

EFHQ: Multi-purpose ExtremePose-Face-HQ dataset

Trung Tuan Dao^{1,*} Duc Hong Vu^{1,2,*} Cuong Pham^{1,3} Anh Tran¹

¹VinAI Research, Vietnam ²Miami University ³Posts & Telecommunications Inst. of Tech., Vietnam

{v.trungdt21, v.ducvh5, v.cuongpv11, v.anhtt152}@vinai.io vudh3@miamioh.edu cuongpv@ptit.edu.vn

*Equal contribution

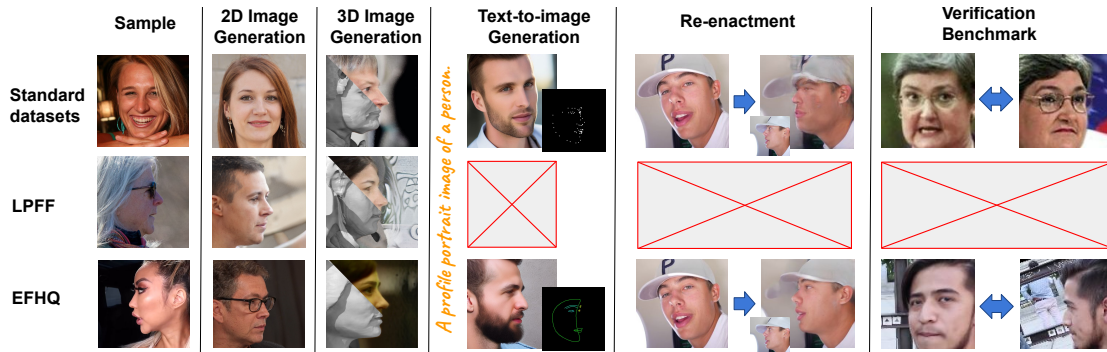


Figure 1. **Benefits of our proposed dataset (EFHQ).** Standard large-scale facial datasets have most images at near frontal views, causing inferior performance of trained models on downstream tasks when dealing with extreme head poses. For instance, the trained 2D image generators and text-to-image ones often produce only near frontal faces, while the 3D face generators and face reenactment methods often show distorted outputs at profile views. The recently proposed dataset LPFF [47] partially handles that issue by providing complementary images at extreme head poses for only 2D and 3D image generation tasks. Our proposed dataset EFHQ provides high-quality extreme-pose images to complement a wide range of face-related tasks. It supports 2D and 3D image generation, with generally better diversity than LPFF. EFHQ also helps correct the outputs of text-to-image generation and face reenactment at extreme views. Finally, EFHQ provides a more challenging pose-based face verification benchmark to better assess the quality of face recognition networks.

Abstract

The existing facial datasets, while having plentiful images at near frontal views, lack images with extreme head poses, leading to the downgraded performance of deep learning models when dealing with profile or pitched faces. This work aims to address this gap by introducing a novel dataset named *Extreme Pose Face High-Quality Dataset (EFHQ)*, which includes a maximum of 450k high-quality images of faces at extreme poses. To produce such a massive dataset, we utilize a novel and meticulous dataset processing pipeline to curate two publicly available datasets, *VFHQ* and *CelebV-HQ*, which contain many high-resolution face videos captured in various settings. Our dataset can complement existing datasets on various facial-related tasks, such as facial synthesis with 2D/3D-aware GAN, diffusion-based text-to-image face generation, and face reenactment. Specifically, training with EFHQ helps models generalize well across diverse poses, significantly improving performance in scenarios involving ex-

treme views, confirmed by extensive experiments. Additionally, we utilize EFHQ to define a challenging cross-view face verification benchmark, in which the performance of SOTA face recognition models drops 5-37% compared to frontal-to-frontal scenarios, aiming to stimulate studies on face recognition under severe pose conditions in the wild.

1. Introduction

The human face is a central focus in computer vision, with extensive research dedicated to tasks like detection, recognition, generation, and manipulation [39, 42]. Numerous large-scale facial datasets [9, 12, 14, 24, 28, 30, 45, 48, 54, 55] have facilitated the face-related studies. Along with deep learning techniques, many tasks have made a considerable leap in performance in recent years. For instance, face recognition systems can recognize near-frontal faces in the wild with near-perfect accuracy [13, 18]. Facial generative models, such as [5, 15, 49], can synthesize realistic

images indistinguishable from real faces. Face reenactment systems, including [41, 43, 51], can animate an input face based on poses and expressions from any driver video.

Despite impressive results, the mentioned models produce inferior results when dealing with faces at extreme head poses. As illustrated in Fig. 1, faces synthesized by 3D face generation or face reenactment models are often distorted at profile views, due to the scarcity of extreme pose images in training datasets. Current extensive datasets for face generation leverage portrait images from Google or stock photo repositories like Flickr, which feature images with upfront views and minimal facial rotations. Face reenactment methods leverage video datasets such as Vox-Celeb [29], primarily featuring frames with frontal faces. The instability of pre-processing tools, particularly face alignment modules when handling extreme pose images, further reduces the proportion of those images in the existing datasets. This issue is critical since faces with large head rotations are common in real life, and dealing with them is unavoidable in practical systems.

