

# VGGTFace: Topologically Consistent Facial Geometry Reconstruction in the Wild

Xin Ming, Yuxuan Han\*, Tianyu Huang, Feng Xu

BNRist and School of Software, Tsinghua University  
1729406968@qq.com, hanyx22@mails.tsinghua.edu.cn, huang-ty21@mails.tsinghua.edu.cn, xufeng2003@gmail.com

## Abstract

Reconstructing topologically consistent facial geometry is crucial for the digital avatar creation pipelines. Existing methods either require tedious manual efforts, lack generalization to in-the-wild data, or are constrained by the limited expressiveness of 3D Morphable Models. To address these limitations, we propose VGGTFace, an automatic approach that innovatively applies the 3D foundation model, *i.e.* VGGT, for topologically consistent facial geometry reconstruction from in-the-wild multi-view images captured by everyday users. Our key insight is that, by leveraging VGGT, our method naturally inherits strong generalization ability and expressive power from its large-scale training and point map representation. However, it is unclear how to reconstruct a topologically consistent mesh from VGGT, as the topology information is missing in its prediction. To this end, we augment VGGT with Pixel3DMM for injecting topology information via pixel-aligned UV values. In this manner, we convert the pixel-aligned point map of VGGT to a point cloud with topology. Tailored to this point cloud with known topology, we propose a novel Topology-Aware Bundle Adjustment strategy to fuse them, where we construct a Laplacian energy for the Bundle Adjustment objective. Our method achieves high-quality reconstruction in 10 seconds for 16 views on a single NVIDIA RTX 4090. Experiments demonstrate state-of-the-art results on benchmarks and impressive generalization to in-the-wild data.

**Code** — <https://github.com/grignarder/vggface>

## Introduction

Reconstructing facial geometry with consistent topology is crucial for digital avatar pipelines (Alexander et al. 2009, 2013), as it enables dense mesh correspondence and transferable animation, rigging, and texture editing. Traditional workflows first reconstruct a detailed facial scan and then enforce topological consistency by fitting a template mesh (Egger et al. 2020; Riviere et al. 2020). However, achieving high quality typically requires substantial manual effort—such as correspondence marking and parameter tuning—making these pipelines difficult to scale to everyday users.

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

There exist some pioneering works that attempt to reconstruct a topologically consistent face mesh automatically. Some works take multi-view images as input and use neural networks to directly predict the results (Li, Kang, and He 2024; Bolkart, Li, and Black 2023; Li et al. 2021). Training on studio-captured data, these methods can only infer images captured in a similar studio. However, they cannot generalize to in-the-wild faces captured by everyday users, as the lighting conditions and camera configurations differ significantly from their training data. Another group of works exploits the 3D Morphable Face Model (3DMM) for topologically consistent facial geometry reconstruction (Tewari et al. 2019; Deng et al. 2019; Feng et al. 2021; Wang et al. 2024). Building on top of the 3DMM prior, these methods can leverage large-scale face image or video datasets for training. However, their expressiveness is confined to the limited 3DMM space. Although some works propose more powerful representations, such as 3DMM with per-vertex displacement (Lei et al. 2023) or corrective basis (Bai et al. 2020), their dependency on 3DMM inevitably constrains the reconstruction quality, as shown by Li et al. (2021) and our experiments.

Recently, 3D foundation models, such as VGGT (Wang et al. 2025), have reshaped the field of geometry reconstruction. Given multi-view images captured in the wild, VGGT can infer the pixel-aligned point map (*i.e.* the 3D position of each pixel) and camera intrinsic and extrinsic for each view in a single forward pass. By training on large-scale datasets, VGGT is robust to diverse lighting conditions and camera configurations that appear in the data captured by everyday users. Meanwhile, the point map representation adopted by VGGT has enough expressiveness to capture high-quality facial geometry. These nice properties motivate us to apply VGGT for topologically consistent facial geometry reconstruction. However, it is unclear how to fulfill this as topology information is missing in the VGGT’s prediction.

To address this, we propose VGGTFace, a novel method for topologically consistent facial geometry reconstruction from in-the-wild multi-view images captured by everyday users. To best maintain the generalization ability of VGGT, we freeze its pretrained weights to infer pixel-aligned point maps and camera parameters from the captured multi-view face images. Next, we augment the point maps with pixel-aligned UV values using Pixel3DMM (Giebenhain et al.




Figure 1: VGGTFace. Given in-the-wild multi-view images (4 of 16 images are shown here) captured by everyday users as input, our method can reconstruct a topologically consistent mesh from these inputs in 10 seconds. Our method demonstrates strong generalization ability across lighting conditions of the capture site, facial expressions, and ethnic groups, opening the door for everyday users to scan themselves into the digital world. By innovatively applying the point map as the facial geometry representation, our method can reconstruct person-specific facial traits, *e.g.* asymmetric expressions, with high fidelity.

2025), where the UV values indicate the corresponding positions on a template face mesh. As each vertex on the template mesh has a unique pair of UV values, we assign the pixel with the same UV values as the 2D projected point of the vertex on this view. In this manner, we get the correspondence information to build the topologically consistent face mesh. However, as neither VGGT nor Pixel3DMM is perfectly accurate, the 3D position of the same template vertex computed from different views is inconsistent. Fusing them to minimize errors and thus reconstruct a topologically consistent mesh is still a challenging task.

