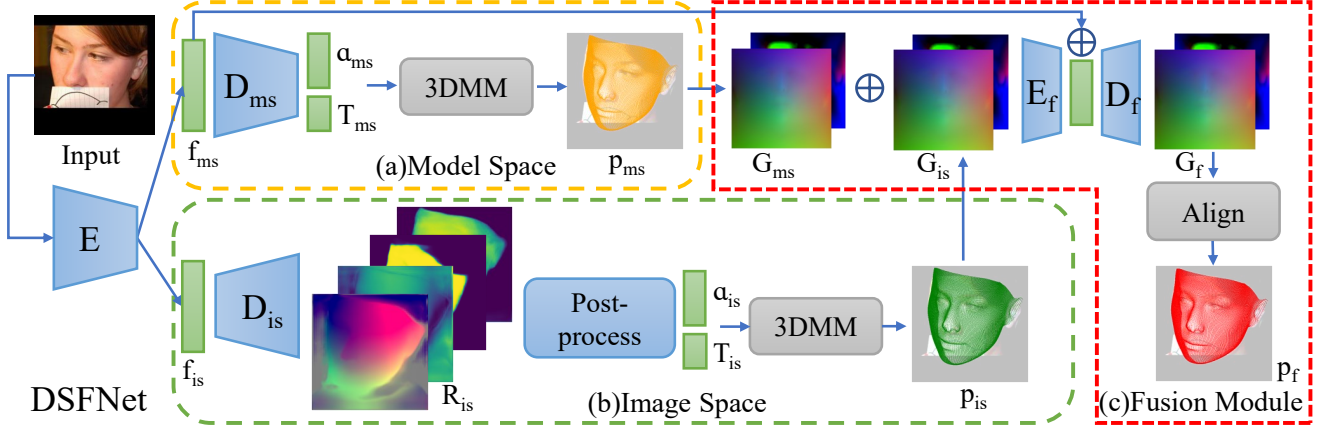


Figure 3. The framework of our proposed Dual Space Fusion Network (DSFNet).

face skin region. Another series of works achieve higher robustness to occlusion by data augmentation. [31] proposed a Self-Supervised Robustifying Guidance framework which takes advantage of the consistency between randomly occluded images and their corresponding unoccluded images. [32] further proposed a Multi-Occlusion Per Identity framework which takes advantage of the consistency between a batch of images from the same identity but occluded with different random occlusion patterns.

Another batch of works [2, 6, 17, 18, 28, 40, 44, 47] predict visible facial geometry in image space first, and the invisible region can be recovered from the predicted visible region using prior knowledge in 3DMM, which is always conducted by an optimization-based algorithm. DenseReg [2] proposed a correspondence map that contains the UV coordinates of vertices on the face model rasterized to the image plane. [44] predicts the correspondences between the input image and a rendered image of the 3DMM mean face geometry in the positive view by predicting 2D flow between them. [6] represents the 3D facial geometry by a projected normalized coordinated code (PNCC) [47] and a 3D offset image. [18] realizes that image-to-image CNNs are more suitable to predict the correspondence map because they reserve more low-level spatial information. [28] represents the face by a depth map and a correspondence map, where each pixel represents its x , y , and z coordinates of the corresponding point on a normalized canonical face. Some recent works don't use correspondence maps.

Wood et al. [40] use hundreds of dense 2D landmarks as the image space representation. [17] predicts the image space 2D-3D correspondence which is used to align predicted shape from canonical view to image view by Perspective-n-Point algorithm. However, these works have two limitations. First, some works have inferior accuracy, which is caused by poor 2D representations. Second, most works use optimization-based algorithms to post-process

the 2D representation, but optimization-based algorithms bring more hyper-parameters which are tricky to tune and slow down the whole pipeline, because optimization-based algorithms update parameters iteratively. By contrast, our method uses a well-designed 3D facial geometry's 2D image space representation to reserve enough geometry information for accuracy and uses a lightweight PointNet-based module as the post-process algorithm instead of an optimization method, therefore avoiding burdened computing.

3. Proposed Method

We present DSFNet, our dual space fusion network, which aims at robust and accurate 3D facial geometry modeling by taking advantages from both model space prediction and image space prediction. Fig. 3 shows our DSFNet's architecture. Our method takes a single face image as input, and uses HRNet [30] as an encoder E to extract model space and image space features. The two space features are processed by model space branch and image space branch separately. Finally, the predictions from two branches are fed into a fusion module for the final prediction.