This paper aims to fill that void by proposing a novel large-scale dataset of extreme-pose and high-quality facial images called ExtremePose-Face-HQ, or EFHQ for short. Our ambition is to build an “one-for-all” dataset that can supplement existing datasets on all mentioned face-related tasks. To achieve that goal, the dataset needs all the following properties: (1) extreme-pose, (2) large-scale, (3) high-quality, (4) in-the-wild, (5) having multiple images for each subject and having subject identity annotated. Given those requirements, we utilize a novel and meticulous dataset processing pipeline to filter and sample frames from two publicly available datasets, including VFHQ [48] and CelebV-HQ [54], which contain high-resolution face videos captured in various settings. This results in a dataset with 450k high-quality images of extreme poses, which are combined with frontal images of the same subject to use as image pairs. These image pairs can support face reenactment and face recognition tasks, while the extreme-pose images alone can complement the existing datasets of face generation training. Note that there is a recent attempt [47] to build a large-pose facial dataset similar to ours. However, that work relies on images crawled from Flickr, thus having limited size (19k images) and missing identity information. Hence, its applications are restricted, covering only several face-generation tasks, unlike our large-scale and multi-purpose dataset.

We perform a comprehensive set of experiments to validate the effectiveness of EFHQ across three distinct face generation sub-tasks and face reenactment scenarios. Our dataset, EFHQ, when combined with standard training sets, significantly enhances the quality of synthesized images in highly profile or pitched views while maintaining synthesis quality in frontal views. EFHQ also provides a more

challenging pose-centric face verification benchmark. It reveals the vulnerability of many state-of-the-art face recognition networks when dealing with extreme-pose images, with TAR@FAR=1e-3 scores dropping 5-37%.

In summary, our contributions include (1) a large-scale, high-quality, in-the-wild, extreme-pose dataset (EFHQ) designed to complement various face-related tasks, (2) a novel and meticulous dataset processing pipeline ensuring the creation and quality of the EFHQ dataset, and (3) a series of experiments, with plan-to-release models, demonstrating the usefulness of EFHQ in enhancing performance on a wide range of tasks when dealing with faces at extreme rotations. The final dataset will be publicly available as metadata files to not violate the two original dataset’s licenses, alongside the pretrained networks.

2. Related Work

Synthetic Face Generation. Existing 2D-based models like StyleGAN [14, 15, 17] and 3D-aware models like EG3D [5] demonstrate impressive fidelity, especially for near-frontal views. Nevertheless, their performance significantly degrades when faced with extreme views, primarily due to insufficient training data for such scenarios. Datasets like FFHQ [14] exhibit a notable bias toward near-frontal poses, heavily limiting the range of poses synthesized by current 2D and 3D models despite the 3D-aware generators’ ability to handle arbitrary poses.

Diffusion-based text-to-image models [32, 33] can synthesize high-quality images from simple text inputs. However, these models lack mechanisms to control the generated pose. Consequently, the generated facial images remain in a frontal view despite terms like “profile” or “looking-up” in the input prompts. ControlNet[49] proposes a finetuning approach allowing these models to take additional context as conditional input. Namely, one can use landmarks to guide the generated face to the desired pose. However, the scarcity of profile-view data in training affects the model’s performance, particularly in extreme angles. Consequently, effectively managing the long-tailed distribution of these extreme views becomes necessary to address this issue.

Face Reenactment. Face reenactment involves transferring facial expressions and movements from a driving image to a source image. Recent advancements claim high-quality results, especially for near-frontal poses [31, 41, 43, 51]. However, pose collapse persists, particularly in keypoint-based models [41, 51]. Common datasets for this task are VoxCeleb1 [29] and HDTF [50], which primarily feature frontal talking face videos. Despite recent potentially high-quality datasets for reenactment such as VFHQ [48] and CelebV-HQ [54], the frontal-view bias persists. As a result, facial reenactment methods often produce artifacts and lower-quality outputs when faced with significant rotations.

Dataset	#images	Attributes			Ready-to-use tasks		
		Multiple images per ID	High Quality	Profile View	Face Verification	Face Synthesis	Face Reenactment
AgeDB	12K	✓			✓		
CPLFW	12K	✓		✓	✓		
CFP-FP	14K	✓		✓	✓		
IJB-x	138K	✓		✓	✓		
CelebA-HQ	30K		✓			✓	
FFHQ	70K		✓			✓	
LPFF	19K		✓	✓		✓	
CelebV-HQ	6M	✓	✓	✓		✓	
VFHQ	3M	✓	✓	✓		✓	✓
VoxCeleb1	4M	✓				✓	✓
Ours	450,538	✓	✓	✓	✓	✓	✓
FFHQ+Ours*	112,671		✓	✓		✓	
Vox+Ours*	4.5M	✓		✓			✓

Table 1. **Representative comparison** of key attributes and supported tasks across existing face datasets and our proposed EFHQ dataset. EFHQ provides large-scale, high-quality images with multiple views per identity, enabling diverse face-related tasks, including synthesis, reenactment, and verification. Combining EFHQ with existing datasets like FFHQ or VoxCeleb can augment available data for selected tasks. Asterisk (*) indicates a combined dataset with a subsampled version of EFHQ.

Face Verification. Despite near-perfect accuracy on benchmark datasets[12, 20, 27, 28, 44], state-of-the-art face verification models[7, 13, 18] fall short in unconstrained environments due to a lack of pose diversity. While CPLFW[52] introduces 6000 cross-pose pairs, its small size and mix of pose samples hinder analyzing performance on specific pose scenarios such as frontal-frontal, frontal-profile, and profile-profile. More granular benchmarking on diverse pose pairs is crucial to identify failure modes and drive progress in robust face verification across poses.

Extreme Pose Face Dataset. Tab. 1 summarizes existing datasets related to face verification, synthesis, and reenactment, showing that existing face datasets have limitations for multi-task use. Recent efforts like LPFF[47] made progress by crawling from Flickr and processing for diverse views, resulting in a dataset of 19,590 images. The dataset could complement high-quality datasets like FFHQ[14] for generative models. However, due to its limited scale comprising only 20% of the combined dataset, the unbalanced distribution between frontal and profile view still needs further improvements. Additionally, since the images are from image hosting services, the dataset fails to provide the identity information necessary for other facial applications.