To solve this problem, our key insight is to regard the corresponding pixel of the template vertex at each view as tracks (the target projected 2D points of a 3D point on different views (Hartley and Zisserman 2003)) and formulate the VGGT point cloud fusion as a Bundle Adjustment problem (Triggs et al. 1999). Specifically, we simultaneously optimize camera parameters and the fused point cloud to minimize the re-projection error *w.r.t* projected distance to the tracks. However, these tracks also contain errors, thus leading to a noisy fused point cloud after Bundle Adjustment (BA). To address this, existing works typically require tedious parameter tweaking to filter out the influence of low-quality tracks (Schönberger et al. 2016; Schönberger and Frahm 2016). In our task, as the VGGT point cloud has one-to-one correspondences to the template mesh after being augmented with UV coordinates, we can exploit the topology information of the template mesh to better perform BA. To be specific, we construct a Laplacian energy (Nicolet, Jacobson, and Jakob 2021) to the BA objective as a regularization to effectively overcome the low-quality tracks. In addition, the Laplacian energy can provide gradients to every 3D point, effectively regularizing unconstrained points that have less than two valid tracks in the BA system. Note that previous BA methods cannot use Laplacian energy as their point cloud does not have topology. Thus, we term our technique Topology-Aware Bundle Adjustment (TopBA). After the TopBA process, we directly connect the fused point cloud to obtain the result. Thanks to the highly optimized

BA solver (Agarwal, Mierle, and Team 2023), our method can reconstruct a high-quality topologically consistent mesh from 16 views in 10 seconds on a single NVIDIA RTX 4090. In conclusion, our main contributions include:

- We propose a novel method that can reconstruct topologically consistent facial geometry from in-the-wild multi-view images captured by everyday users in 10 seconds.
- We propose a novel Topology-Aware Bundle Adjustment (TopBA) technique to robustly distill a topologically consistent mesh from the raw point cloud predicted by VGGT.
- We demonstrate state-of-the-art geometry reconstruction results on various benchmarks and impressive generalization ability on in-the-wild captured data.

We hope our method provides a starting point for applying 3D foundation models to face-related tasks, helping overcome the scarcity of 3D facial data and enabling high-quality reconstruction for everyday users in the wild.

## Related Works

### Monocular Face Reconstruction

As reconstructing geometry from a single image is an ill-posed problem, most methods rely on the 3DMM prior (Blaiz and Vetter 1999; Zollhöfer et al. 2018). Given a single face image, they either optimize (Dib et al. 2021; Thies et al. 2016; Giebenhain et al. 2025) or learn a neural network to predict 3DMM parameters (Deng et al. 2019; Feng et al. 2021; Wang et al. 2024). However, their reconstruction quality is largely constrained by 3DMM itself. To improve the representation power, some works propose to augment 3DMM by adding per-vertex displacement (Lei et al. 2023), predicting a displacement map (Lei et al. 2023; Feng et al. 2021; Yang et al. 2020; Dib et al. 2024), or fine-tuning the 3DMM model on large-scale datasets (Tewari et al. 2018, 2019; R. et al. 2021; Han, Wang, and Xu 2023). Although their results are consistent with the 3DMM topology, the reconstruction quality is still confined by the

3DMM prior. Another group of works proposes to reconstruct facial geometry without 3DMM (Guo et al. 2023; Sela, Richardson, and Kimmel 2017; Zhang et al. 2021). Among these methods, the most relevant to us is Sela, Richardson, and Kimmel (2017); they learn a network to predict a UV coordinate image along with a depth image from the input image and then fit a template mesh to it to ensure topologically consistent reconstruction. Although we adopt a similar representation, our method works on a multi-view setup. Thus, we need to fuse the prediction from each view to a consistent mesh, which is challenging but not explored by Sela, Richardson, and Kimmel (2017). In addition, we apply a 3D foundation model, *i.e.* VGGT, to infer the point map instead of using a small model specially trained on facial data. Thus, we achieve more robust results as shown in our experiments.

### Multi-view Face Reconstruction

From multi-view images, traditional methods typically apply multi-view stereo to reconstruct a detailed scan and then register a template to it (Riviere et al. 2020), or simultaneously perform these steps (Fyffe et al. 2017). However, this process is often time-consuming and requires complex manual efforts to ensure high-quality results. Recent works apply a neural network to predict 3DMM parameters from multi-view images (Bai et al. 2020; Wu et al. 2019). Even with corrective basis (Bai et al. 2020), the construction quality is inevitably limited by the usage of 3DMM. More recently, some works (Li et al. 2021; Bolkart, Li, and Black 2023; Li, Kang, and He 2024; Xiao et al. 2022) propose to infer a topologically consistent mesh directly from the multi-view images. Although high-quality results are demonstrated, they cannot generalize to in-the-wild data captured by daily users, as they are trained on limited 3D facial data captured in the studio. In this paper, we innovatively apply the point map representation for multi-view topologically consistent facial geometry reconstruction. Compared to prior arts, our representation not only has powerful representation ability but also inherits strong generalization ability from VGGT.

### Foundation Models in Face-Related Problems

By training on large-scale datasets, Vision foundation models (Rombach et al. 2021; Caron et al. 2021; Oquab et al. 2023; Wang et al. 2025) have demonstrated impressive generalization ability on various downstream tasks. To improve the generalization ability, recent works augment the feature maps with these foundation models during training on face datasets (Giebenhain et al. 2025; Kirschstein et al. 2025; Ren et al. 2025). Instead, we distill high-quality topologically consistent facial geometry from raw predictions of VGGT in a zero-shot manner via test-time optimization.

### Method

Given in-the-wild multi-view images captured around a face by everyday users, our goal is to reconstruct a topologically consistent face mesh from these inputs. To ensure generalization to in-the-wild capture, we build our method on top of VGGT (Wang et al. 2025), a powerful 3D foundation model.

Specifically, we augment the VGGT raw predictions with facial topology and propose a novel Topology-Aware Bundle Adjustment technique to fuse the prediction in each view to a high-quality topologically consistent mesh. Before diving into the details, we first introduce the necessary background of VGGT and Bundle Adjustment.

### Preliminary

**VGGT** VGGT (Wang et al. 2025) is a recently proposed foundation model for 3D reconstruction. Given multi-view images  $\{I_i\}_{i=1}^V$ ,  $I_i \in \mathbb{R}^{H \times W \times 3}$ , it learns a transformer  $\mathcal{H}(\cdot)$  to infer the point map and camera intrinsic and extrinsic parameters for each view in a single network pass:

$$\{K_i, R_i, t_i, X_i\}_{i=1}^V = \mathcal{H}(\{I_i\}_{i=1}^V) \quad (1)$$

Here,  $K_i \in \mathbb{R}^{3 \times 3}$  denotes the camera intrinsic matrix,  $R_i \in \mathbb{R}^{3 \times 3}$  and  $t_i \in \mathbb{R}^3$  are the camera extrinsic parameters, and  $X_i \in \mathbb{R}^{H \times W \times 3}$  is the pixel-aligned point map which records the 3D position of each pixel; the subscript  $i$  means the prediction is corresponding to the  $i$ -th view. All the predictions of VGGT are in the coordinate frame of the first camera. Please refer to their paper for more details.