Facial Geometry Prior We adopt the 3D Morphable Model as the geometry prior to our 3D face modeling. 3DMM represents face shape $S \in \mathbb{R}_{3n}$ with n vertices by identity and expression PCA bases and coefficients:

$$S = \bar{S} + B_{id}\alpha_{id} + B_{exp}\alpha_{exp}, \quad (1)$$

where \bar{S} is the mean face geometry; B_{id} and B_{exp} are the PCA bases of identities and expressions. Variables α_{id} and α_{exp} are the corresponding coefficient vectors. We adopt the widely-used 2009 Basel Face Model [22], where $\alpha_{shape} \in \mathbb{R}_{199}$ and $\alpha_{exp} \in \mathbb{R}_{29}$. Orthographic projection is used to align the 3D face shape to the image view: $G = fRS + t$, where $G \in \mathbb{R}_{3n}$ is aligning with the image view. Variables f , R and t are scale factor, 3D rotation

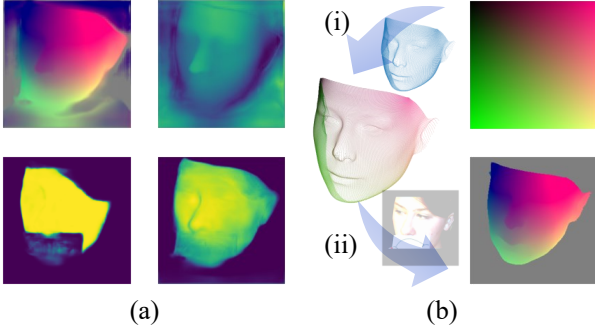


Figure 4. (a) An example of the four components of the image space representation in the inference stage. From left to right, from top to down, the four maps are correspondence map, depth map, segmentation map, and confidence map. (b) The illustration of the construction of a correspondence map. (i) Register UV locations to their corresponding 3D points. (ii) Rasterize the UV locations to the image plane to obtain the correspondence map.

matrix and 3D translation, respectively. We combine these three transformation parameters together as a pose vector $T \in \mathbb{R}_{12}$.

3.1. Model Space Prediction

As shown in Fig. 3 (a), we use a contractive CNN module as a model space decoder D_{ms} to regress model space prediction. The network takes f_{ms} as input, which is the smallest feature map (8×8) from last stage of encoder E , and output the 3DMM coefficients α_{ms} and pose vector T_{ms} :

$$\alpha_{\text{ms}}, T_{\text{ms}} = D_{\text{ms}}(f_{\text{ms}}). \quad (2)$$

With ground truth notation $*$ hereafter, the model space loss is shown as follows:

$$L_{\text{ms}} = w_{\alpha} \|\alpha_{\text{ms}} - \alpha^*\|^2 + w_T \|T_{\text{ms}} - T^*\|^2, \quad (3)$$

where w denotes the weight of the corresponding item.

3.2. Image Space Prediction

As depicted in Fig. 3 (b), an image space decoder D_{is} takes f_{is} as input, which is the feature map (64×64) concatenated from 4 scales of the output from the last stage of encoder E . D_{is} predicts the image space representation R_{is} in the input size via upsampling and convolutions layers. Then, a PointNet-based post-process algorithm Post converts the predicted representation into 3DMM coefficients α_{is} and pose vector T_{is} :

$$R_{\text{is}} = D_{\text{is}}(f_{\text{is}}), \quad (4)$$

$$\alpha_{\text{is}}, T_{\text{is}} = \text{Post}(R_{\text{is}}). \quad (5)$$

3.2.1 Image Space Representation

Our image space representation R_{is} consists of four 2D maps: correspondence map Cor, depth map Dep, segmentation map Seg, and confidence map Cf, shown in Fig. 4 (a). We represent a pixel X in the input image by $(x, y) \in [1, H] \times [1, W]$, where H and W are the height and width of the input image I .

Correspondence Map $\text{Cor} \in \mathbb{R}_{2 \times H \times W}$ represents pixels' UV coordinates on the 3DMM. $\text{Cor}(x, y)$ bridges the correspondence between X 's location (x, y) in image and its UV location (u, v) on our face model. As shown in Fig. 4 (b), our correspondence map is constructed in the following steps. First, similar to representing a mesh's texture by a 2D texture map, we unwrap the 3DMM to a 2D UV space $M \in \mathbb{R}_{2 \times h \times w}$, where h and w are the height and width of the 2D map. Now, every vertex has a UV location $(u, v) \in [1, h] \times [1, w]$. The UV location is then normalized to $(u, v) \in [-1, 1] \times [-1, 1]$. Every vertex's UV location is registered to its 3D location on G , and the correspondence map is obtained by rasterizing the UV locations to the image plane via G .

Depth Map $\text{Dep} \in \mathbb{R}_{H \times W}$ represents pixels' relative depth, which contains the geometry information in the non-occluded region. Using relative depth is feasible due to orthogonal projection, which helps the model by narrowing the prediction space. The nose tip's depth is used as zero point to compute the relative depth.

Segmentation Map $\text{Seg} \in \mathbb{R}_{H \times W}$ indicates the visible facial regions. Its values range from 0 to 1. The larger the value is, the more certain the pixel in visible facial regions is. The visible region V is defined as all pixels where $\text{Seg}(x, y) > \theta$. θ denotes a threshold set to 0.5 in our experiments.