3. Dataset Process

In this section, we discuss the process of producing our novel dataset EFHQ. To fulfill our goals, we carefully curated frames from the two recent facial video datasets, VFHQ[48] and CelebV-HQ[54], as seen in Fig. 2.

3.1. Attributes Preparation

Attributes Extraction. To curate accurate facial poses and high-quality frames from existing datasets, we first ex-

tract (1) face bounding boxes, (2) facial landmarks (5/68 keypoints), (3) image quality score, and (4) face identity. For bounding boxes and landmarks, we employ the popular RetinaFace[36] and SynergyNet[46]. We use HyperIQA [40] to grade image quality. For VFHQ, we only generate 68 landmarks and image quality scores since bounding boxes and 5 keypoints are provided. After extracting these attributes, we match them with existing annotations using an IoU-based Hungarian matching algorithm [21]. Regarding identity, VFHQ-wise, we source labels from the annotations. As CelebV-HQ lacks individual identity labels within the video, we use[35] to identify the identities associated with each bounding box.

Due to the instability of head pose estimators when handling extreme-pose images, we strategically apply multiple pose estimators relying on various methodologies: (1) SynergyNet[46], a 3DMM-based model, (2) DirectMHP[53], a landmark-free joint head detector and pose estimator, and (3) an in-house FacePoseNet[6] improved through extensive data augmentation focused on extreme poses and extra training data. We then categorize each estimated pose into a hierarchical binning scheme illustrated in Fig. 2. Next, we gather bin predictions and perform majority voting to arrive at consensus labels. When no agreement emerges, we put it into the 5th bin for “confusing” cases. As depicted in Fig. 3, this ensemble approach provides robustness in cases where a single estimator incorrectly categorizes a sample. Please refer to the supplementary for statistics of the extracted bins and the hyperparameters for each model used for attribute extraction.

Annotation Review. Next, we manually review the results to ensure high-quality labels. Given the large dataset, we developed a graphical user interface tool to streamline the review process. Our focus is verifying the binning annota-

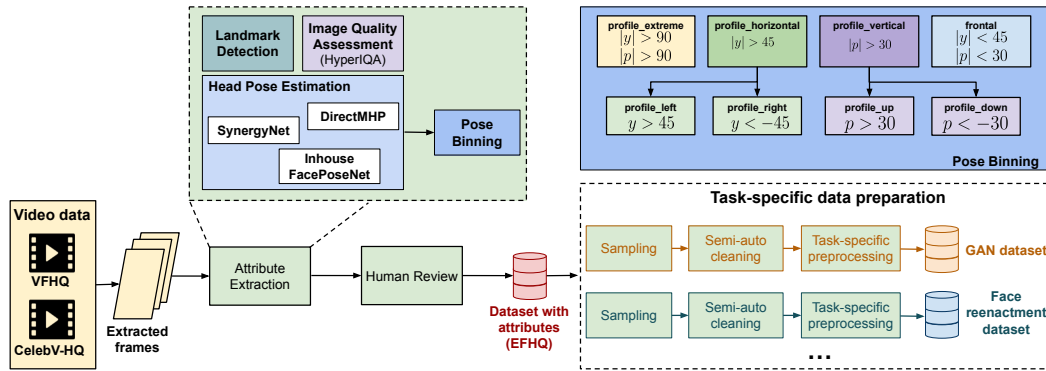


Figure 2. **The pipeline of EFHQ dataset creation.** Starting with high-quality videos from the VFHQ[48] and CelebV-HQ[54] datasets, single-frame attributes are extracted then manually reviews. Task-specific preprocessing is then applied to generate specialized versions of the dataset for tasks such as face generation, reenactment, and verification.

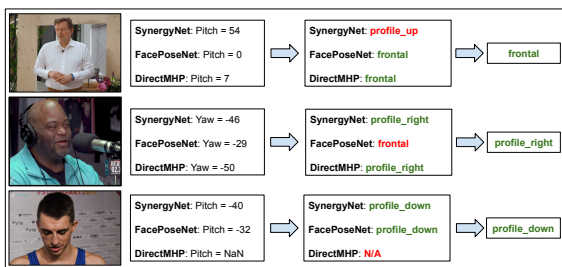


Figure 3. **Example cases** where a pose estimator fails to categorize the sample to the correct bin.

tions, which most directly impact label accuracy. Concurrently, we randomly subsample the data based on pose angle and image quality to validate facial landmark quality further and discard cases of landmark prediction failure.

3.2. EFHQ dataset

Our final EFHQ dataset comprises up to 450k frames with extreme poses extracted from approximately 5,000 clips. Most of the clips in our dataset include at least one frame with a frontal face and multiple frames with extreme pose angles. We also include extreme pose-only clips, representing profile-to-profile pose transfer cases.




Figure 4. **Pose distribution** comparison between our sampled dataset and other datasets, including FFHQ[14], LFFF[47], CPLFW[52] and 40K random samples from VoxCeleb1 [29]. Our sampled dataset demonstrates greater pose diversity, with increased sample counts across high angle bins.

