

AvatarMAV: Fast 3D Head Avatar Reconstruction Using Motion-Aware Neural Voxels

Yuelang Xu Tsinghua University Beijing, China xll20@mails.tsinghua.edu.cn Lizhen Wang Tsinghua University NNKosmos Technology Beijing, China wlz18@mails.tsinghua.edu.cn Xiaochen Zhao Tsinghua University NNKosmos Technology Beijing, China zhaoxc19@mails.tsinghua.edu.cn

Hongwen Zhang Tsinghua University Beijing, China zhanghongwen@mail.tsinghua.edu.cn Yebin Liu Tsinghua University Beijing, China liuyebin@mail.tsinghua.edu.cn



Figure 1: We propose AvatarMAV, a fast 3D head avatar reconstruction method. Given a monocular video, our method can recover photo-realistic head avatar in 5 minutes.

ABSTRACT

With NeRF widely used for facial reenactment, recent methods can recover photo-realistic 3D head avatar from just a monocular video. Unfortunately, the training process of the NeRF-based methods is quite time-consuming, as MLP used in the NeRF-based methods is inefficient and requires too many iterations to converge. To overcome this problem, we propose AvatarMAV, a fast 3D head avatar reconstruction method using Motion-Aware Neural Voxels. AvatarMAV is the first to model both the canonical appearance and the decoupled expression motion by neural voxels for head avatar. In particular, the motion-aware neural voxels is generated from the weighted concatenation of multiple 4D tensors. The 4D tensors semantically correspond one-to-one with 3DMM expression basis and share the same weights as 3DMM expression coefficients. Benefiting from our novel representation, the proposed Avatar-MAV can recover photo-realistic head avatars in just 5 minutes (implemented with pure PyTorch), which is significantly faster



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License

SIGGRAPH '23 Conference Proceedings, August 06–10, 2023, Los Angeles, CA, USA © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0159-7/23/08. https://doi.org/10.1145/3588432.3591567 than the state-of-the-art facial reenactment methods. Project page: https://www.liuyebin.com/avatarmav.

CCS CONCEPTS

• Computing methodologies \rightarrow Volumetric models.

KEYWORDS

Neural Radiance Field, Facial Reenactment

ACM Reference Format:

Yuelang Xu, Lizhen Wang, Xiaochen Zhao, Hongwen Zhang, and Yebin Liu. 2023. AvatarMAV: Fast 3D Head Avatar Reconstruction Using Motion-Aware Neural Voxels. In Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings), August 06–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3588432.3591567

1 INTRODUCTION

Facial reenactment and head avatar reconstruction from a monocular video have been research hotspots recently, which have a very broad application prospect in VR/AR, digital human, holographic communication, webcast, etc.

However, current methods cannot reconstruct 3D head avatars in minutes, which has become a main limitation of applications. Recent works [Gafni et al. 2021; Gao et al. 2022a; Grassal et al. 2022;

Zheng et al. 2022] can recover photo-realistic 3D head avatars using easily available data, such as a monocular video. These methods can be divided into two categories: geometry-based methods and NeRFbased methods. Geometry-based [Grassal et al. 2022; Khakhulin et al. 2022; Zheng et al. 2022] methods can obtain well-defined 3D face geometry. For example, Neural Head Avatar [Grassal et al. 2022] constructs the avatar model through non-rigid deformation of the mesh template but is limited by the topology of template itself. IMAvatar [Zheng et al. 2022] solves this problem by optimizing an implicit signed distance field upon a mesh template. These methods model the motion by linear blend skinning of template, which leads to slow training and rendering speed. NeRF-based methods [Gafni et al. 2021; Gao et al. 2022a; Guo et al. 2021; Liu et al. 2022] achieve photo-realistic and view-consistent rendering and are not limited by the topology and coarse expressiveness of face templates. NeR-Face [Gafni et al. 2021] proposes to use an expression conditioned dynamic NeRF to model a head avatar and generates photo-realistic portrait images. However, training a NeRF model often takes hours or even days. The concurrent work on NeRFBlendShape [Gao et al. 2022a] reduces the training time consumption to 20 minutes by introducing the multi-level voxel field representation with multiresolution hash tables storing features. However, the coupling of motion and appearance still limits their efficiency.

On the other hand, a number of recent approaches [Müller et al. 2022; Sara Fridovich-Keil and Alex Yu et al. 2022; Sun et al. 2022a] propose to use explicit voxel data structures to represent a static NeRF scene, and their experiments prove that this representation plays a dramatic role in accelerating NeRF training. Meanwhile, some other methods extend explicit voxel representations to dynamic NeRF. The original D-NeRF [Pumarola et al. 2021] represents dynamic scenes by decoupling deformation field and canonical appearance. TineuVox [Fang et al. 2022] further accelerates D-NeRF by introducing a voxel grid to model the canonical component. But the deformation field of TineuVox is still built by a deep MLP, which still takes much time to converge. Based on these methods, it is still not trivial to accelerate dynamic NeRF training in head avatar reconstruction. Specifically, as the human face motions are more complex, a deep MLP to model the complex motions is usually inevitable. As a result, it takes much longer time for the MLP to converge during model training.