By training on large-scale datasets, VGGT demonstrates superior generalization ability and achieves state-of-the-art results on various benchmarks. To best maintain its generalization ability, we freeze its pretrained weights and propose novel techniques to distill a topologically consistent mesh from its raw predictions, *i.e.*  $\{K_i, R_i, t_i, X_i\}_{i=1}^V$ .

**Bundle Adjustment** Bundle Adjustment (Triggs et al. 1999) is a process that simultaneously optimizes the initial estimation of the camera parameters  $\{K_i, R_i, t_i\}_{i=1}^V$  and  $N$  3D points  $\{p_j\}_{j=1}^N$  to make them better match the 2D tracks  $q_{ij}$  after the perspective projection  $\pi(\cdot)$ :

$$\min_{\{K_i, R_i, t_i\}_{i=1}^V, \{p_j\}_{j=1}^N} \sum_{i=1}^V \sum_{j=1}^N v_{ij} \cdot \|\pi(K_i, R_i, t_i, p_j) - q_{ij}\|_2^2 \quad (2)$$

Here,  $p_j \in \mathbb{R}^3$  is the 3D point,  $q_{ij} \in \mathbb{R}^2$  is the target screen coordinate of  $p_j$  on the  $i$ -th view, and  $v_{ij} \in \{0, 1\}$  is a flag indicating whether the track is valid. Note we use the subscript  $i$  to index the view and  $j$  to index the 3D points; we will also adhere to this convention in the following paper. Eq. (2) is a non-linear least square problem, which is well-studied in the literature and can be solved efficiently by the Ceres solver (Agarwal, Mierle, and Team 2023).

### Augment VGGT Predictions with Topology

As shown in Figure 2, given  $V$  multi-view face images captured by everyday users, we first feed them to VGGT to obtain the point map  $X_i$  and camera parameters  $\{K_i, R_i, t_i\}$  for each view using Eq. (1).

However, applying VGGT for topologically consistent facial geometry reconstruction is not easy, as the raw prediction of VGGT does not contain any topology information, as shown in Eq. (1). Our key observation is that the geometry of VGGT is represented in a pixel-aligned fashion, *i.e.* the point map  $X_i$ . Thus, we can inject extra information into the

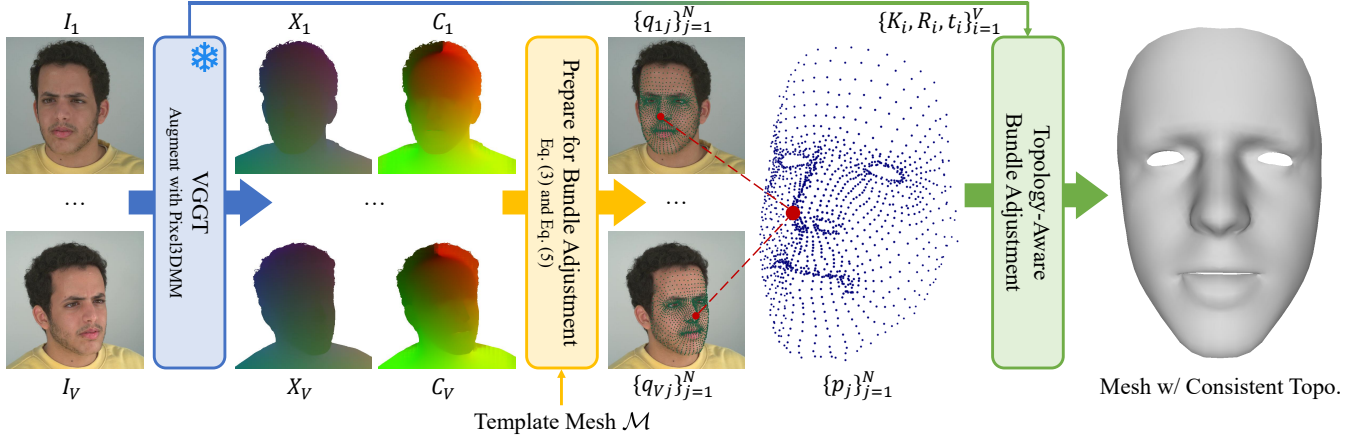


Figure 2: The framework of VGGTFace. Given  $V$  multi-view images  $\{I_i\}_{i=1}^V$  as input, we first augment VGGT with Pixel3DMM to obtain the camera parameters  $\{K_i, R_i, t_i\}_{i=1}^V$ , point map  $\{X_i\}_{i=1}^V$ , and UV coordinate image  $\{C_i\}_{i=1}^V$  for each view. We then convert these raw predictions to a set of tracks  $\{q_{ij}\}$  and an initial point cloud  $\{p_j\}_{j=1}^N$  according to the template mesh  $\mathcal{M}$ . Next, we propose a novel Topology-Aware Bundle Adjustment technique to simultaneously optimize the camera parameters  $\{K_i, R_i, t_i\}_{i=1}^V$  and the point cloud  $\{p_j\}_{j=1}^N$  to better match the tracks. After that, we connect the point cloud with the topology of  $\mathcal{M}$  to obtain a mesh with consistent topology.

prediction of VGGT using any screen-space method, which is a format widely adopted in face-related problems such as face parsing (Zheng et al. 2022) and facial intrinsic decomposition (Kim et al. 2024).