3.3. Dataset for subtasks

Supplementary datasets for face generation. To address FFHQ’s pose distribution gap, we compile a dataset encompassing diverse poses, varied identities, and FFHQ-comparable image quality from EFHQ. Additionally, we also embed an image brightness filter[3] in the sampling process, given that HyperIQA doesn’t explicitly assess this aspect. The sampled images are then processed using the same pipeline in [5, 16] and manually reviewed. The aim is to improve coverage of under-represented poses without degrading frontal performance of trained generation models. This yields a dataset of 42,671 images with equitably distributed poses as visualized in Fig. 4. Compared to the original FFHQ distribution, yaw and pitch distributions are significantly improved.

Additionally, for 3D-aware GAN subtasks, we extract camera parameters to serve as conditional pose information, following [5]. We convert these parameters into yaw, pitch angles to validate against our existing annotations. When disagreements emerge, we re-examine the images and filter if needed, to ensure accurate, quality labels.

For diffusion-based text-to-image generation, a specialized dataset is curated to refine Stable Diffusion models[33] by accommodating landmark conditional input through ControlNet[49]. Based on our 3D-aware face generation dataset, we integrate landmark-based conditional images and tailored text prompts for each image. To capture fine facial nuances, we extract 478 facial landmarks and draw the condition image with connected edges between matched keypoints, following Mediapipe framework[26], better-conveying face expression and pose. For text prompts, detailed facial attributes like gender, race, and emotion are derived from existing labels or inferred via BLIP-2 pretrained captioning model [23] and a face attribute estimator [37]. The resulting prompt adheres to a structured format: “A profile portrait image of a [emotion] [race] [gender].”

Model	Reference Dataset	FID↓	Recall↑
G^{FFHQ}	FFHQ	2.84	0.49
$G^{\text{FFHQ+LPFF}}$	FFHQ+LPFF	3.43	0.44
$G^{\text{FFHQ+EFHQ}}$	FFHQ+EFHQ	3.44	0.46
$G^{\text{FFHQ+LPFF}}$	FFHQ+EFHQ	8.26	0.43
$G^{\text{FFHQ+EFHQ}}$	FFHQ+LPFF	7.04	0.44
$G_c^{\text{FFHQ+EFHQ}}$	FFHQ+EFHQ	3.33	0.46
$G_c^{\text{FFHQ+EFHQ}}$	FFHQ	3.12	0.44
$G_c^{\text{FFHQ+EFHQ}}$	FFHQ+LPFF	4.47	0.42

Table 2. **StyleGAN2-ADA models trained on different datasets.** Lower FIDs indicate better fidelity, while higher Recalls indicate better diversity. The second and third blocks show comparisons between unconditional models trained on FFHQ+LPFF and FFHQ+EFHQ when evaluating on the same and cross-dataset settings, with the better metrics in **bold**.

Supplementary dataset for face reenactment. As EFHQ is crafted, we can employ the entire dataset to complement established training datasets, e.g., VoxCeleb1 [29]. We also built an extra evaluation set with 1200 EFHQ clips.

Benchmarking dataset for face verification. We curate a dataset covering three distinct scenarios: frontal-to-frontal, frontal-to-profile, and profile-to-profile. To enable rigorous benchmarking, we sample 10,000 pairs each for both negative and positive cases per scenario, resulting in a balanced benchmark dataset containing 60,000 image pairs. Images follow the same established preprocessing pipelines in [7]. Next, we filter misaligned images with RetinaFace[36], randomly review, and replace equivalent samples from our diverse corpus, if needed. Finally, to simulate varying image quality, we randomly applied downscaling and compression to samples in the dataset.

4. Face Generation Subtask

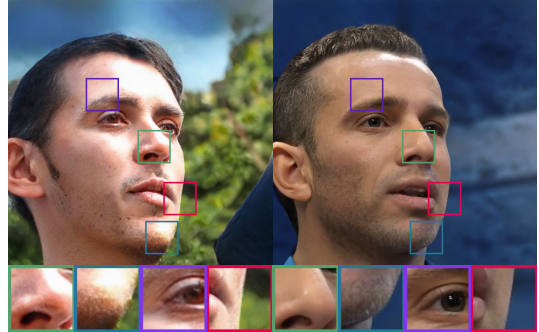
4.1. StyleGAN

Training details. We train StyleGAN2-ADA[15] models from scratch after combining FFHQ with our own curated dataset (denoted FFHQ+EFHQ) at 1024×1024 resolution. We also train a conditional version, allowing direct control over frontal versus profile synthesis and enhancing model performance. Specifically, we categorize the dataset into two classes: “frontal” and “profile” based on pose angles. All models are trained for 6 days on 8 Nvidia A100 GPUs.

Evaluation protocol. We utilize the Fréchet Inception Distance[11] (FID) metric and improved Recall[22] to quantify the generation fidelity and diversity. Regarding the two unconditional models trained with the supplemented FFHQ+LPFF and FFHQ+EFHQ, we evaluate each model with both of these datasets. For conditional mod-



(a) Samples from the model trained with FFHQ+EFHQ.



(b) Profile-view comparison of LPFF (left) and EFHQ (right).

Figure 5. **Qualitative result and comparison** of generated samples, with truncation $\psi = 0.7$, from StyleGAN2-ADA training with other dataset and ours.

els, we evaluate performance on FFHQ, FFHQ+LPFF, and FFHQ+EFHQ to demonstrate the broader coverage of the model trained with our proposed dataset. For FFHQ, we generate 50K images with the “frontal” label, whereas for the other two datasets, we generate 50K following their label distributions. This comprehensive quantitative analysis examines the fidelity and variety of the generated images.