To overcome the above challenges, we propose AvatarMAV, a fast 3D head avatar reconstruction method using Motion-Aware Neural Voxels. Our approach can achieve 5-minute 3D head avatar reconstruction which means the training can be completed immediately just after collecting the training data. The key idea consists of two points: First, inspired by previous dynamic NeRF reconstruction approaches[Athar et al. 2023; Park et al. 2021a,b], we decouple the complex expression-related motion from the canonical appearance in NeRF-based head avatar reconstruction. Second, We propose an efficient voxel-based representation instead of deep MLPs, for both the appearance and the motion field. Note that NeRFBlendshape [Gao et al. 2022a] only introduces the voxel-based representation but does not decouple the motion and appearance. We utilize the prior information provided by 3DMM expression basis to model the expression-related motion using only voxel grids and a tiny MLP. Specifically, we use an expression-conditioned

neural voxel grid to describe the motion field, and further decompose this neural voxel grid as a linearly weighted concatenation of multiple neural voxel grids basis. The number of the voxel grids basis is the same as the dimension of 3DMM expression basis and the weights of the concatenation is exactly the 3DMM expression coefficients. We define this representation as motion-aware neural voxels. For the canonical appearance, we simply use a single neural voxel grid. Overall, our method can reconstruct a photo-realistic 3D head avatar in 5 minutes implemented by pure Pytorch code. Both quantitative and qualitative experiments demonstrate the superiority of our method in training speed while preserving a comparable rendering result.

2 RELATED WORKS

Portrait Video Synthesis. In the past period of time, a large number of portrait video synthesis methods have been proposed. The earliest approaches [Li et al. 2012; Thies et al. 2015, 2016; Vlasic et al. 2005; Weise et al. 2011] track to reconstruct textured face mesh with expressions and then re-render it to synthesize portrait images with desired expressions. Areas such as mouth, hair and background need to be blended later. Such methods heavily rely on accurate tracking methods and high-quality reconstructed mesh models etc, which severely limits the broad applications. Warping-based methods [Averbuch-Elor et al. 2017; Geng et al. 2018; Nirkin et al. 2019; Siarohin et al. 2019; Wiles et al. 2018; Yin et al. 2022] model the 2D motion flow map as a common representation for expression transfer, but fail to deal with extreme head poses or expressions. Some recent approaches [Drobyshev et al. 2022; Wang et al. 2021] lift the 2D motion flow map into 3D space to overcome such artifacts to a certain extent. Based on the widely-used facial parametric models [Blanz and Vetter 1999; Gerig et al. 2017; Li et al. 2017; Wang et al. 2022b], template-based methods [Buehler et al. 2021; Doukas et al. 2020, 2021; Kim et al. 2018; Koujan et al. 2020; Thies et al. 2019; Wang et al. 2023] that combine face template and neural rendering has gradually become the mainstream. These methods first render coarse images using face templates with controllable expression and pose coefficients, and then utilize convolutional neural networks to generate mouths, hair and rich details to synthesize high-quality portrait images. Other similar methods [Chen et al. 2020; Zakharov et al. 2019] use semantic maps or landmarks for coarse-level representation to replace face models.

Monocular 3D Head Avatar Reconstruction. It has always been a challenging task to reconstruct 3D full head avatar from monocular video. Early methods [Cao et al. 2015, 2016; Hu et al. 2017; Ichim et al. 2015; Nagano et al. 2018] use blendshape-based templates to fit portraits in input videos to model human heads. For parts such as eyeball and mouth interior that are difficult to model by the templates, these methods require additional post-processing or leverage neural networks to learn geometry or textures [Grassal et al. 2022; Khakhulin et al. 2022]. Inspired by the 3D scene reconstruction method IDR [Yariv et al. 2020] based on implicit representation [Park et al. 2019], IMAvatar [Zheng et al. 2022] proposes to optimize an occupancy field and a color field to represent the head model based on the FLAME model. As NeRF [Mildenhall et al. 2020] representation shows strong ability to synthesis high-fidelity photo-realistic images, Recent methods [Park et al. 2021a,b; Shao



Figure 2: Overview. Given a portrait video, we first track the expression and head pose using a 3DMM template. After the pre-processing, given expression coefficients, we use motion voxel grid basis to represent motions caused by each expression basis and sum them weighted as an entire motion voxel grid. The entire motion voxel grid and the following 2-layer MLP will then transfer an input point x to $x + \delta x$ by adding all expression-related deformations. Finally, we will query point $x + \delta x$ in the appearance voxel grid and generate a final portrait image using volumetric rendering.

et al. 2023] propose to model a 4D scene by a deformable NeRF which can not be animated. As the solution, avatar reconstruction methods [Athar et al. 2023, 2022; Gafni et al. 2021; Zheng et al. 2023] use a face template as the condition to learn a controllable NeRF model. Furthermore, NeRFBlendShape [Gao et al. 2022a] use a voxel-based representation to replace the MLP, and model the dynamic NeRF by linear combination of multiple NeRF basis oneto-one corresponding to semantic blendshape coefficients. In the field of the audio-driven avatar, latest methods [Guo et al. 2021; Liu et al. 2022] also leverage dynamic NeRF as the representation of head avatars. Besides, some other methods [Chan et al. 2022; Hong et al. 2022; Sun et al. 2022b, 2023, 2022c; Wang et al. 2022a; Zhuang et al. 2022] are not limited to training person-specific avatars, but to train a general human head model as prior on large-scale datasets. NeRF representation is used in our method as well, but we make great progress in accelerating training speed by both using explicit representation and the decomposition of dynamic motion and static appearance.