Confidence Map $\text{Cf} \in \mathbb{R}_{H \times W}$ represents the reliability of Cor and Dep. Since there is no ground truth of it, we compute the ground truth $\text{Cf}(x, y)^*$ as follows:

$$\text{Cf}(x, y)^* = \sqrt{\text{Cf}_{\text{Cor}}(x, y)^* \text{Cf}_{\text{Dep}}(x, y)^*}, \quad (6)$$

where:

$$\text{Cf}_{\text{Cor}}(x, y)^* = \exp\left(-\frac{|\text{Cor}(x, y) - \text{Cor}(x, y)^*|}{a}\right), \quad (7)$$

$$\text{Cf}_{\text{Dep}}(x, y)^* = \exp\left(-\frac{|\text{Dep}(x, y) - \text{Dep}(x, y)^*|}{b}\right), \quad (8)$$

with $\text{Cf}_{\text{Cor}}(x, y)^*$ and $\text{Cf}_{\text{Dep}}(x, y)^*$ denote the ground truth confidence values of Cor(x, y) and Dep(x, y), respectively. Variables a and b are parameters that control the tolerance to the error of predicted Cor(x, y) and Dep(x, y), respectively. Fig. 4 illustrates that the confidence map effectively

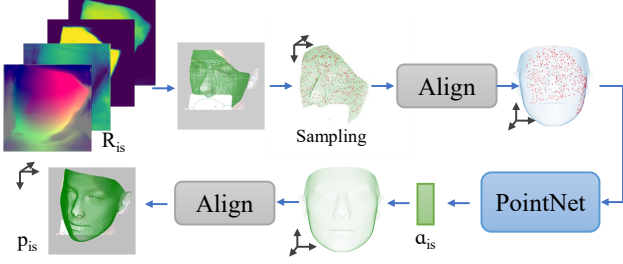


Figure 5. The post-process algorithm

learns that facial attributes region is more reliable, but UV-island boundaries [40] and side face area are less reliable.

The image space representation loss L_{Rep} is computed as follows:

$$L_{Rep} = w_1 L_{Cor} + w_2 L_{Dep} + w_3 L_{Seg} + w_4 L_{Cf}, \quad (9)$$

where:

$$L_{Cor} = \frac{1}{s(V)} \sum_{(x,y) \in V} \|\text{Cor}(x,y) - \text{Cor}(x,y)^*\|_2, \quad (10)$$

$$L_{Dep} = \frac{1}{s(V)} \sum_{(x,y) \in V} \|\text{Dep}(x,y) - \text{Dep}(x,y)^*\|_2, \quad (11)$$

$$L_{Seg} = \frac{1}{HW} \sum_{(x,y) \in I} \|\text{Seg}(x,y) - \text{Seg}(x,y)^*\|_2, \quad (12)$$

$$L_{Cf} = \frac{1}{s(V)} \sum_{(x,y) \in V} \|\text{Cf}(x,y) - \text{Cf}(x,y)^*\|_2, \quad (13)$$

with $s(V)$ denotes the area of V . L_{Seg} teaches the network to extract facial region from non-facial region. L_{Dep} teaches the network to learn and represent visible facial geometry. L_{Cor} teaches the network how to register visible facial geometry to the face model. With L_{Cf} , the network learns the reliability of the predicted Cor and Dep.

3.2.2 A PointNet-based Post-process Algorithm

The image space representation contains only the predicted geometry in the visible regions. To infer the geometry in the invisible regions due to extra-occlusion or self-occlusion, we need to leverage the prior knowledge of 3D facial geometry in 3DMM.

Previous works use optimization-based post-process algorithms to achieve the conversion. However, the optimization-based algorithms are slower due to their iterative updating parameters. Moreover, the optimized results might have distortion or might be too rigid due to the inappropriate weight of the regularization items. To avoid these problems, we convert the image space representation into a point cloud format. We then propose a PointNet-based post-process algorithm, which is fast and reliable. Illustrated in Fig. 5, the algorithm has four steps:

Step 1: Convert Image Space Representation to Point Cloud We randomly sample m pixels from the visible regions V . Using orthographic projection, the 3D location P of every selected pixel X in the image view can be represented as (x, y, z) , where z is $\text{Dep}(x, y)$. P_c denotes the point cloud consisting of these selected points.

Step 2: Align to Canonical View [26] proposes a self-alignment module to align the face mesh to the image view by 68 facial landmarks. We use a similar algorithm Align to align the point cloud from the image view to the canonical view, but using a set of randomly selected points:

$$\hat{K}, T = \text{Align}(K, D, W), \quad (14)$$

where K and \hat{K} are the point clouds before and after the alignment, T is the transform vector of the alignment, D is the goal point cloud used to align to, W is the corresponding weights. We can find the values of P corresponding to the 3D point \hat{P} on a mean face's geometry in the canonical view by searching the point that has the nearest UV location to $\text{Cor}(x, y)$. The corresponding point cloud on the mean face is denoted as \bar{P}_c . At the same time, $\text{Cf}(x, y)$ indicates the reliability of P . Hence, the weight set of P_c is denoted as $\text{Cf}(P_c)$. Therefore, the transformed point cloud in the canonical view \hat{P}_c and the transformed vector from the image view to the canonical view T_c are represented as:

$$\hat{P}_c, T_c = \text{Align}(P_c, \bar{P}_c, \text{Cf}(P_c)). \quad (15)$$

Step 3: Compute 3DMM Coefficients from Point Cloud PointNet [24] is a commonly used model to process point clouds. We here propose a PointNet-based module that converts the point cloud into the 3DMM geometry coefficients α_{is} . Since the input point cloud \hat{P}_c is in the canonical view, the transformation modules in the PointNet are removed, which makes the model more lightweight. For every point \hat{P} , we concatenate its 3D location with $\text{Cor}(x, y)$ and $\text{Cf}(x, y)$ as the input of the PointNet-based module Pnet. This process is represented as:

$$\alpha_{is} = \text{Pnet}(\hat{P}_c, \text{Cor}(P_c), \text{Cf}(P_c)). \quad (16)$$

A previous work [41] also makes use of the PointNet-based module to regress the 3DMM coefficients and transform vectors by 3D points. Their module however only makes use of 68 sparse facial landmarks. Unlike this method, our module makes use of all the point clouds on the facial geometry in the visible regions, which contains much more geometry information. Our method separates the prediction of T_{is} and α_{is} by the previous alignment algorithm and this PointNet-based module instead of mixing them together. Our ablation studies in Sec. 4.7 demonstrates this disentanglement helps both of them.

Step 4: Align to Image View As α_{is} is obtained, the whole face mesh in the canonical view \hat{G}_{is} can be generated by 3DMM. The 3D facial geometry in the image view G_{is} can be acquired by aligning \hat{G}_{is} back to the image view:

$$P'_c, T_i = \text{Align}(\hat{P}_c, P_c, \text{Con}(P_c)), \quad (17)$$

$$G_{is} = T_i^m \hat{G}_{is}, \quad (18)$$

where P'_c is the point cloud which is transformed back to image view, T_i^m is the transform matrix reshaped from T_i which is the transformed vector from canonical view to image view.

The loss of training the PointNet-based post-process module is then expressed as:

$$L_{\text{post}} = w_\alpha \|\alpha_{is} - \alpha^*\|^2 + w_5 L_{\text{cons}}, \quad (19)$$

$$L_{\text{cons}} = \frac{1}{s(V)} \sum_{P \in V} \|P - P'\|^2, \quad (20)$$

where L_{cons} is the geometry consistency in the visible facial region between all 3D points in V and their corresponding 3D points P' in the image space prediction. In the training phase, the image space loss consists of two items, image space representation loss L_{Rep} and post-process loss L_{Post} as: $L_{is} = L_{\text{Rep}} + L_{\text{Post}}$.

3.3. Dual Space Fusion

In our investigation, we realize that model space prediction and image space prediction are complementary (Sec. 4.6). Hence, it is possible to fuse them together by a fusion module in order to take advantage of both of them. We fuse the predictions from the two branches in the UV space in a model-free way. Similar to [26], we represent a 3D face geometry G by an offset map Off which indicates the offset of the face's shape from the mean face shape and a position map Pos , which indicates the face's pose in image view:

$$S = \bar{S} + \text{Off}, \quad (21)$$

$$G = \text{Align}(S, \text{Pos}, \text{Cf}). \quad (22)$$

The fusion module, as shown in Fig. 3 (c) (E_f and D_f), is a lightweight U-Net [25], which has two heads for predicting fused offset map Off_f and fused position map Pos_f separately. Model space feature f_{ms} is concatenated with the feature map of the center layer of the U-Net, in order to be aware of which branch's prediction is more reliable. The fused prediction G_f is represented as:

$$\text{Off}_f, \text{Pos}_f = F(\text{Off}_{ms}, \text{Pos}_{ms}, \text{Off}_{is}, \text{Pos}_{is}, f_{ms}), \quad (23)$$

$$G_f = \text{Align}(\bar{S} + \text{Off}_f, \text{Pos}_f, \text{Cf}). \quad (24)$$

The training loss of the whole pipeline is the combination of the model space loss L_{ms} , image space loss L_{is} , and fusion module loss L_f which consists of L_2 loss of Off and

Pos :

$$L_f = w_6 \|\text{Off} - \text{Off}^*\|_2 + w_7 \|\text{Pos} - \text{Pos}^*\|_2, \quad (25)$$

$$L = L_{ms} + L_{is} + L_f. \quad (26)$$

4. Experiments

In this section, we evaluate our method's ability to model 3d facial geometry qualitatively and quantitatively. We mainly compare our method with methods that have state-of-the-art robustness to occlusion or large view angle, including 3DDFA-V2 [12], SADRNet [26], SynergyNet [41].