Experimental results. We compared the performance of our proposed unconditional models ($G^{\text{FFHQ+EFHQ}}$), against models trained on the FFHQ and FFHQ+LPFF datasets (G^{FFHQ} and $G^{\text{FFHQ+LPFF}}$). Models trained conditionally are specifically subscripted, e.g., $G_c^{\text{FFHQ+EFHQ}}$. The qualitative results in Fig. 5 illustrate that our models generate high-quality frontal faces comparable to the FFHQ model while producing realistic and varied profile faces. Additionally, our method has less noise in extreme poses than the FFHQ+LPFF model. Specifically, Fig. 5b shows LPFF’s generated face contains more noise and less photorealistic details than our model, such as the pixelated eye and nose patches. Meanwhile, ours achieves greater realism with smoother skin, well-defined features, noise reduction, and natural lighting. The quantitative results in Tab. 2 show that our unconditional model achieves competitive FID scores and improved Recall to other models, indicating similar fidelity but greater sample pose diversity. In the cross-examination (third block of Tab. 2), $G^{\text{FFHQ+EFHQ}}$ has both better FID and improved Recall on the opposite dataset than $G^{\text{FFHQ+LPFF}}$. Finally, our conditional model demonstrates substantially improved FID scores over multiple reference datasets, highlighting its versatility.

Model	Reference Dataset	FID↓	Pose↓	Depth↓	ID↑
E^{FFHQ}	FFHQ/Trainset	4.8	$5e^{-4}$	0.234	<u>0.74</u>
E^{bug}	FFHQ/Trainset	3.7	$5e^{-4}$	0.288	<u>0.74</u>
$E^{\text{FFHQ+LPFF}}$	FFHQ	6.2	$6e^{-4}$	0.316	0.77
$E^{\text{FFHQ+EFHQ}}$	FFHQ	3.9	$4e^{-4}$	0.224	0.77
$E^{\text{FFHQ+LPFF}}$	Trainset	7.0	$6e^{-4}$	0.316	0.65
$E^{\text{FFHQ+EFHQ}}$	Trainset	4.6	<u>$4e^{-4}$</u>	0.223	0.65

Table 3. **EG3D models trained on different datasets.** In the second and third blocks, we compare models trained on FFHQ+LPFF and FFHQ+EFHQ with the testset is either FFHQ or their training datasets; the better metrics are in **bold**. The best metrics on the FFHQ testset are in **red**, while the best metrics when evaluating on training dataset are underlined.




Figure 6. **Comparison** between multiview generated samples, with truncation $\psi = 0.8$, of EG3D model trained with various datasets.

4.2. EG3D

Training details. We followed the same pipeline and hyper-parameters as the original EG3D model [5]. However, based on empirical results, we adjusted $\gamma = 5$. We also fixed the triplane implementation bug (XY, XZ, ZX) [1] to replicate the original idea in the papers (XY, XZ, ZY). We trained the model on 8 Nvidia A100 40GB GPUs for approximately 7 days until convergence.

Evaluation Protocol. To evaluate the models, mirroring [5], we measure image quality with FID; multi-view consistency and identity preservation with mean cosine similarity of pre-trained face recognition model[7]; pose accuracy and geometry by calculating MSE against pseudo labels from [8]. We evaluate all models on 1,024 generated images, except for FID, which we compute on 50k images.

Experimental results. To evaluate the benefits of our dataset, we trained the same backbone architecture on different datasets: FFHQ (E^{FFHQ}), FFHQ+LPFF ($E^{\text{FFHQ+LPFF}}$), and FFHQ+EFHQ (our model, $E^{\text{FFHQ+EFHQ}}$). The quantitative results are presented in Table 3, and qualitative examples are shown in Figure 6. Qualitatively, we observe that $E^{\text{FFHQ+EFHQ}}$ generates improved face shapes in extreme poses compared to E^{FFHQ} , similar to the improvements seen with $E^{\text{FFHQ+LPFF}}$. Quantitatively, the FID

Metrics	Dataset	Pretrained ControlNet	Ours
NME↓	LPFF	0.085	0.084
FID↓	LPFF	119.74	83.39
FID _{eyepatch} ↓	LPFF	103.77	49.45
NME↓	EFHQ	0.106	0.081
FID↓	EFHQ	111.23	81.90
FID _{eyepatch} ↓	EFHQ	137.43	58.60

Table 4. **Quantitative results** of fine-tuning StableDiffusion v1.5 with ControlNet trained on various datasets, using Normalized Mean Error (NME), FID, and FID on eye region.

of $E^{\text{FFHQ+EFHQ}}$ is comparable to E^{FFHQ} when evaluated on FFHQ, indicating no degradation in frontal view performance. Meanwhile, the competitive FID achieved by $E^{\text{FFHQ+EFHQ}}$ on the combined FFHQ+EFHQ dataset also suggests proficient modeling of profile views. Regarding identity consistency, we see a significant drop when the reference dataset contains profile-view images, due to the weak confidence of the pretrained face recognition model when handling profile-view images, as confirmed by its low average cosine similarity of around 0.6 on the simple frontal-to-profile verification set CFP-FP[34]. Furthermore, our model also outperforms in pose accuracy and geometry quality with both frontal and cross-pose dataset.

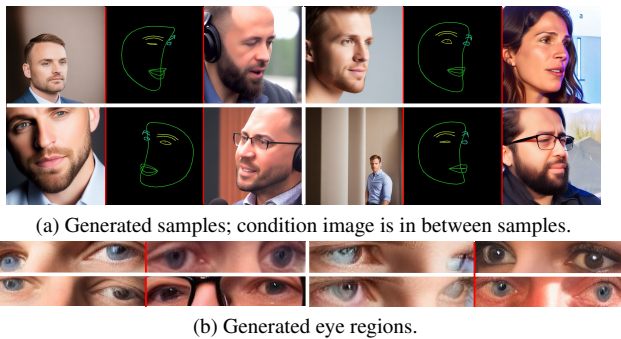


Figure 7. **Comparison** profile-view generated samples of pre-trained ControlNet (left) and our fine-tuned ControlNet (right) with the prompt: “A profile portrait image of a person”.