Training Acceleration for NeRF. As vanilla NeRF[Mildenhall et al. 2020] usually takes hours or even days to complete the training of a static scene, a lot of works focus on speeding up NeRF training process. DVGO [Sun et al. 2022a] proposes to accelerate NeRF training by directly replacing most MLPs with voxel grids, which significantly reduces the time required for training convergence. Plenoxels [Sara Fridovich-Keil and Alex Yu et al. 2022] goes a step further and proposes to only use a sparse grid with density and spherical harmonic coefficients at each voxel without any neural network. TensoRF [Chen et al. 2022] proposes to decompose the voxel grid into sum of vector-matrix outer products. It reduces the size of the model while ensuring training efficiency and rendering quality. Instant-NGP [Müller et al. 2022] introduces multi-resolution

hash embedding for the voxel grid structure. Combined with customized CUDA implementation, they enable extremely fast NeRF training. These methods mainly focus on static scenes reconstruction. For dynamic scenes, TiNeuVox [Fang et al. 2022] proposes to replace the MLPs for canonical scene in D-NeRF [Pumarola et al. 2021] with an explicit voxel grid. Different from [Fang et al. 2022] which still represent the time-dependent deformation field with a MLP, our method propose to utilize the explicit data structure for both the static content and the expression-dependent deformation field.

3 METHOD

AvatarMAV can generate a fast 3D head avatar by decoupling complex expression motion from the canonical appearance and voxel-based representation. In this section, we will introduce our voxel-based representation and the training process.

3.1 Representation

3.1.1 Formulation. Generally, previous methods [Gafni et al. 2021; Gao et al. 2022a; Guo et al. 2021; Liu et al. 2022] formulate NeRF-based head avatar as a expression-dependent NeRF model:

$$(c,\sigma) = \Phi(x,d,\theta). \tag{1}$$

Given a query point *x* with view direction *d* and expression coefficients θ , the color and the density denoted by *c* and σ respectively are computed for volumetric rendering [Mildenhall et al. 2020].

However, their practice of mixing the dynamics expression motion with the appearance of the human head brings obstacles to the fast convergence of avatar learning. In contrast, our method decouples the NeRF-based head avatar into the canonical appearance and expression motion based on the physical assumption that a face

geometry with expressions can be deformed from the neutral face geometry. Specifically, we define $\Phi(\cdot)$ as expression-independent NeRF and additionally introduce an expression-dependent deformation $\Omega(\cdot)$, which can be formulated as:

$$(c,\sigma) = \Phi(x + \delta x, d), \tag{2}$$

with
$$\delta x = \Omega(x, \theta)$$
. (3)

Note that $\Phi(\cdot)$ represents a static NeRF and does not contain the expression θ in its parameters. Since the degrees of freedom of the parameters are greatly reduced, the training speed can be significantly improved as shown in our experiments.

To further improve training efficiency, we enhance our representation with an explicit voxel-based strategy. However, using a single voxel grid is infeasible to model the dynamic expression motion. Previous methods [Fang et al. 2022; Pumarola et al. 2021] utilize deep MLPs to alleviate this issue but lead to a quite time-consuming convergence process. In our solution, we take the expression prior from 3DMM and use its expression coefficients to condition multiple voxel grids, which is the key to efficient and effective avatar creation.

3.1.2 Expression Motion. In the traditional PCA-based 3DMM [Blanz and Vetter 1999; Gerig et al. 2017], the face model can be deformed through a linear combination of its expression PCA basis. To leverage the expression prior of 3DMM and the expressive power of neural voxels, we use the expression coefficients as the weights to concatenate multiple Motion Voxel Grid (MVG) basis into an entire MVG to represent the expression motion. Specifically, each dimension of MVG basis uses the same expression coefficient $\theta \in \mathbb{R}^N$ as the corresponding 3DMM expression basis, where N denotes the dimension of 3DMM expression basis. From another point of view, MVG basis can be considered as "neural" expression basis for head motion, which is able to present more detailed motion in a wider range (including the ears and hair regions). The process can be formulated as:

$$V_d(\theta) = \theta^1 V_d^1 \oplus \theta^2 V_d^2 \oplus \dots \oplus \theta^N V_d^N, \tag{4}$$

where $\{V_d^1, V_d^2, \cdots, V_d^N\} \in \mathbb{R}^{N \times C_d \times L_d \times L_d}$ denotes the MVG basis in the latent space. $V_d(\theta) \in \mathbb{R}^{NC_d \times L_d \times L_d \times L_d}$ denotes the MVG and $\theta = \{\theta^1, \theta^2, \cdots, \theta^N\} \in \mathbb{R}^N$ denotes the expression coefficients derived from 3DMM. C_d denotes the channel number of voxel features in each voxel grid base.

Finally, for a query point x, we adopt multi-distance interpolation [Fang et al. 2022] to sample the corresponding feature vector v_d and feed it into a 2-layer MLP f_d to predict the motion offsets of input points. This process can be formulated as:

$$v_d = T_K(x, V_d(\theta)) \tag{5}$$

$$\delta x = f_d(v_d),\tag{6}$$

where $T_K(\cdot)$ denotes K-distance interpolation [Fang et al. 2022] and δx denotes the motion offset of a query point x. We have described in detail how we formulate the expression-dependent deformation $\Omega(\cdot)$ in Eq.3 using Motion-Aware Neural Voxels representation.