4.1. Implement Details

We train our model end-to-end on the 300W-LP dataset [47], which contains more than 122k synthetic face images across different angles. We follow the preprocess and data augmentation of PRN [11] and SADRNet [26], where input images are cropped by their ground truth bounding box and rescaled to size 256×256 . We select around 45K vertices on the front face region of the model. HRNet-W18 [30] is adopted as the backbone of the encoder. The whole pipeline has 10.58M trainable parameters, while the image space branch has 5.85M parameters, comparable to the most lightweight 3D dense face alignment algorithms [12, 41], but with better performance. On average, our method processes one image in 70 ms using a Nvidia RTX 2080 Ti GPU. During training, the learning rate starts at $5e-6$ and reaches $5e-5$ after 4 epochs of warm-up. Then, the learning rate decays by a factor of 0.85 every epoch using an exponential scheduler. The training process is optimized by Adam optimizer and lasts 30 epochs where the batch size is set to 12 and weight decay is set to $1e-4$. Please refer to the supplementary material for more details.

4.2. Evaluation Datasets

We use multiple datasets to evaluate our method's performance on 3D dense face alignment, 3D face reconstruction, and head pose estimation tasks.

AFLW2000-3D [47] is a widely used testing set for 3D dense face alignment, reconstruction and head pose estimation. It contains 2000 unconstrained images. Part of them have occlusion and large poses, so they can be used to evaluate the model's ability to handle challenging cases. We use this dataset to evaluate our model on all three tasks.

AFLW2000-3D-occlusion is a variant of AFLW2000-3D collected by us to especially evaluate a model's robustness to occlusion. It has three subsets: 1. Naturally Occluded Dataset (NOD) contains 127 automatically selected images from AFLW2000-3D. 2. Color Synthetically Occluded Dataset (CSOD) contains 6000 images where every image in AFLW2000-3D is occluded by three different types of color, which are similar to the occlusion patterns in [32]. 3. NatOcc Synthetically Occluded Dataset (NSOD)

Method	2D Sparse Face Alignment			
	0 to 30	30 to 60	60 to 90	Mean
Dense Corr [44]	3.62	6.06	9.56	6.41
3DDFA [47]	3.78	4.54	7.93	5.42
N3DMM [35]	-	-	-	4.12
3DSTN [3]	3.15	4.33	5.98	4.49
PRN [11]	2.75	3.51	4.61	3.62
DeFA [20]	-	-	-	4.50
DAMDN [15]	2.90	3.83	4.95	3.89
SPDT [23]	3.56	4.06	4.11	3.88
3DDFA-V2 [12]	2.63	3.42	4.48	3.51
2DASL [36]	2.75	3.51	4.61	3.62
SADRNet [26]	2.66	3.30	4.42	3.46
SynergyNet [41]	2.65	3.30	4.27	3.41
DSFNet-f (ours)	2.46	3.20	4.16	3.27

Table 1. Sparse face alignment (68 landmarks) on AFLW2000-3D. The NME (%) for faces with different yaw angles are reported.

is generated by Naturalistic Occlusion Generation (NatOcc) technique from [38] using full AFLW2000-3D dataset. Details and samples of these subsets are shown in the supplementary material.

4.3. Face Alignment and Reconstruction

For this task, we evaluate our method on the AFLW2000-3D benchmark. Following [47], we adopt normalized mean error (NME) where vertices are normalized by the geometric mean of the height and width of the bounding box of the evaluated vertices, as the evaluation metric of face alignment. For 2D sparse alignment, we evaluate the NME of 68 facial landmarks on the AFLW2000-3D dataset, shown in Tab. 1, where our method has the best performance.

For 3D dense alignment, we evaluate the NME of the core facial region which has around 45K points on AFLW2000-3D. Tab. 2 shows that our image space branch solely reaches a new state-of-the-art result. Although the model space branch itself does not have top performance, fusing it with the image space branch is able to achieve a better result. Our whole pipeline outperforms existing methods by a large margin, which is a 5.5% improvement compared to the existing SOTA method [26]. We notice that the model space feature f_{ms} helps the fusion module to learn how to fuse them together in different cases. Without f_{ms} , the fusion module will even hurt the image space branch’s performance.

For 3D face reconstruction, we evaluate our method on AFLW2000-3D by NME normalized by outer interocular distance after aligning to the ground truth. Tab. 2 shows our image space branch achieves the best result.

4.4. Head Pose Estimation

For head pose estimation task, following the metric in [5, 41], we compute the mean absolute error (MAE) of each Euler angle and the MAE of all three Euler angles. We evaluate our method on images with yaw angle ranged from -99

Method	AFLW2000-3D 3D Facial Geometry	
	Dense Alignment	Reconstruction
3DDFA [47]	6.55	5.36
DeFA [20]	6.04	5.64
PRN [11]	4.40	3.96
3DDFA-V2 [12]	4.18	-
SynergyNet [41]	4.06	-
SADRNet [26]	4.02	3.25
DSFNet-ms (ours)	4.15	3.46
DSFNet-is (ours)	3.89	3.16
DSFNet-f (ours) w/o f_{ms}	3.96	3.27
DSFNet-f (ours)	3.80	3.24

Table 2. 3D dense face alignment and reconstruction results on AFLW2000-3D. NME (%) normalized by bounding box side length and NME (%) normalized by 3D outer interocular distance are reported separately for 3D dense alignment and 3D reconstruction tasks. $-ms$, $-is$, $-f$ denote our method’s model space prediction, image space prediction, and fusion model prediction separately, where all modules are trained jointly in the fusion model.