4.3. ControlNet

Training details. We finetuned Stable Diffusion v1.5 [33] using ControlNet [49] on the EFHQ dataset to enhance its ability to generate extreme facial poses. We use AdamW optimizer [25] with learning rate $1 \times e^{-5}$, $\beta_0 = 0.9$, $\beta_1 = 0.999$, and $\epsilon = 0.01$. We train with batch size 24 on 1 Nvidia A100 40GB GPU in 2 days.

Evaluation Protocol. To assess image qualities across diverse poses, we measure the Normalized Mean Error (NME) of 478-point facial landmarks between the generated and source image. Specifically, we evaluate performance across the full pose distribution in the reference

Model	Test Dataset (\mathcal{L}_1 ↓ / AKD ↓ / AED ↓)	
	VoxCeleb1	EFHQ (Ours)
$TPS^{\text{VoxCeleb1}}$	0.041/1.214/0.132	0.061/3.814/0.275
$TPS^{\text{VoxCeleb1+EFHQ}}$	0.041/1.250/0.133	0.051/1.963/0.214
$LIA^{\text{VoxCeleb1}}$	0.044/1.342/0.138	0.059/3.807/0.266
$LIA^{\text{VoxCeleb1+EFHQ}}$	0.044/1.398/0.142	0.054/2.291/0.233

Table 5. **Quantitative results** of two state-of-the-art reenactment methods (TPS and LIA) when training on VoxCeleb1 and the combined set VoxCeleb1+EFHQ. We present results evaluated either on VoxCeleb1 (left) or EFHQ (right) test sets.

dataset. For the prompt, we use: “A profile portrait image of a person.” with extra positive and negative keywords, further detailed in supplementary. Utilizing LPFF and EFHQ datasets as the reference datasets, we compute FID to assess synthesized face quality and extend evaluation to cropped eye regions due to unrealistic gazes observed from the original ControlNet model. Note that the models receive condition images created akin to its training data.

Experimental results. Fig. 7a visually compares the baseline ControlNet OpenPose face and our fine-tuned model. Our fine-tuned model adeptly conditions diverse extreme poses, generating high-quality results with precise details while maintaining fidelity to the condition, such as mouth shape, and exhibiting minimal artifacts. These patterns indicate the superiority of the fine-tuned model over the baseline model in handling extreme viewing angles. Additionally, Fig. 7b contrasts the eye regions of both models. Our fine-tuned model accurately captures natural eye gaze and shape from condition images, a notable improvement over the baseline. Quantitative results in Tab. 4 consistently favor our fine-tuned model across all metrics. FID scores on full and eye patches exhibit significant enhancement, and NME performance remains stable on training and reference datasets (EFHQ and LPFF). These findings showcase the fine-tuned model’s capability to generate high-quality, geometrically consistent faces across diverse poses.

5. Face Reenactment Subtask

5.1. Evaluation Protocol

We follow the same preprocessing and evaluation protocol as FOMM [38]. Specifically, all the frames are cropped to 256×256 , and we use the first frame of a clip as the source image to reconstruct the remaining frames. In the case of the EFHQ test set, most source images have a frontal pose, while the target frames have an extreme pose. Some test cases only contain extreme pose to evaluate self-reconstruction on these cases. We run the experiments using two recent methods, Thin-Plate Spline Motion Model (TPS) [51] and Latent Image Animator (LIA) [43]. For training details, please refer to the supplementary. We select the following three metrics to compare the generated frames and the driving ground truth. \mathcal{L}_1 measures the av-

Model	Trainset	CPLFW	EFHQ		
			f2f	f2p	p2p
$R18_{\text{arcface}}$	Glint360k	0.80	0.94	0.57 _{0.37}	0.88 _{0.06}
$R18_{\text{arcface}}$	MS1MV3	0.75	0.93	0.57 _{0.36}	0.88 _{0.05}
$R50_{\text{arcface}}$	Glint360k	0.91	0.97	0.77 _{0.20}	0.92 _{0.05}
$R100_{\text{arcface}}$	Glint360k	0.93	0.98	0.83 _{0.15}	0.82 _{0.16}
$R100_{\text{adaface}}$	WebFace4M	0.94	0.99	0.94 _{0.05}	0.97 _{0.02}

Table 6. **Quantitative evaluation** of state-of-the-art face recognition models on the cleansed CPLFW and the proposed benchmark dataset. We use the metric TAR@FAR=1e-3 (higher is better). Our EFHQ benchmark includes frontal-frontal (f2f), frontal-profile (f2p), and profile-profile (p2p) subsets. For f2p and p2p, we report the score drop when comparing to f2f in the red subscription. The lowest benchmark score for each model is in bold.

erage distance among pixel values, indicating reconstruction faithfulness. Average Keypoint Distance (AKD) calculates the mean distance between keypoints, thus assessing the pose accuracy. Average Euclidean Distance (AED) measures identity preservation by computing the average \mathcal{L}_2 distance between identity embeddings.