Figure 3: We visualize our learned canonical appearance and expression motion. The color value reflecting the 3D vector of flow

3.1.3 Canonical Appearance. Since we model the canonical appearance $\Phi(\cdot)$ as a static NeRF, we just directly represent the canonical appearance field with a single appearance voxel grid $V_a \in \mathbb{R}^{C_a \times L_a \times L_a}$ and a two-layer MLP F_a , where C_a denotes the channel number of voxel features and L_a denotes the resolution of the appearance voxel grid. The explicit appearance voxel grid V_a contains information for appearance in its per-voxel features, while the two-layer MLP F_a translates the sampled features to the color and the density.

Specifically, for each canonical query point x, its corresponding feature vector v_a is first sampled from the appearance voxel grid V_a by multi-distance interpolation [Fang et al. 2022]. Then, we feed this feature vector together along with the view direction d into the 2-layer MLP F_a to predict the color and the density:

$$(c(x), \sigma(x)) = F_a(\gamma(v_a), \gamma(d), \theta), \tag{7}$$

where $c(\cdot)$ and $\sigma(\cdot)$ denote the RGB value and the density value of query point *x* respectively, $\gamma(\cdot)$ is the positional encoding [Mildenhall et al. 2020] mapping the feature vector v_a and the view direction *d* to their periodic formulation. As it is difficult to model dynamic details and topology changes of the human face such as wrinkles given only motion information, we additionally feed the expression coefficients θ to F_a to describe the slightly dynamic appearance.

3.2 Training

For more effective training, we first remove the background [Lin et al. 2022], the neck, and the body part of the human body [zllrunning 2019] from each video frame during the data preprocessing phase. Then we detect face landmarks [emilianavt 2020] and obtain expression coefficients and head poses by tracking 3DMM models on each frame. During training, we optimize the voxel grids and MLPs with direct photometric supervision. Moreover, we empirically find that the motion-aware neural voxels may also learn to model the information of canonical appearances such as global constant offset and non-zero offsets in static areas without regularization. Hence, we add a regularization term to punish all non-zero offsets of sampled points. Meanwhile, as all the offsets in the motion voxel grid are pushed to zeros, the canonical appearance tends to present a neutral expression. The total loss function can be formulated as:

$$\mathcal{L} = \sum_{r \in \mathcal{R}} ||I(r) - I_{gt}(r)||_1 + \lambda \sum_{r \in \mathcal{R}} \sum_{t \in \mathcal{T}(r)} ||\delta(r(t))||_2, \qquad (8)$$

where \mathcal{R} denotes the sampled rays in a batch. I_{gt} denotes the preprocessed ground-truth image, $\mathcal{T}(r)$ denotes the distances set of the



AvatarMAV: Fast 3D Head Avatar Reconstruction Using Motion-Aware Neural Voxels

sampled points on rays r, and λ denotes the weight of the regular term.

4 EXPERIMENTS

4.1 Implementation Details

We implement our whole framework with pure PyTorch. For expression coefficients, we directly use the first 32 expression basis of Basel Face Model [Blanz and Vetter 1999; Gerig et al. 2017]. For appearance voxel grid of canonical appearance, we set the channel number of feature as 4 and resolution as 64^3 . For expression motion, the number of MVG basis is set as 32 (equal to the number of expression coefficients). The channel number of the features is set as 2 and the resolution of each grid is set as 16^3 . All the MLPs in our framework are 2-layer with 64 neurons for each hidden layer. We uniformly set frequency number as 4 for positional encoding. For multi-distance interpolation, we set *K* as 3. Under this hyperparameter setting, an average of 140ms to render a 256×256 resolution image is required on one RTX 3090 GPU.

During optimization, we use Adam optimizer. The initial learning rate is set as 1×10^{-2} for MVG basis and appearance voxel grid, and 1×10^{-3} for all the tiny MLPs. At the 500 and 2000 iterations, we reduce the learning rate by a third. For the loss function, we set $\lambda = 0.01$. For ray sampling, 4096 rays and 64 points are sampled along each ray in each iteration. The batch size is set as 1. Thus, the time consumption of one iteration is 34ms. We train a model for 10000 iterations in total. For the first 6000 iterations, we use training images with 256 × 256 resolution and for the last 4000 iterations, we use 512 × 512 resolution.

We collected 9 training videos for our experiments. with 4 of them are from HDTF dataset [Zhang et al. 2021], 1 from NeRF-BlendShape [Gao et al. 2022a]. We additionally collect 4 videos by a hand-hold mobile phone. Each video contains about 2000-4000 frames and we reserved the last 15% of each video solely for evaluation. Meanwhile, we collect more several videos for cross-identity reenactment task and more visualization results.

4.2 Comparisons on Render Quality

We conduct qualitative and quantitative comparisons on render qualities between our AvatarMAV and three state-of-the-art methods DVP (Deep Video Portraits) [Kim et al. 2018], IMAvatar [Zheng et al. 2022], and NeRFace [Gafni et al. 2021] on both self reenactment task and cross-identity reenactment task. DVP does not reconstruct the full head model, but directly synthesizes 2D images through an image-to-image translation network with a controllable coarse 3DMM model as guidance. IMAvatar reconstructs an implicit signed distance field based on the FLAME model. NeRFace reconstructs a NeRF head model with 3DMM expression coefficients as condition. In our experiments, we spend enough time for training to ensure complete convergence for all the methods. Respectively, 1 day for IMAvatar [Zheng et al. 2022], 1 day for DVP [Kim et al. 2018], 12 hours for NeRFace [Gafni et al. 2021], and 5 minutes for Avatar-MAV. Since the concurrent work on NeRFBlendShape [Gao et al. 2022a] doesn't release their training code, we only make qualitative comparisons on training speed to their demo video in the next section.