To augment the predictions of VGGT with topology, we use an off-the-shelf method, Pixel3DMM (Giebenhain et al. 2025), to predict the UV coordinate image  $\{C_i\}_{i=1}^V$ ,  $C_i \in \mathbb{R}^{H \times W \times 2}$  for each view as shown in Figure 2. This way, for each pixel in each view, we have not only its 3D position predicted by VGGT, but also its UV coordinate predicted by Pixel3DMM. Then, we can use the UV coordinate to build correspondence between the vertex on the template mesh and the pixel in each view. After that, we can transfer the topology information defined on the template mesh to the VGGT’s prediction.

Specifically, given a template mesh  $\mathcal{M}^1$  that has  $N$  vertices with the corresponding UV coordinate  $\{c_j\}_{j=1}^N$ ,  $c_j \in \mathbb{R}^2$  and the topology  $\mathcal{F}$  which defines the connection information of the  $N$  vertices, we can lookup its corresponding pixel on the  $i$ -th view via nearest neighbor search in the UV distance space:

$$\{q_{ij}^x, q_{ij}^y\} = \operatorname{argmin} \|C_i[q_{ij}^x][q_{ij}^y] - c_j\|_2 \quad (3)$$

We denote  $q_{ij} = (q_{ij}^x, q_{ij}^y)$ ,  $q_{ij} \in \mathbb{R}^2$  as the corresponding screen coordinate of the  $j$ -th template vertex on the  $i$ -th view. Note that not all the vertices are visible in each view; we realize this visibility using a flag  $v_{ij}$  by thresholding the nearest distance with a tolerance  $\epsilon$ :

$$v_{ij} = \begin{cases} 0 & \|C_i[q_{ij}^x][q_{ij}^y] - c_j\|_2 \geq \epsilon \\ 1 & \|C_i[q_{ij}^x][q_{ij}^y] - c_j\|_2 < \epsilon \end{cases} \quad (4)$$

<sup>1</sup> $\mathcal{M}$  should have the same UV definition as the UV predictor, *i.e.* Pixel3DMM.

Then, we can transfer the topology  $\mathcal{F}$  to the VGGT’s predicted 3D point  $\{X_i[q_{ij}^x][q_{ij}^y]\}_{j=1}^N$  on the  $i$ -th view.

However, as VGGT are not perfectly accurate, the 3D position of the same template vertex computed from different views are often inconsistent, *i.e.*  $X_{i_1}[q_{i_1j}^x][q_{i_1j}^y] \neq X_{i_2}[q_{i_2j}^x][q_{i_2j}^y]$  where  $i_1$  and  $i_2$  are the index of two different views. Naively fusing the 3D position computed from different views does not guarantee valid results, as shown in Figure 5 (the *w/o BA* baseline). That is because the fusion results are not supervised during the training of VGGT. To solve this problem, our key insight is to regard the screen coordinates of the template vertex, *i.e.*  $q_{ij}$ , as tracks and formulate the point cloud fusion as the Bundle Adjustment problem.

### Topology-Aware Bundle Adjustment

To enable Bundle Adjustment, we first fuse the same template vertex computed from different views, *i.e.*  $\{X_i[q_{ij}^x][q_{ij}^y]\}_{i=1}^V$ , to a single 3D point  $p_j$  to serve as the initial point cloud of Eq. (2); recall that  $X_i[q_{ij}^x][q_{ij}^y]$  is the 3D position of the  $j$ -th template vertex computed from the  $i$ -th view. We define the fusing operator as a simple weighted average by the visibility  $v_{ij}$ ; this is not a critical step and can be implemented in other ways, as the fused point cloud only acts as an initialization in BA:

$$p_j = \frac{\sum_{i=1}^V v_{ij} \cdot X_i[q_{ij}^x][q_{ij}^y]}{\sum_{i=1}^V v_{ij}} \quad (5)$$

We visualize  $\{p_j\}_{j=1}^N$  and the corresponding tracks  $q_{ij}$  in Figure 2. In practice, some template vertices are invisible to all the views because they cannot pass the visibility test in Eq. (4), *i.e.* for some  $j$ ,  $\sum_{i=1}^V v_{ij}$  is zero; for these points, we simply set their fused position to be the origin.

Now, by augmenting with Pixel3DMM, we have transferred the raw prediction of VGGT to the combination of

per-view camera parameters  $\{K_i, R_i, t_i\}$ , a set of tracks  $q_{ij}$  which denotes the screen coordinates of the  $j$ -th template vertex on the  $i$ -th view along with their visibility  $v_{ij}$ , and a point cloud  $\{p_j\}_{j=1}^N$  which denotes the initial 3D positions of the template vertices. Then, we can directly send these parameters to Eq. (2) to optimize the camera parameters  $\{K_i, R_i, t_i\}$  and the position of the template vertices  $\{p_j\}_{j=1}^N$  to best match the tracks  $q_{ij}$ . After the Bundle Adjustment, we can connect  $\{p_j\}_{j=1}^N$  via the topology  $\mathcal{F}$  to obtain a topologically consistent mesh.

However, we observe that the resulting mesh is noisy if naively following the optimization in Eq. (2), as shown in Figure 5 (the *w/o Laplacian* baseline). That is because the tracks  $q_{ij}$  are computed from Pixel3DMM, which is a neural network and thus not perfectly accurate. In addition, if a template vertex has only one or zero visible tracks, *i.e.*  $\sum_{i=1}^V v_{ij} < 2$ , it is unconstrained in Eq. (2) (Triggs et al. 1999). To address the inaccurate track problem, traditional works like COLMAP (Schönberger et al. 2016; Schönberger and Frahm 2016) typically require tedious parameter tweaking to filter out the influence of low-quality tracks. Although more robust to the inaccurate tracks, they still cannot handle points with fewer than two visible tracks. Ignoring these points without enough visible tracks is unacceptable in our method, as all the points correspond to the template mesh; deleting any of them will destroy the topology.