AFLW2000-3D	Yaw	Pitch	Roll	Mean
FSANet [42]	4.50	6.08	4.64	5.07
TriNet [5]	4.20	5.77	4.04	3.97
RankPose [7]	2.99	4.75	3.25	3.66
3DDFA [47]	4.33	5.98	4.30	4.87
2DASL [36]	3.85	5.06	3.50	4.13
FDN [45]	3.78	5.61	3.88	4.42
MNN [37]	3.34	4.69	3.48	3.83
3DDFA-V2 [12]	4.06	5.26	3.48	4.27
WHENet-V [46]	4.44	5.75	4.31	4.83
img2pose [1]	3.43	5.03	3.28	3.91
SADRNet [26]	2.93	5.00	3.54	3.82
SynergyNet [41]	3.42	4.09	2.55	3.35
6DRepNet [13]	3.63	4.91	3.37	3.97
DSFNet-f (ours)	2.65	4.28	2.82	3.25

Table 3. Head pose estimation results on AFLW2000-3D dataset.

degrees to 99 degrees in AFLW2000-3D. Tab. 3 shows the superiority of our method. In particular, our method infers much more accurate yaw angle prediction compared to existing methods.

4.5. Occlusion Robustness

Benefiting from the image space branch, our method is robust to occlusion and large view angles as shown in Fig. 1. Quantitative evaluation on the three occlusion-specific subsets (i.e., one natural occlusion, NOD, and two synthetic occlusion subsets, CSOD and NSOD) is performed to further demonstrate the superior robustness of the proposed method compared with the SOTA methods as shown Tab. 4.

4.6. Complementary Relationship between Image Space and Model Space Branches

Although the image space branch solely already achieves the SOTA accuracy and occlusion robustness, it has some drawbacks. In our framework, the model space branch complements the image space branch as they focus on different

Method	NOD	CSOD	NSOD
SynergyNet [41]	5.49	7.95	14.66
SADRNet [26]	6.70	5.31	8.70
DSFNet-ms (ours)	5.23	5.06	6.62
DSFNet-is (ours)	4.78	4.59	5.83
DSFNet-f (ours)	4.80	4.44	5.40

Table 4. Occlusion resistance results on AFLW2000-3D-occlusion. NME (%) of 3d dense face alignment is reported.




Figure 6. Failure case of our method.

information. The image space can handle large pitch and roll angles where existing methods fail, while it has more difficulty handling large yaw angles. The reason is straightforward: In large yaw angle cases, most of the visible region is in the side face where fewer facial attributes lay, which leads the image space branch hard to deduce the accurate UV locations of visible pixels. Second, we train our model on 300W-LP which is a synthetic dataset consisting of images generated by 3D rotation around the yaw angle from the positive images, resulting in artifacts in the side face region. Since the image space branch relies more on low-level image information, it is sensitive to this data problem. In contrast, the model space branch focuses more on high-level information, and is more robust to large yaw angles. Similarly, due to sensitivity to low-level information, the image space is easier to be interfered by noise and blur. In this situation, model space prediction is an important complementary for the robustness of the whole pipeline.

4.7. Ablation Study of the Post-process Algorithm

To evaluate the contribution of each component in the proposed post-process algorithm, an ablation study is performed using the image space branch on AFLW2000-3D dataset as shown in Tab. 5. It is observed that disentangling the prediction of T_{is} and α_{is} by the Align algorithm and the PointNet separately is crucial to the high performance of the post-process algorithm.

4.8. Failure Cases and Limitations

Fig. 6 shows a typical failure case of our method. The pose and scale of our image space prediction rely on the number of the points sampled from the visible facial region. If the visible facial region is too small, the valid points will be too few. Thus, while our model space predicts the coarse pose based on the context information like the hair, image space prediction fails due to the small visible facial area,

xyz	Cor	Cf	L_{cons}	$Align$	NME	$diff$
✓			✓	✓	4.01	0.12
✓	✓		✓	✓	3.93	0.04
✓	✓	✓		✓	3.94	0.05
✓	✓	✓	✓		4.12	0.23
✓	✓	✓	✓	✓	3.89	-

Table 5. 3D dense face alignment results of our image space branch on AFLW2000-3D dataset with different post-process modules. The first three are the ablation study of the input components of the PointNet. xyz , Cor , Cf denotes using the point cloud Pc 's 3D locations from depth map, UV locations from correspondence map, and values from confidence map as the PointNet's input. L_{cons} denotes using L_{cons} in training. $Align$ denotes if the point cloud is aligned to the canonical view or not before being fed to the PointNet.

resulting in the wrong final prediction.

Besides, our method adopts a highly supervised training framework and relies on the quality of the labeled data. To address the limitation, we intend to adopt semi-supervised or self-supervised learning to make use of the unlabeled data in future work. Another issue is the assumption of orthographic projection, where most methods trained on 300W-LP [47], including ours, use weak perspective or orthographic projection. This limits our method in handling images taken from a close distance. Exploring the use of perspective projection [17] is an area for future work.