SIGGRAPH '23 Conference Proceedings, August 06-10, 2023, Los Angeles, CA, USA

Qualitative comparisons on self reenactment task are shown in Fig. 4. The results validate that AvatarMAV achieves the highest render quality while the training time is far less than the other methods. IMAvatar reconstructs an implicit model based on the mesh of a FLAME template, yet the expressiveness is limited. DVP inherits the GAN framework to generate images. But in many cases, the generated details are not appropriate. The performance of NeRFace is comparable to AvatarMAV, but the training time is much longer.

Table 1: Quantitative evaluation results of AvatarMAV and other three state-of-the-art methods on self reenactment task. We calculate mean values on three different subjects.

Method	MSE ↓	PSNR ↑	SSIM	LPIPS
DVP [Kim et al. 2018]	0.0047	24.2	0.89	0.072
IMavatar [Zheng et al. 2022]	0.0041	24.6	0.92	0.131
NeRFace [Gafni et al. 2021]	0.0015	30.2	0.96	0.038
AvatarMAV	0.0014	30.4	0.96	0.038

Table. 1 shows the quantitative evaluation results. We evaluate on four metrics: Mean Squared Error (MSE), Peak Signal-to-Noise Ratio (PSNR), Structure Similarity Index (SSIM) and Learned Perceptual Image Patch Similarity (LPIPS). Our AvatarMAV achieves the best results in MSE and PSNR metrics, and comparable results to NeRFace on both SSIM and LPIPS metrics.

Next, we compare AvatarMAV with these three state-of-the-art methods on cross-identity reenactment task. Qualitative results are shown in Fig. 5. In cases where the expression from the source video is out of the distribution in the training data, NeRFace might produce floating artifacts. Benefiting from our decoupling of the dynamic and the static elements, AvatarMAV maintains better stability.

4.3 Comparisons on Training Speed

We qualitatively compare the training speed of our AvatarMAV and two other NeRF-based methods: NeRFBlendShape [Gao et al. 2022a] and NeRFace [Gafni et al. 2021]. We train each model from scratch, and render the corresponding image at 5s, 15s, 30s, 1min, 2min. Since our model has almost reached complete convergence within the first 2 minutes. We do not visualize the render results after 2 minutes. NeRFBlendshape claims that their method takes 20 minutes to converge. In our experiments, we found that NeRFace actually only takes a few hours to converge. Note that, We use the subject in the demo video of NeRFBlendShape[Gao et al. 2022b] for comparison.

Qualitatively comparisons are shown in Fig. 7. Our model converges with very high efficiency. In the first 30 seconds of training, our model can complete most of the convergence after very few iterations. At the time of 2 minutes, our model has almost converged (5 minutes for the complete convergence). In contrast, much more training time is required for the other two methods.

Yuelang Xu, et al.



Figure 4: Qualitative comparisons between AvatarMAV and other three state-of-the-art methods on self reenactment task. From left to right: IMAvatar [Zheng et al. 2022], DVP [Kim et al. 2018], NeRFace [Gafni et al. 2021], AvatarMAV and Ground Truth. Our approach is able to converge and learn details within a short time. YouTube video IDs are listed: tUQlOFFCeOY for Chirs Van Hollen; tvA_jbuL0Jo for Barack Obama.

5 ABLATION STUDY

5.1 Complete Voxel-based Representation

To accelerate training, our idea consists of two key points in improving the representation of NeRF-based head avatar: decoupling expression motion from canonical appearance; modeling the expression motion by motion-aware neural voxels representation. In the following, we ablate these two components.

The first baseline no longer decouples expression motion and canonical appearance (Without Decoupling). Given N dimension expression coefficients, we also establish N corresponding voxel grids basis but the weighted concatenation of the voxel grids directly represent the final expressive head avatar. A 2-layer percepton is followed to predict the color and the density values. In our expriments, we set the resolution of the voxel grid basis as 64^3 , which is equal to resolution of the appearance voxel grid in our full method.

In the second baseline, we directly use a MLP to implicitly model the expression motion (MLP Deformation). Specifically, for a query point, we feed the coordinate concatenated with the expression coefficients in to a MLP with positional encoding to obtain the offset. In our expriments, we use a 4-layer percepton with 128 neurons for each hidden layer. For the canonical appearance, we still use a neural voxel grid with a 2-layer percepton. In this section, we evaluate our full method AvatarMAV and the two variant baselines above. Qualitative and quantitative results are shown in Fig. 7 and Table. 2. As the without decoupling baseline does not decouple the expression motion from the canonical appearance, the rendered images tend to be blurred. the MLP deformation baseline benefits from decoupling and converges faster. However, it can be observed from the images at 1min or 2min in Fig. 7 that only voxel-based canonical appearance converges to the average state rapidly, while MLP-based motion converges slowly. Our Avatar-MAV combines the advantages of both decoupling and voxel-based representation to achieve the best training efficiency. We also mark the corresponding quantitative evaluation results (PSNR) of the two frames, which also verified the above conclusions.