Fortunately, our key insight is that our 3D points  $\{p_j\}_{j=1}^N$  have topology, which is not satisfied in the scenario of existing works. Motivated by this observation, we propose adding a Laplacian term to the Bundle Adjustment objective:

$$\min_{\{K_i, R_i, t_i\}_{i=1}^V, \{p_j\}_{j=1}^N} \sum_{i=1}^V \sum_{j=1}^N v_{ij} \cdot \|\pi(K_i, R_i, t_i, p_j) - q_{ij}\|^2 + \sum_{j=1}^N \left\| p_j - \frac{1}{|\mathcal{N}(j)|} \cdot \sum_{k \in \mathcal{N}(j)} p_k \right\| \quad (6)$$

Here,  $\mathcal{N}(j)$  denotes the 1-ring neighborhood list of the  $j$ -th vertex on the template mesh  $\mathcal{M}$ , which can be computed efficiently given the topology  $\mathcal{F}$ , and  $|\mathcal{N}(j)|$  denotes the number of neighborhoods of the  $j$ -th vertex. This Laplacian term not only enforces a smooth regularization on the 3D points to make the optimization more robust to inaccurate tracks, but also adds a direct constraint on the 3D points that lack enough visible tracks. We emphasize that the latter benefit, *i.e.*, ensuring every 3D point is supervised in the BA process, is important in our problem, but has not been explored in the literature. As shown in Figure 5, our Topology-Aware Bundle Adjustment method obtains significantly better results than the naive baseline.

In practice, we modify the original BA Ceres solver to solve Eq. (6). Specifically, we iteratively minimize the original re-projection error and the Laplacian constraint until convergence. Thanks to the highly optimized Ceres solver, the whole method takes less than 10 seconds to process 16 views on a single NVIDIA RTX 4090, including the inference time of VGGT and Pixel3DMM, the time to prepare inputs for BA, and the optimization process of Eq. (6).

Method	H3DS	NeRsemble			Time	# Views
	Mean ↓	Mean ↓	Median ↓	Std ↓		
DECA	1.99	1.58	1.55	1.18	< 1s	1
3DDFA-V3	1.81	1.52	1.46	1.08	< 1s	1
Pixel3DMM	1.69	1.42	1.38	1.01	~ 70s	1
DFNRMVS	1.78	1.62	1.58	1.09	~ 30s	16
HRN	1.46	1.49	1.43	1.10	~ 160s	16
NeuS2	2.11	2.58	2.52	2.39	~ 300s	16
Ours	<b>1.18</b>	<b>0.98</b>	<b>0.98</b>	<b>0.65</b>	< 10s	16
w/o BA	1.52	1.44	1.37	1.04	< 10s	16
w/o Laplacian	1.20	0.99	0.98	0.65	< 10s	16

Table 1: Quantitative comparison on H3DS and NeRsemble benchmark. All metrics are in millimeters (mm).

## Experiments

We first provide implementation details of our method. We then compare our method to baseline methods on various benchmarks. Next, we evaluate several key design choices in our method. Lastly, we discuss the limitations and future works. Due to the page limit, we urge the reader to check our *suppl. material* for more results.

### Implementation Details

Our system takes 16 input views unless otherwise specified. Although it can process arbitrary numbers of images, we find 16 provides a good balance between view coverage and runtime; for in-the-wild videos, we uniformly sample 16 frames. Input images are center-cropped and resized to  $518 \times 518$  before being fed to VGGT. Following Wang et al. (2025), we compute the point map using VGGT’s depth-branch output, which is empirically more accurate than its point-map head. The visibility threshold  $\epsilon$  in Eq. (4) is set as the 70-th percentile of the nearest-distance distribution in each view, which performs slightly better than using a fixed value. We adopt FLAME (Li et al. 2017) as the template mesh  $\mathcal{M}$  for its UV compatibility with Pixel3DMM. Importantly, FLAME is used only for topology (vertex connectivity and UV layout); no FLAME shape or expression bases are used. The facial region is subdivided to approximately 5000 vertices. All experiments are performed on an NVIDIA RTX 4090.

### Comparisons

We first introduce the baseline methods we used for comparison. Then, we report the comparison results on the H3DS (Ramon et al. 2021; Caselles et al. 2025) and NeRsemble benchmark (Kirschstein et al. 2023).

**Baseline Methods** We compare state-of-the-art methods that can reconstruct a topologically consistent mesh from single-view (DECA (Feng et al. 2021), 3DDFA-V3 (Wang et al. 2024), and Pixel3DMM (Giebenhain et al. 2025)) or multi-view face images (HRN (Lei et al. 2023) and DFNRMVS (Bai et al. 2020)). We also involve Sela, Richardson, and Kimmel (2017) for comparison in the *suppl. material*, as their method adopts a similar representation as our method, but in a single-view setup. In addition, we compare



Figure 3: Qualitative comparison on the H3DS benchmark (the first row) and the NeRSemble benchmark (the last two rows). For each result, we pair it with the error map against the GT scan. Methods marked with \* adopt the BFM topology with around 40000 vertices, while Ours and Pixel3DMM follow the FLAME topology with around 5000 and 1500 vertices, respectively.

to NeuS2 (Wang et al. 2023), a well-known multi-view geometry reconstruction method for generic objects. We do not compare to ToFu (Li et al. 2021) and TEMPEH (Bolkart, Li, and Black 2023) as their method cannot generalize to camera configurations different from their training studio. Although GRAPE (Li, Kang, and He 2024) claims it can process camera configurations different from their training studio, we do not compare to it as its code is not released.

**Comparison on H3DS Benchmark** The H3DS benchmark (Ramon et al. 2021; Caselles et al. 2025) contains 59 subjects with neutral expression. Following SIRA++ (Caselles et al. 2025), we report the mean chamfer distance *w.r.t* to the ground truth mesh in Table 1 and present qualitative comparisons in Figure 3. Compared to existing works, our method obtains the best results. In addition, as an optimization-based method, ours is still significantly faster than other optimization-based methods such as Pixel3DMM, HRN, DFNRMV5, and NeuS2. That is because our method is based on the highly optimized BA solver.