5. Conclusions

We have presented a dual space fusion network (DSFNet) for robust 3D facial geometry modeling in unconstrained scenarios. Instead of directly regressing the whole face, our image space branch predicts a 3D facial geometry's 2D image space representation, which decentralizes the instance-level prediction to pixel-level dense prediction. A novel PointNet-based post-process algorithm is proposed to recover the face's whole 3D geometry from the visible region information contained in the image space representation. Since it makes better use of low-level 2D spatial and semantic information, it achieves high robustness to occlusion and large view angles. We further proposed a fusion module to combine advantages from both image and model space branches, resulting in a new state-of-the-art performance in mainstream benchmarks on 3D dense face alignment and reconstruction, and head pose estimation tasks.

Acknowledgments This research is supported by the National Research Foundation, Singapore under its Strategic Capability Research Centres Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

- [1] Vitor Albiero, Xingyu Chen, Xi Yin, Guan Pang, and Tal Hassner. img2pose: Face alignment and detection via 6dof, face pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7617–7627, 2021. 7
- [2] Riza Alp Guler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, and Iasonas Kokkinos. Densereg: Fully convolutional dense shape regression in-the-wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6799–6808, 2017. 3
- [3] Chandrasekhar Bhagavatula, Chenchen Zhu, Khoa Luu, and Marios Savvides. Faster than real-time facial alignment: A 3d spatial transformer network approach in unconstrained poses. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3980–3989, 2017. 7
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. 2
- [5] Zhiwen Cao, Zongcheng Chu, Dongfang Liu, and Yingjie Chen. A vector-based representation to enhance head pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1188–1197, 2021. 7
- [6] Daniel Crispell and Maxim Bazik. Pix2face: Direct 3d face model estimation. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 2512–2518, 2017. 3
- [7] Donggen Dai, Wangkit Wong, and Zhuojun Chen. Rankpose: Learning generalised feature with rank supervision for head pose estimation. *arXiv preprint arXiv:2005.10984*, 2020. 7
- [8] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *IEEE Computer Vision and Pattern Recognition Workshops*, 2019. 1, 2
- [9] Abdallah Dib, Junghyun Ahn, Cedric Thebault, Philippe-Henri Gosselin, and Louis Chevallier. S2f2: Self-supervised high fidelity face reconstruction from monocular image. *arXiv preprint arXiv:2203.07732*, 2022. 1, 2
- [10] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. volume 40, 2021. 1, 2
- [11] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. Joint 3d face reconstruction and dense alignment with position map regression network. In *ECCV*, 2018. 1, 2, 6, 7
- [12] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3d dense face alignment. In *European Conference on Computer Vision*, pages 152–168. Springer, 2020. 1, 2, 6, 7
- [13] Thorsten Hempel, Ahmed A Abdelrahman, and Ayoub Al-Hamadi. 6d rotation representation for unconstrained head pose estimation. *arXiv preprint arXiv:2202.12555*, 2022. 7
- [14] Aaron S Jackson, Adrian Bulat, Vasileios Argyriou, and Georgios Tzimiropoulos. Large pose 3d face reconstruction from a single image via direct volumetric cnn regression. In *Proceedings of the IEEE international conference on computer vision*, pages 1031–1039, 2017. 2
- [15] Lei Jiang, Xiao-Jun Wu, and Josef Kittler. Dual attention mobdensenet (damdnet) for robust 3d face alignment. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, pages 0–0, 2019. 7
- [16] Amin Jourabloo and Xiaoming Liu. Large-pose face alignment via cnn-based dense 3d model fitting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4188–4196, 2016. 1
- [17] Yueying Kao, Bowen Pan, Miao Xu, Jiangjing Lyu, Xiangyu Zhu, Yuanzhang Chang, Xiaobo Li, Zhen Lei, and Zixiong Qin. Single-image 3d face reconstruction under perspective projection. *arXiv preprint arXiv:2205.04126*, 2022. 3, 8
- [18] Tatsuro Koizumi and William AP Smith. “look ma, no landmarks!”—unsupervised, model-based dense face alignment. In *European Conference on Computer Vision*, pages 690–706. Springer, 2020. 2, 3
- [19] Chunlu Li, Andreas Morel-Forster, Thomas Vetter, Bernhard Egger, and Adam Kortylewski. To fit or not to fit: Model-based face reconstruction and occlusion segmentation from weak supervision. *arXiv preprint arXiv:2106.09614*, 2021. 1, 2
- [20] Yaojie Liu, Amin Jourabloo, William Ren, and Xiaoming Liu. Dense face alignment. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1619–1628, 2017. 7
- [21] Yanda Meng, Xu Chen, Dongxu Gao, Yitian Zhao, Xiaoyun Yang, Yihong Qiao, Xiaowei Huang, and Yalin Zheng. 3d dense face alignment with fused features by aggregating cnns and gcns. *arXiv preprint arXiv:2203.04643*, 2022. 1
- [22] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance*, pages 296–301. Ieee, 2009. 3
- [23] Jingtian Piao, Chen Qian, and Hongsheng Li. Semi-supervised monocular 3d face reconstruction with end-to-end shape-preserved domain transfer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9398–9407, 2019. 7
- [24] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 5
- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 6
- [26] Zeyu Ruan, Changqing Zou, Longhai Wu, Gangshan Wu, and Limin Wang. Sadrnet: Self-aligned dual face regression networks for robust 3d dense face alignment and reconstruc-