5.2 Latent Voxel

An alternative idea is to directly store offset value δx rather than feature vector in the described MVG basis in Sec. 3.1, such that the tiny MLP f_d followed can be removed to achieve higher efficiency. We call the architecture direct motion voxels to distinguish it from the neural motion voxels. Meanwhile, the MVG is calculated by weighted blending of MVG basis to match the subsequent process. We evaluate this variant baseline qualitatively and quantitatively as shown in Fig. 8 and Table. 2. As a result, the direct motion voxels AvatarMAV: Fast 3D Head Avatar Reconstruction Using Motion-Aware Neural Voxels

SIGGRAPH '23 Conference Proceedings, August 06-10, 2023, Los Angeles, CA, USA



Figure 5: Qualitative results of AvatarMAV and three other state-of-the-art methods on cross-identity reenactment task. From left to right: DVP [Kim et al. 2018], IMAvatar [Zheng et al. 2022], NeRFace [Gafni et al. 2021] and AvatarMAV. YouTube video IDs are listed: uBWJJvynXpA for Chris Talor; 5rTk831pQFM for Rand Paul; hSvTbV4wW8A for Ben Cardin; jXkU9PYwQSE for Sam Brownback.





Figure 6: We make qualitative comparisons on training speed among NeRFace [Gafni et al. 2021], NeRFBlendShape [Gao et al. 2022a] and our AvatarMAV. Our model converges rapidly within the first 2 minutes.

Figure 7: Qualitative results of two baselines and our Avatar-MAV on self reenactment task at different points in time. Our full method AvatarMAV can learn both high-quality appearance and expression motion faster.

leads to a decrease in the performance. To explain, the motion field in AvatarMAV is a backward warp field that maps the points in the live space to the canonical space. Such a backward warp field

is ambiguous as there exist non-linear many-to-one mappings. To

Table 2: Quantitative evaluation results of all the ablation baselines and AvatarMAV on self reenactment task at 30 seconds and 5 minutes. We calculate mean values on three different subjects.

Method (30s)	MSE	PSNR	SSIM	LPIPS
Wo decoupling	0.0032	26.0	0.926	0.073
MLP deformation	0.0037	25.4	0.919	0.070
Direct motion voxels	0.0030	26.4	0.931	0.061
AvatarMAV	0.0025	27.1	0.938	0.062
Method (5min)	MSE	PSNR	SSIM	LPIPS
Wo decoupling	0.0021	28.0	0.946	0.062
MLP deformation	0.0028	26.7	0.933	0.065
Direct motion voxels	0.0023	27.9	0.945	0.053
AvatarMAV	0.0020	28.5	0.950	0.052

handle this issue, an MLP is necessary to achieve such non-linear mapping.



Figure 8: Qualitative results of direct motion voxels and AvatarMAV on self reenactment task (5min). AvatarMAV can learn finer expressions.

6 LIMITATION

Although we accelerate the training speed of NeRF-based head avatar to almost instant, there is no significant improvement in rendering quality. As the camera poses and expression coefficients obtained by fitting a face template are not accurate enough. Artifacts and blurriness appear when the rendering viewpoint is away from the front view or the expression is exaggerated as shown in Fig. 9. On the other hand, limited by the face template, our approach can only handle human heads, without the ability to reconstruct other regions such as neck and upper body. In the future, we will try to use more general parameterized representation to solve this problem.



Figure 9: Failure cases when the rendering viewpoint is away from the front view (left) or the expression is exaggerated (right).

CONCLUSION

In this paper, we have presented AvatarMAV, a fast 3D head avatar reconstruction method using motion-aware neural voxels. The proposed neural-voxel-based representation of both the canonical appearance and the motion field is able to greatly accelerate the training process. Furthermore, we have also demonstrated superiority and efficiency of the proposed motion-aware neural voxels in reconstructing complex expression-related motion. In future work, on-line creating a 3D head avatar during live capturing process might become feasible, which can greatly reduce the cost of generating digital human faces. We believe AvatarMAV will inspire the following facial reenactment researches, and the proposed motion-aware neural voxels could further broaden the applications of dynamic NeRF.

ACKNOWLEDGMENTS

This paper is supported by National Key R&D Program of China (2022YFF0902200), the NSFC project No.62125107 and No.61827805.

REFERENCES

7

- ShahRukh Athar, Zhixin Shu, and Dimitris Samaras. 2023. Flame-in-nerf: Neural control of radiance fields for free view face animation. In IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG). 1–8.
- ShahRukh Athar, Zexiang Xu, Kalyan Sunkavalli, Eli Shechtman, and Zhixin Shu. 2022. RigNeRF: Fully Controllable Neural 3D Portraits. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Hadar Averbuch-Elor, Daniel Cohen-Or, Johannes Kopf, and Michael F. Cohen. 2017. Bringing Portraits to Life. ACM Trans. Graph. 36, 6, Article 196 (nov 2017), 13 pages.
- Volker Blanz and Thomas Vetter. 1999. A Morphable Model for the Synthesis of 3D Faces. In Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99). ACM Press/Addison-Wesley Publishing Co., USA, 187–194.
- Marcel C. Buehler, Abhimitra Meka, Gengyan Li, Thabo Beeler, and Otmar Hilliges. 2021. VariTex: Variational Neural Face Textures. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV).
- Chen Cao, Derek Bradley, Kun Zhou, and Thabo Beeler. 2015. Real-Time High-Fidelity Facial Performance Capture. ACM Trans. Graph. 34, 4, Article 46 (jul 2015), 9 pages.
- Chen Cao, Hongzhi Wu, Yanlin Weng, Tianjia Shao, and Kun Zhou. 2016. Real-Time Facial Animation with Image-Based Dynamic Avatars. ACM Trans. Graph. 35, 4, Article 126 (jul 2016), 12 pages.
- Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. 2022. Efficient Geometry-aware 3D Generative Adversarial Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 16102–16112.
- Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. 2022. TensoRF: Tensorial Radiance Fields. In Proceedings of the European Conference on Computer Vision (ECCV).
- Zhuo Chen, Chaoyue Wang, Bo Yuan, and Dacheng Tao. 2020. PuppeteerGAN: Arbitrary Portrait Animation With Semantic-Aware Appearance Transformation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 13515–13524.
- Michail Christos Doukas, Mohammad Rami Koujan, Viktoriia Sharmanska, Anastasios Roussos, and Stefanos Zafeiriou. 2020. Head2Head++: Deep Facial Attributes Re-Targeting. IEEE Transactions on Biometrics, Behavior, and Identity Science 3 (2020), 31–43.
- Michail Christos Doukas, Stefanos Zafeiriou, and Viktoriia Sharmanska. 2021. HeadGAN: One-shot Neural Head Synthesis and Editing. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV).
- Nikita Drobyshev, Jenya Chelishev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. 2022. MegaPortraits: One-shot Megapixel Neural Head Avatars. Proceedings of the 30th ACM International Conference on Multimedia. emilianavt. 2020. OpenSeeFace. https://github.com/emilianavt/OpenSeeFace.
- Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. 2022. Fast Dynamic Radiance Fields with Time-Aware Neural Voxels. In SIGGRAPH Asia 2022 Conference Papers.
- Guy Gafni, Justus Thies, Michael Zollhofer, and Matthias Niessner. 2021. Dynamic Neural Radiance Fields for Monocular 4D Facial Avatar Reconstruction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 8645–8654.