5.2. Experimental results

We compare TPS and LIA models trained with our supplementary EFHQ dataset ($TPS^{\text{VoxCeleb1+EFHQ}}$ and $LIA^{\text{VoxCeleb1+EFHQ}}$) against models trained solely on VoxCeleb1 ($TPS^{\text{VoxCeleb1}}$ and $LIA^{\text{VoxCeleb1}}$). Tab. 5 displays quantitative results, revealing that models trained with EFHQ maintain comparable performance when evaluated on VoxCeleb1, emphasizing high-quality performance in frontal-to-frontal reenactment scenarios. Notably, on the EFHQ test set, both EFHQ-trained models show significant improvement in all metrics, indicating enhanced extreme pose and identity reenactment performance. In Fig. 8, qualitative results prove that EFHQ-trained models excel in transferring facial expressions and motion across diverse poses. Models trained on EFHQ showcase improved image quality, reduced artifacts, and well-preserved shapes of driving faces. These results highlight our models’ effectiveness in capturing and reproducing facial motions across varying poses while enhancing fidelity to the original facial features.

6. Face Recognition Subtask

6.1. Evaluation Protocol

To evaluate face recognition performance on extreme poses, we selected two state-of-the-art methods: ArcFace[7] and AdaFace[18]. We further explored various backbone sizes (ResNet18, ResNet50, ResNet101)[10] and training datasets (MS1MV2[9], Glint360K[2], WebFace4M[55]) of increasing size. For metrics, we used True Acceptance Rate (TAR) at False Acceptance Rate (FAR) of $1 \times e^{-3}$.




Figure 8. **Comparison** of performance for frontal-profile face reenactment on EFHQ. From left to right: source image, TPS on VoxCeleb1, TPS with EFHQ, LIA on VoxCeleb1, LIA with EFHQ, driving image. Additional EFHQ training data improves synthesis of extreme poses.

6.2. Experimental results

Before assessing face recognition networks on EFHQ, we first clean the commonly-used version of CPLFW, processed by [7], by removing misaligned images (please refer to the Supplementary), and evaluate the models on it. Models performed well on this cleansed set, with minimal gaps between ArcFace and AdaFace.

Our proposed evaluation benchmark involves three pose-based scenarios. Results in Tab. 6 reveal performance drops of 5-37% on the frontal-profile subset compared to frontal-frontal, highlighting the challenge of matching faces in different poses. Larger ResNet models outperformed ResNet18 across subsets, showcasing the advantages of increased capacity in handling pose variation. However, the ResNet100 model trained on Glint360K performed worse than smaller models on the profile-profile subset, indicating the need for sufficient profile data to prevent overfitting. Our results also indicate AdaFace’s suitability for handling pose variations. In summary, pose mismatch brings significant challenges for face recognition, necessitating larger models trained on diverse data covering profile faces.

7. Discussion




Figure 9. **User study results.** We reported the percentage of times testers rank whose output is best or tied for each criterion.

User survey. We conduct user surveys to evaluate the face generation improvement of our method through human per-

ceptual assessment. In this survey, we compare the outputs of the models mentioned in Secs. 4 and 5. The setup of this user survey is included in the supplementary. In all human evaluation results (Fig. 9), our models outperform previous ones, notably on identity and motion test of face reenactment models (74.80% and 64.22%, respectively).




Figure 10. **Representative examples** of the raw LPPF dataset that got excluded from the final dataset, either from misdetections or pose filtering process. **Red** keypoints represents landmarks prediction of selected samples from LPPF, **green** keypoints represents landmarks prediction from our pipeline.

Robustness of data processing pipeline. As shown in Fig. 10, our data processing pipeline can detect additional profile-view faces filtered out by the LPPF pipeline, likely due to misdetections or incorrect pose filtering. This confirms the robustness of our pipeline, and we hope it will enable the community to develop larger and higher-quality in-the-wild facial datasets.

8. Conclusion

This work has introduced a large-scale, diverse facial dataset to address performance gaps between frontal and profile faces, powered by a novel and robust data processing pipeline. Crucially, we provided tailored sub-datasets for advancing essential tasks like face synthesis and reenactment. Moreover, our new face verification benchmark reveals gaps between techniques, granting valuable insights. Ultimately, we hope this high-quality, diverse facial dataset opens up new and exciting opportunities to push forward cross-pose tasks.

References

- [1] Fix for triplane projection · Issue #67 · NVlabs/eg3d. <https://github.com/NVlabs/eg3d/issues/67>, 2023. 6
- [2] Xiang An, Xuhan Zhu, Yuan Gao, Yang Xiao, Yongle Zhao, Ziyong Feng, Lan Wu, Bin Qin, Ming Zhang, Debing Zhang, et al. Partial fc: Training 10 million identities on a single machine. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1445–1449, 2021. 7
- [3] Sergey Bezryadin, Pavel Bourov, and Dmitry Ilinih. Brightness calculation in digital image processing. *International Symposium on Technologies for Digital Photo Fulfillment*, 2007:10–15, 2007. 4, 1
- [4] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 4
- [5] 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. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022. 1, 2, 4, 6
- [6] Feng-Ju Chang, Anh Tuan Tran, Tal Hassner, Iacopo Masi, Ram Nevatia, and Gerard Medioni. Faceposenet: Making a case for landmark-free face alignment. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1599–1608, 2017. 3
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. 3, 5, 6, 7, 8, 1, 2
- [8] Yu Deng, Jialong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *IEEE Computer Vision and Pattern Recognition Workshops*, 2019. 6, 3
- [9] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 87–102. Springer, 2016. 1, 7
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778. IEEE, 2016. 7
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017. 5
- [12] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008. 1, 3
- [13] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: adaptive curriculum learning loss for deep face recognition. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5901–5910, 2020. 1, 3
- [14] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 1, 2, 3, 4
- [15] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Proc. NeurIPS*, 2020. 1, 2, 5
- [16] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 4
- [17] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Proc. NeurIPS*, 2021. 2
- [18] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 3, 7
- [19] Davis E. King. Dlib-ml: A machine learning toolkit. *J. Mach. Learn. Res.*, 10:1755–1758, 2009. 3
- [20] Brendan F Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, and Anil K Jain. Pushing the frontiers of unconstrained face detection and recognition: Iarpa janus benchmark a. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1931–1939, 2015. 3
- [21] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 3
- [22] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2019. Curran Associates Inc. 5
- [23] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. 4
- [24] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015. 1
- [25] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2017. 6
- [26] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuoling Chang, Ming Guang Yong, Juhyun Lee, Wan-Teh