**Comparison on NeRSemble Benchmark** As the H3DS benchmark only contains subjects with neutral expressions, we conduct comparisons on the more challenging NeRSemble dataset (Kirschstein et al. 2023) with facial expression variation. The original NeRSemble dataset only contains 16-view videos of multiple subjects performing expression sequences. We build a benchmark for facial geometry evaluation on top of it<sup>2</sup>. Specifically, we manually select 76

frames from the data of 40 subjects; see more details in our *suppl. material*. For each subject, we use MetaShape<sup>3</sup> to reconstruct a mesh from the provided 16-view images, which serves as the ground truth facial geometry (GT scan). We then manually annotate 6 landmarks on the GT scan for alignment. Given a reconstructed face mesh and the corresponding 6 landmarks, we first rigidly align it to the GT scan, and then compute the mean/median/std of the chamfer distance *w.r.t* the GT scan.

As shown in Table 1, our method is the only one that achieves an error less than 1 millimeter. As shown in Figure 3, our method can reconstruct challenging expressions with high fidelity. Compared to single-view methods, *e.g.* 3DDFA-V3 and Pixel3DMM, our method can better reconstruct the mouth shape due to the exploitation of multi-view information, which inherently has ambiguity in the single-view setup. Compared to multi-view methods, *e.g.* HRN and DFNRMV5, our method has stronger representation power stemming from the point map representation to model person-specific facial characteristics such as puffing and pouting (highlighted with green boxes in Figure 3); despite per-vertex displacement or corrective basis being adopted, HRN and DFNRMV5 are inevitably constrained by their dependency on the 3DMM.

**Comparison on In-the-Wild Data** As shown in Figure 4, we conduct qualitative and quantitative comparisons on in-the-wild data captured by ourselves. We involve

<sup>2</sup>Although Pixel3DMM also builds a similar geometry reconstruction benchmark on NeRSemble, upon the submission of our

paper, their benchmark is not publicly available.

<sup>3</sup><https://www.agisoftmetashape.com>

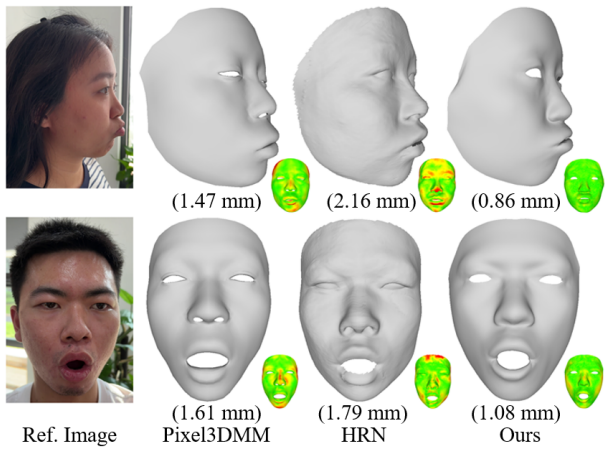


Figure 4: Comparisons on in-the-wild data. For each result, we pair it with the error map against the GT scan with the same colorbar as Figure 3. We report the mean chamfer distance below each method’s result.

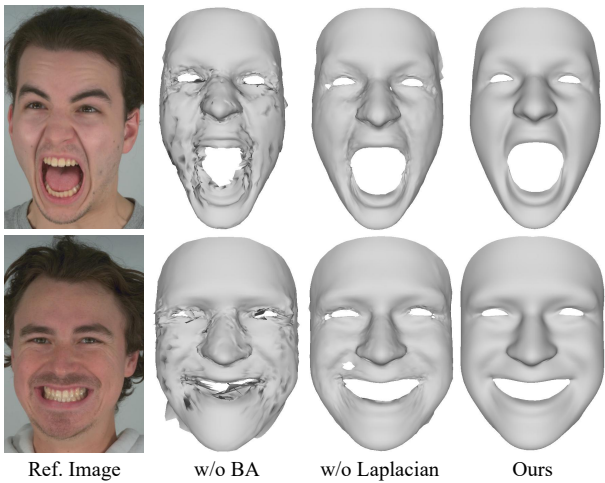


Figure 5: Qualitative evaluation of the key design choices in our method.

Pixel3DMM and HRN here, as these two methods perform better in the H3DS and NeRSemble benchmark. We capture a smartphone video around the subject and sample 16 frames as the inputs to our method. We run MetaShape on all the frames to reconstruct a mesh as the GT scan. We then align the mesh to compute the chamfer distance the same way as we do in building the NeRSemble benchmark. Again, our method obtains the best mean chamfer distance.

## Evaluations

We first evaluate the effectiveness of applying Bundle Adjustment to fuse the initial point cloud estimated by VGGT. For the *w/o BA* baseline, we directly connect the points in Eq. (5) according to the template topology  $\mathcal{F}$ . As shown in Figure 5, this naive fusion leads to noticeable surface artifacts, as the predictions of VGGT and Pixel3DMM are not




Figure 6: Limitations of our method.

fully accurate and the fused result is not supervised during their training. The quantitative results in Table 1 also verify the benefit of BA.

We then evaluate the role of the Laplacian term in the BA objective. The *w/o Laplacian* baseline applies the naive BA in Eq. (2). Although it improves upon *w/o BA*, the reconstructed surfaces remain noisy due to imperfect correspondences, and some vertices disappear when fewer than two valid tracks are available. With the proposed TopBA, the Laplacian energy enforces local consistency and results in a clean and complete mesh. Again, Table 1 confirms the superiority of our full method.

## Limitations and Discussions

First, as a mesh-based method, our method has limitations in representing hair strands, as shown in Figure 6. Secondly, constrained by the inaccuracy of the tracks predicted by Pixel3DMM (visualized in the Ref. Tracks image in Figure 6), we observe inaccurate reconstruction results on subjects with eyes closed, as shown in Figure 6. Incorporating special losses for eye regions is our future work (Wang et al. 2024; Giebenhain et al. 2025). Besides, our method focuses on reconstructing a large-scale static facial geometry without texture. Predicting a displacement or normal map on top of it (Li et al. 2021) is an interesting direction, but beyond the scope of this paper. Combining existing works to build a complete avatar with blendshapes (Ming et al. 2024), and reflectance maps (Han, Lyu, and Xu 2024; Han et al. 2025) is an interesting future direction.