AvatarMAV: Fast 3D Head Avatar Reconstruction Using Motion-Aware Neural Voxels

SIGGRAPH '23 Conference Proceedings, August 06-10, 2023, Los Angeles, CA, USA

- Xuan Gao, Chenglai Zhong, Jun Xiang, Yang Hong, Yudong Guo, and Juyong Zhang. 2022a. Reconstructing Personalized Semantic Facial NeRF Models From Monocular Video. ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 41, 6 (2022).
- Xuan Gao, Chenglai Zhong, Jun Xiang, Yang Hong, Yudong Guo, and Juyong Zhang. 2022b. Reconstructing Personalized Semantic Facial NeRF Models From Monocular Video. https://ustc3dv.github.io/NeRFBlendShape.
- Jiahao Geng, Tianjia Shao, Youyi Zheng, Yanlin Weng, and Kun Zhou. 2018. Warp-Guided GANs for Single-Photo Facial Animation. ACM Trans. Graph. 37, 6, Article 231 (dec 2018), 12 pages.
- Thomas Gerig, Andreas Forster, Clemens Blumer, Bernhard Egger, Marcel Lüthi, Sandro Schönborn, and Thomas Vetter. 2017. Morphable Face Models - An Open Framework. 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018), 75–82.
- Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. 2022. Neural Head Avatars From Monocular RGB Videos. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 18632–18643.
- Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. 2021. AD-NeRF: Audio Driven Neural Radiance Fields for Talking Head Synthesis. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 5764–5774.
- Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. 2022. HeadNeRF: A Real-Time NeRF-Based Parametric Head Model. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 20374–20384.
- Liwen Hu, Shunsuke Saito, Lingyu Wei, Koki Nagano, Jaewoo Seo, Jens Fursund, Iman Sadeghi, Carrie Sun, Yen-Chun Chen, and Hao Li. 2017. Avatar Digitization from a Single Image for Real-Time Rendering. ACM Trans. Graph. 36, 6, Article 195 (nov 2017), 14 pages.
- Alexandru Eugen Ichim, Sofien Bouaziz, and Mark Pauly. 2015. Dynamic 3D Avatar Creation from Hand-Held Video Input. ACM Trans. Graph. 34, 4, Article 45 (jul 2015), 14 pages.
- Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. 2022. Realistic One-shot Mesh-based Head Avatars. In Proceedings of the European Conference on Computer Vision (ECCV).
- Hyeongwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. 2018. Deep Video Portraits. ACM Trans. Graph. 37, 4, Article 163 (jul 2018), 14 pages.
- Mohammad Rami Koujan, Michail Christos Doukas, Anastasios Roussos, and Stefanos Zafeiriou. 2020. Head2Head: Video-based Neural Head Synthesis. 2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020), 16–23.
- Kai Li, Feng Xu, Jue Wang, Qionghai Dai, and Yebin Liu. 2012. A data-driven approach for facial expression synthesis in video. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition (CVPR). 299–310.
- Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. 2017. Learning a Model of Facial Shape and Expression from 4D Scans. *ACM Trans. Graph.* 36, 6, Article 194 (nov 2017), 17 pages.
- Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. 2022. Robust High-Resolution Video Matting With Temporal Guidance. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). 3132–3141.
- Xian Liu, Yinghao Xu, Qianyi Wu, Hang Zhou, Wayne Wu, and Bolei Zhou. 2022. Semantic-Aware Implicit Neural Audio-Driven Video Portrait Generation. In Proceedings of the European Conference on Computer Vision (ECCV).
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In Proceedings of the European Conference on Computer Vision (ECCV).
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.* 41, 4, Article 102 (July 2022), 15 pages.
- Koki Nagano, Jaewoo Seo, Jun Xing, Lingyu Wei, Zimo Li, Shunsuke Saito, Aviral Agarwal, Jens Fursund, and Hao Li. 2018. PaGAN: Real-Time Avatars Using Dynamic Textures. ACM Trans. Graph. 37, 6, Article 258 (dec 2018), 12 pages.
- Yuval Nirkin, Yosi Keller, and Tal Hassner. 2019. FSGAN: Subject Agnostic Face Swapping and Reenactment. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 7183–7192.
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 165–174.
- Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. 2021a. Nerfies: Deformable neural radiance fields. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 5845–5854.
- Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. 2021b. HyperNeRF: A Higher-Dimensional Representation for Topologically Varying Neural Radiance Fields. ACM Trans. Graph. 40, 6, Article 238 (dec 2021).

- Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. 2021. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 10313–10322.
- Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2022. Plenoxels: Radiance Fields without Neural Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. 2023. Tensor4D: Efficient Neural 4D Decomposition for High-fidelity Dynamic Reconstruction and Rendering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. 2019. First Order Motion Model for Image Animation. In Conference on Neural Information Processing Systems (NeurIPS).
- Cheng Sun, Min Sun, and Hwann-Tzong Chen. 2022a. Direct Voxel Grid Optimization: Super-fast Convergence for Radiance Fields Reconstruction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. 2022b. IDE-3D: Interactive Disentangled Editing for High-Resolution 3D-aware Portrait Synthesis. ACM Transactions on Graphics (TOG) 41, 6, Article 270 (2022), 10 pages.
- Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. 2023. Next3D: Generative Neural Texture Rasterization for 3D-Aware Head Avatars. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. 2022c. FENeRF: Face Editing in Neural Radiance Fields. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 7662–7672.
- Justus Thies, Michael Zollhöfer, and Matthias Nießner. 2019. Deferred Neural Rendering: Image Synthesis Using Neural Textures. ACM Trans. Graph. 38, 4, Article 66 (jul 2019), 12 pages.
- Justus Thies, Michael Zollhöfer, Matthias Nießner, Levi Valgaerts, Marc Stamminger, and Christian Theobalt. 2015. Real-Time Expression Transfer for Facial Reenactment. ACM Trans. Graph. 34, 6, Article 183 (oct 2015), 14 pages.
- Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Niessner. 2016. Face2Face: Real-Time Face Capture and Reenactment of RGB Videos. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2387–2395.
- Daniel Vlasic, Matthew Brand, Hanspeter Pfister, and Jovan Popović. 2005. Face Transfer with Multilinear Models. ACM Trans. Graph. 24, 3 (jul 2005), 426–433.
- Daoye Wang, Prashanth Chandran, Gaspard Zoss, Derek Bradley, and Paulo Gotardo. 2022a. MoRF: Morphable Radiance Fields for Multiview Neural Head Modeling. In ACM SIGGRAPH 2022 Conference Proceedings (Vancouver, BC, Canada) (SIGGRAPH '22). Association for Computing Machinery, New York, NY, USA, Article 55, 9 pages.
- Lizhen Wang, Zhiyuan Chen, Tao Yu, Chenguang Ma, Liang Li, and Yebin Liu. 2022b. FaceVerse: a Fine-grained and Detail-controllable 3D Face Morphable Model from a Hybrid Dataset. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Lizhen Wang, Xiaochen Zhao, Jingxiang Sun, Yuxiang Zhang, Hongwen Zhang, Tao Yu, and Yebin Liu. 2023. StyleAvatar: Real-time Photo-realistic Portrait Avatar from a Single Video. In ACM SIGGRAPH 2023 Conference Proceedings.
- Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-Shot Free-View Neural Talking-Head Synthesis for Video Conferencing. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 10034–10044.
- Thibaut Weise, Sofien Bouaziz, Hao Li, and Mark Pauly. 2011. Realtime Performance-Based Facial Animation. ACM Trans. Graph. 30, 4, Article 77 (jul 2011), 10 pages.
- Olivia Wiles, A. Sophia Koepke, and Andrew Zisserman. 2018. X2Face: A network for controlling face generation using images, audio, and pose codes. In Proceedings of the European Conference on Computer Vision (ECCV).
- Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. 2020. Multiview Neural Surface Reconstruction by Disentangling Geometry and Appearance. In Conference on Neural Information Processing Systems (NeurIPS), Vol. 33.
- Fei Yin, Yong Zhang, Xiaodong Cun, Ming Cao, Yanbo Fan, Xuanxia Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujiu Yang. 2022. StyleHEAT: One-Shot High-Resolution Editable Talking Face Generation via Pre-trained StyleGAN. In Proceedings of the European Conference on Computer Vision (ECCV).
- Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. 2019. Few-Shot Adversarial Learning of Realistic Neural Talking Head Models. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 9458–9467.
- Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. 2021. Flow-Guided One-Shot Talking Face Generation With a High-Resolution Audio-Visual Dataset. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 3661–3670.
- Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C. Bühler, Xu Chen, Michael J. Black, and Otmar Hilliges. 2022. I M Avatar: Implicit Morphable Head Avatars From Videos. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 13535–13545.

- Zerong Zheng, Xiaochen Zhao, Hongwen Zhang, Boning Liu, and Yebin Liu. 2023. AvatarReX: Real-time Expressive Full-body Avatars. ACM Transactions on Graphics (TOG) 42, 4 (2023), 1–19. https://doi.org/10.1145/3592101
 Yiyu Zhuang, Hao Zhu, Xusen Sun, and Xun Cao. 2022. MoFaNeRF: Morphable Facial Neural Radiance Field. In Proceedings of the European Conference on Computer Vision

(ECCV).

zllrunning. 2019. face-parsing.PyTorch. https://github.com/zllrunning/face-parsing. PyTorch.