- Chang, Wei Hua, Manfred Georg, and Matthias Grundmann. Mediapipe: A framework for building perception pipelines. *CoRR*, abs/1906.08172, 2019. 4
- [27] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *2018 international conference on biometrics (ICB)*, pages 158–165. IEEE, 2018. 3
- [28] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: the first manually collected, in-the-wild age database. In *proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 51–59, 2017. 1, 3
- [29] Arsha Nagrani, Joon Son Chung, and Andrew Senior. VoxCeleb: A large-scale speaker identification dataset. In *Interspeech 2017*. ISCA, 2017. 2, 4, 5
- [30] Aaron Nech and Ira Kemelmacher-Shlizerman. Level playing field for million scale face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1
- [31] Youxin Pang, Yong Zhang, Weize Quan, Yanbo Fan, Xiaodong Cun, Ying Shan, and Dong-Ming Yan. Dpe: Disentanglement of pose and expression for general video portrait editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 427–436, 2023. 2
- [32] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv e-prints*, pages arXiv–2204, 2022. 2
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 4, 6
- [34] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Castillo, Vishal M Patel, Rama Chellappa, and David W Jacobs. Frontal to profile face verification in the wild. In *2016 IEEE winter conference on applications of computer vision (WACV)*, pages 1–9. IEEE, 2016. 6
- [35] Sefik Ilkin Serengil and Alper Ozpinar. Lightface: A hybrid deep face recognition framework. In *2020 Innovations in Intelligent Systems and Applications Conference (IASU)*, pages 23–27. IEEE, 2020. 3
- [36] Sefik Ilkin Serengil and Alper Ozpinar. Hyperextended lightface: A facial attribute analysis framework. In *2021 International Conference on Engineering and Emerging Technologies (ICEET)*, pages 1–4. IEEE, 2021. 3, 5, 1, 2
- [37] Sefik Ilkin Serengil and Alper Ozpinar. Hyperextended lightface: A facial attribute analysis framework. In *2021 International Conference on Engineering and Emerging Technologies (ICEET)*, pages 1–4. IEEE, 2021. 4
- [38] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in neural information processing systems*, 32, 2019. 7, 2
- [39] L Sirovich and M Kirby. Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America A*, 4(3):519–524, 1987. 1
- [40] Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [41] Jiale Tao, Biao Wang, Borun Xu, Tiezheng Ge, Yuning Jiang, Wen Li, and Lixin Duan. Structure-aware motion transfer with deformable anchor model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3637–3646, 2022. 2
- [42] Matthew A Turk and Alex P Pentland. Face recognition using eigenfaces. In *Proceedings. 1991 IEEE computer society conference on computer vision and pattern recognition*, pages 586–587. IEEE Computer Society, 1991. 1
- [43] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. Latent image animator: Learning to animate images via latent space navigation. In *ICLR 2022-The International Conference on Learning Representations*, 2022. 2, 7, 3
- [44] Cameron Whitelam, Emma Taborsky, Austin Blanton, Brianna Maze, Jocelyn Adams, Tim Miller, Nathan Kalka, Anil K Jain, James A Duncan, Kristen Allen, et al. Iarpa janus benchmark-b face dataset. In *proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 90–98, 2017. 3
- [45] L Wolf, T Hassner, and I Maoz. Face recognition in unconstrained videos with matched background similarity. In *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, pages 529–534, 2011. 1
- [46] Cho-Ying Wu, Qiangeng Xu, and Ulrich Neumann. Synergy between 3dmm and 3d landmarks for accurate 3d facial geometry. In *2021 International Conference on 3D Vision (3DV)*, 2021. 3
- [47] Yiqian Wu, Jing Zhang, Hongbo Fu, and Xiaogang Jin. Lpff: A portrait dataset for face generators across large poses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20327–20337, 2023. 1, 2, 3, 4
- [48] Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. Vfhq: A high-quality dataset and benchmark for video face super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2022. 1, 2, 3, 4
- [49] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 1, 2, 4, 6
- [50] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. 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*, pages 3661–3670, 2021. 2
- [51] Jian Zhao and Hui Zhang. Thin-plate spline motion model for image animation. In *Proceedings of the IEEE/CVF Con-*

- ference on Computer Vision and Pattern Recognition*, pages 3657–3666, 2022. [2](#), [7](#), [3](#)
- [52] Tianyue Zheng and Weihong Deng. Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments. *Beijing University of Posts and Telecommunications, Tech. Rep*, 5(7), 2018. [3](#), [4](#), [2](#)
- [53] Huayi Zhou, Fei Jiang, and Hongtao Lu. Directmhp: Direct 2d multi-person head pose estimation with full-range angles. *arXiv preprint arXiv:2302.01110*, 2023. [3](#)
- [54] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. CelebV-HQ: A large-scale video facial attributes dataset. In *ECCV*, 2022. [1](#), [2](#), [3](#), [4](#)
- [55] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Jiwen Lu, Dalong Du, and Jie Zhou. Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#), [7](#)