## Conclusion

We propose VGGTFace, a novel method for topologically consistent facial geometry reconstruction from in-the-wild multi-view face images captured by everyday users. The core idea is to apply the recently-proposed 3D foundation model, *i.e.* VGGT, for face reconstruction. As the prediction of VGGT does not contain topology information, we propose to augment it with pixel-aligned UV coordinate images predicted by Pixel3DMM. To fuse the VGGT point cloud from different views, we propose a novel Topology-Aware Bundle Adjustment strategy, where we construct a Laplacian energy to the Bundle Adjustment objective, making our method robust to the inaccurate tracks. Experiments demonstrate our method obtains state-of-the-art results on various benchmarks and strong generalization ability on in-the-wild data captured by everyday users.

## Acknowledgments

This work was supported by the National Key R&D Program of China (2023YFC3305600). This work was also supported by THUIBCS, Tsinghua University, and BLBCI, Beijing Municipal Education Commission.

## References

- Agarwal, S.; Mierle, K.; and Team, T. C. S. 2023. Ceres Solver.
- Alexander, O.; Fyffe, G.; Busch, J.; Yu, X.; Ichikari, R.; Jones, A.; Debevec, P.; Jimenez, J.; Danvoye, E.; Antoniazzi, B.; et al. 2013. Digital ira: Creating a real-time photoreal digital actor. In *ACM SIGGRAPH 2013 Posters*, 1–1.
- Alexander, O.; Rogers, M.; Lambeth, W.; Chiang, M.; and Debevec, P. 2009. The digital emily project: photoreal facial modeling and animation. In *Acm siggraph 2009 courses*, 1–15.
- Bai, Z.; Cui, Z.; Rahim, J. A.; Liu, X.; and Tan, P. 2020. Deep Facial Non-Rigid Multi-View Stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5850–5860.
- Blanz, V.; and Vetter, T. 1999. A Morphable Model for the Synthesis of 3D Faces. In Waggenspack, W. N., ed., *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1999, Los Angeles, CA, USA, August 8-13, 1999*, 187–194. ACM.
- Bolkart, T.; Li, T.; and Black, M. J. 2023. Instant Multi-View Head Capture through Learnable Registration. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 768–779.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*.
- Caselles, P.; Ramon, E.; Garcia, J.; Triginer, G.; and Moreno-Noguer, F. 2025. Implicit shape and appearance priors for few-shot full head reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Deng, Y.; Yang, J.; Xu, S.; Chen, D.; Jia, Y.; and Tong, X. 2019. Accurate 3D Face Reconstruction with Weakly-Supervised Learning: From Single Image to Image Set. In *IEEE Computer Vision and Pattern Recognition Workshops*.
- Dib, A.; Bharaj, G.; Ahn, J.; Thébault, C.; Gosselin, P.; Romeo, M.; and Chevallier, L. 2021. Practical face reconstruction via differentiable ray tracing. In *Computer Graphics Forum*, volume 40, 153–164. Wiley Online Library.
- Dib, A.; Hafemann, L. G.; Got, E.; Anderson, T.; Fadaeinejad, A.; Cruz, R. M.; and Carbonneau, M.-A. 2024. MoSAR: Monocular semi-supervised model for avatar reconstruction using differentiable shading. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1770–1780.
- Egger, B.; Smith, W. A.; Tewari, A.; Wuhler, S.; Zollhoefer, M.; Beeler, T.; Bernard, F.; Bolkart, T.; Kortylewski, A.; Romdhani, S.; et al. 2020. 3d morphable face models—past, present, and future. *ACM Transactions on Graphics (ToG)*, 39(5): 1–38.
- Feng, Y.; Feng, H.; Black, M. J.; and Bolkart, T. 2021. Learning an Animatable Detailed 3D Face Model from In-The-Wild Images. volume 40.
- Fyffe, G.; Nagano, K.; Huynh, L.; Saito, S.; Busch, J.; Jones, A.; Li, H.; and Debevec, P. E. 2017. Multi-View Stereo on Consistent Face Topology. *Comput. Graph. Forum*, 36(2): 295–309.
- Giebenhain, S.; Kirschstein, T.; Rünz, M.; Agapito, L.; and Nießner, M. 2025. Pixel3DMM: Versatile Screen-Space Priors for Single-Image 3D Face Reconstruction.
- Guo, L.; Zhu, H.; Lu, Y.; Wu, M.; and Cao, X. 2023. RA-FaRe: Learning Robust and Accurate Non-parametric 3D Face Reconstruction from Pseudo 2D&3D Pairs. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Han, Y.; Lyu, J.; Sheng, K.; Que, M.; Zhang, Q.; Xu, L.; and Xu, F. 2025. Facial Appearance Capture at Home with Patch-Level Reflectance Prior. In *SIGGRAPH*.
- Han, Y.; Lyu, J.; and Xu, F. 2024. High-Quality Facial Geometry and Appearance Capture at Home. In *CVPR*.
- Han, Y.; Wang, Z.; and Xu, F. 2023. Learning a 3D Morphable Face Reflectance Model from Low-cost Data. In *CVPR*.
- Hartley, R.; and Zisserman, A. 2003. *Multiple view geometry in computer vision*. Cambridge university press.
- Kim, H.; Jang, M.; Yoon, W.; Lee, J.; Na, D.; and Woo, S. 2024. Switchlight: Co-design of physics-driven architecture and pre-training framework for human portrait relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25096–25106.
- Kirschstein, T.; Qian, S.; Giebenhain, S.; Walter, T.; and Nießner, M. 2023. NeRsemble: Multi-View Radiance Field Reconstruction of Human Heads. *ACM Trans. Graph.*, 42(4).
- Kirschstein, T.; Romero, J.; Sevastopolsky, A.; Nießner, M.; and Saito, S. 2025. Avat3r: Large Animatable Gaussian Reconstruction Model for High-fidelity 3D Head Avatars. arXiv:2502.20220.
- Lei, B.; Ren, J.; Feng, M.; Cui, M.; and Xie, X. 2023. A Hierarchical Representation Network for Accurate and Detailed Face Reconstruction from In-The-Wild Images. arXiv:2302.14434.
- Li, J.; Kang, D.; and He, Z. 2024. GRAPE: Generalizable and Robust Multi-view Facial Capture. In *European Conference on Computer Vision*, 403–418. Springer.
- Li, T.; Bolkart, T.; Black, M. J.; Li, H.; and Romero, J. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6).
- Li, T.; Liu, S.; Bolkart, T.; Liu, J.; Li, H.; and Zhao, Y. 2021. Topologically Consistent Multi-View Face Inference Using Volumetric Sampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3824–3834.
- Ming, X.; Li, J.; Ling, J.; Zhang, L.; and Xu, F. 2024. High-quality mesh blendshape generation from face videos via neural inverse rendering. In *European Conference on Computer Vision*, 106–125. Springer.