- tion. *IEEE Transactions on Image Processing*, 30:5793–5806, 2021. 1, 2, 5, 6, 7, 8
- [27] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael J Black. Learning to regress 3d face shape and expression from an image without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7763–7772, 2019. 1, 2
- [28] Matan Sela, Elad Richardson, and Ron Kimmel. Unrestricted facial geometry reconstruction using image-to-image translation. *arxiv*, 2017. 3
- [29] Jiayang Shang, Tianwei Shen, Shiwei Li, Lei Zhou, Mingmin Zhen, Tian Fang, and Long Quan. Self-supervised monocular 3d face reconstruction by occlusion-aware multi-view geometry consistency. *arXiv preprint arXiv:2007.12494*, 2020. 1, 2
- [30] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 3, 6
- [31] Hitika Tiwari, Min-Hung Chen, Yi-Min Tsai, Hsien-Kai Kuo, Hung-Jen Chen, Kevin Jou, KS Venkatesh, and Yong-Sheng Chen. Self-supervised robustifying guidance for monocular 3d face reconstruction. *arXiv preprint arXiv:2112.14382*, 2021. 1, 3
- [32] Hitika Tiwari, Vinod K Kurmi, KS Venkatesh, and Yong-Sheng Chen. Occlusion resistant network for 3d face reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 813–822, 2022. 1, 3, 6
- [33] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gerard Medioni. Regressing robust and discriminative 3d morphable models with a very deep neural network. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [34] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard Medioni. Extreme 3D face reconstruction: Seeing through occlusions. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [35] Luan Tran and Xiaoming Liu. On learning 3d face morphable model from in-the-wild images. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):157–171, 2019. 7
- [36] Xiaoguang Tu, Jian Zhao, Mei Xie, Zihang Jiang, Akshaya Balamurugan, Yao Luo, Yang Zhao, Lingxiao He, Zheng Ma, and Jiashi Feng. 3d face reconstruction from a single image assisted by 2d face images in the wild. *IEEE Transactions on Multimedia*, 23:1160–1172, 2020. 7
- [37] Roberto Valle, José M Buenaposada, and Luis Baumela. Multi-task head pose estimation in-the-wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(8):2874–2881, 2020. 7
- [38] Kenny T. R. Voo, Liming Jiang, and Chen Change Loy. Delving into high-quality synthetic face occlusion segmentation datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2022. 1, 7
- [39] Yandong Wen, Weiyang Liu, Bhiksha Raj, and Rita Singh. Self-supervised 3d face reconstruction via conditional estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13289–13298, 2021. 1, 2
- [40] Erroll Wood, Tadas Baltrusaitis, Charlie Hewitt, Matthew Johnson, Jingjing Shen, Nikola Milosavljevic, Daniel Wilde, Stephan Garbin, Toby Sharp, Ivan Stojiljkovic, et al. 3d face reconstruction with dense landmarks. *arXiv preprint arXiv:2204.02776*, 2022. 3, 5
- [41] Cho-Ying Wu, Qiangeng Xu, and Ulrich Neumann. Synergy between 3dmm and 3d landmarks for accurate 3d facial geometry. In *2021 International Conference on 3D Vision (3DV)*, 2021. 1, 2, 5, 6, 7, 8
- [42] Tsun-Yi Yang, Yi-Ting Chen, Yen-Yu Lin, and Yung-Yu Chuang. Fsa-net: Learning fine-grained structure aggregation for head pose estimation from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1087–1096, 2019. 7
- [43] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341, 2018. 2
- [44] Ronald Yu, Shunsuke Saito, Haoxiang Li, Duygu Ceylan, and Hao Li. Learning dense facial correspondences in unconstrained images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4723–4732, 2017. 3, 7
- [45] Hao Zhang, Mengmeng Wang, Yong Liu, and Yi Yuan. Fdn: feature decoupling network for head pose estimation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 12789–12796, 2020. 7
- [46] Yijun Zhou and James Gregson. Whenet: Real-time fine-grained estimation for wide range head pose. *arXiv preprint arXiv:2005.10353*, 2020. 7
- [47] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z Li. Face alignment across large poses: A 3d solution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 146–155, 2016. 1, 2, 3, 6, 7, 8
- [48] Xiangyu Zhu, Fan Yang, Di Huang, Chang Yu, Hao Wang, Jianzhu Guo, Zhen Lei, and Stan Z Li. Beyond 3dmm space: Towards fine-grained 3d face reconstruction. In *ECCV*, pages 343–358. Springer, 2020. 2
- [49] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. *arXiv preprint arXiv:2204.06607*, 2022. 1, 2