- Nicolet, B.; Jacobson, A.; and Jakob, W. 2021. Large Steps in Inverse Rendering of Geometry. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 40(6).
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; Howes, R.; Huang, P.-Y.; Xu, H.; Sharma, V.; Li, S.-W.; Galuba, W.; Rabbat, M.; Assran, M.; Ballas, N.; Synnaeve, G.; Misra, I.; Jegou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2023. DINOv2: Learning Robust Visual Features without Supervision. R., M. B.; Tewari, A.; Seidel, H.; Elgharib, M.; and Theobalt, C. 2021. Learning Complete 3D Morphable Face Models From Images and Videos. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 3361–3371. Computer Vision Foundation / IEEE.
- Ramon, E.; Triginer, G.; Escur, J.; Pumarola, A.; Garcia, J.; Giro-i Nieto, X.; and Moreno-Noguer, F. 2021. H3D-Net: Few-Shot High-Fidelity 3D Head Reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5620–5629.
- Ren, X.; Deng, J.; Cheng, Y.; Zhu, W.; Yan, Y.; Yang, X.; Zafeiriou, S.; and Ma, C. 2025. S<sup>3</sup>-Face: SSS-Compliant Facial Reflectance Estimation via Diffusion Priors. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 16051–16060.
- Riviere, J.; Gotardo, P. F.; Bradley, D.; Ghosh, A.; and Beeler, T. 2020. Single-shot high-quality facial geometry and skin appearance capture. *ACM Trans. Graph.*, 39(4): 81.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752.
- Schönberger, J. L.; and Frahm, J.-M. 2016. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Schönberger, J. L.; Zheng, E.; Pollefeys, M.; and Frahm, J.-M. 2016. Pixelwise View Selection for Unstructured Multi-View Stereo. In *European Conference on Computer Vision (ECCV)*.
- Sela, M.; Richardson, E.; and Kimmel, R. 2017. Unrestricted Facial Geometry Reconstruction Using Image-to-Image Translation. *arxiv*.
- Tewari, A.; Bernard, F.; Garrido, P.; Bharaj, G.; Elgharib, M.; Seidel, H.-P.; Pérez, P.; Zollhofer, M.; and Theobalt, C. 2019. Fml: Face model learning from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10812–10822.
- Tewari, A.; Zollhöfer, M.; Garrido, P.; Bernard, F.; Kim, H.; Pérez, P.; and Theobalt, C. 2018. Self-Supervised Multi-Level Face Model Learning for Monocular Reconstruction at Over 250 Hz. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 2549–2559. Computer Vision Foundation / IEEE Computer Society.
- Thies, J.; Zollhofer, M.; Stamminger, M.; Theobalt, C.; and Nießner, M. 2016. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2387–2395.
- Triggs, B.; McLauchlan, P. F.; Hartley, R. I.; and Fitzgibbon, A. W. 1999. Bundle adjustment—a modern synthesis. In *International workshop on vision algorithms*, 298–372. Springer.
- Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Ruppel, C.; and Novotny, D. 2025. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5294–5306.
- Wang, Y.; Han, Q.; Habermann, M.; Daniilidis, K.; Theobalt, C.; and Liu, L. 2023. NeuS2: Fast Learning of Neural Implicit Surfaces for Multi-view Reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Wang, Z.; Zhu, X.; Zhang, T.; Wang, B.; and Lei, Z. 2024. 3D Face Reconstruction with the Geometric Guidance of Facial Part Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1672–1682.
- Wu, F.; Bao, L.; Chen, Y.; Ling, Y.; Song, Y.; Li, S.; Ngan, K. N.; and Liu, W. 2019. Mvf-net: Multi-view 3d face morphable model regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 959–968.
- Xiao, Y.; Zhu, H.; Yang, H.; Diao, Z.; Lu, X.; and Cao, X. 2022. Detailed Facial Geometry Recovery from Multi-view Images by Learning an Implicit Function. *CoRR*, abs/2201.01016.
- Yang, H.; Zhu, H.; Wang, Y.; Huang, M.; Shen, Q.; Yang, R.; and Cao, X. 2020. FaceScape: A Large-Scale High Quality 3D Face Dataset and Detailed Riggable 3D Face Prediction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhang, Z.; Ge, Y.; Chen, R.; Tai, Y.; Yan, Y.; Yang, J.; Wang, C.; Li, J.; and Huang, F. 2021. Learning To Aggregate and Personalize 3D Face From In-the-Wild Photo Collection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14214–14224.
- Zheng, Y.; Yang, H.; Zhang, T.; Bao, J.; Chen, D.; Huang, Y.; Yuan, L.; Chen, D.; Zeng, M.; and Wen, F. 2022. General facial representation learning in a visual-linguistic manner. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18697–18709.
- Zollhöfer, M.; Thies, J.; Garrido, P.; Bradley, D.; Beeler, T.; Pérez, P.; Stamminger, M.; Nießner, M.; and Theobalt, C. 2018. State of the art on monocular 3D face reconstruction, tracking, and applications. In *Computer graphics forum*, volume 37, 523–550. Wiley Online Library.