

AUEditNet: Dual-Branch Facial Action Unit Intensity Manipulation with Implicit Disentanglement

Shiwei Jin¹, Zhen Wang², Lei Wang², Peng Liu², Ning Bi², Truong Nguyen¹

¹ ECE Dept. UC San Diego, ² Qualcomm Technologies, Inc.

{sjin, tqn001}@eng.ucsd.edu, {zhewang, wlei, peli, nbi}@qti.qualcomm.com

Abstract

Facial action unit (AU) intensity plays a pivotal role in quantifying fine-grained expression behaviors, which is an effective condition for facial expression manipulation. However, publicly available datasets containing intensity annotations for multiple AUs remain severely limited, often featuring a restricted number of subjects. This limitation places challenges to the AU intensity manipulation in images due to disentanglement issues, leading researchers to resort to other large datasets with pretrained AU intensity estimators for pseudo labels. In addressing this constraint and fully leveraging manual annotations of AU intensities for precise manipulation, we introduce AUEditNet. Our proposed model achieves impressive intensity manipulation across 12 AUs, trained effectively with only 18 subjects. Utilizing a dual-branch architecture, our approach achieves comprehensive disentanglement of facial attributes and identity without necessitating additional loss functions or implementing with large batch sizes. This approach offers a potential solution to achieve desired facial attribute editing despite the dataset's limited subject count. Our experiments demonstrate AUEditNet's superior accuracy in editing AU intensities, affirming its capability in disentangling facial attributes and identity within a limited subject pool. AUEditNet allows conditioning by either intensity values or target images, eliminating the need for constructing AU combinations for specific facial expression synthesis. Moreover, AU intensity estimation, as a downstream task, validates the consistency between real and edited images, confirming the effectiveness of our proposed AU intensity manipulation method.

1. Introduction

Facial action units (AUs), serving as anatomical indicators of facial muscle movements, have been effectively utilized as conditions for fine-grained facial expression editing in images [20, 29]. The manipulation of AU intensities offers advantages such as objective quantification on a six-integer-

level ordinal scale defined by the Facial Action Coding System (FACS) [8], the ability to generate over 7000 combinations in observed facial expressions with a small number of AUs (30) [33], and the potentials for continuous intensity manipulation, instead of the category-based expression editing [6]. However, public datasets containing intensity annotations for over 10 AUs are constrained by limited subject counts, and frame-level AU intensity annotation requires expert involvement and extensive works. As a result, current AU intensity manipulation methods [20, 29, 40] often resort to the pretrained AU intensity estimator [3] to obtain predicted annotations for datasets with larger subject pools, sidestepping the reliance on expert-labeled datasets with a restricted number of subjects.

On the other hand, the semantic richness in latent space and high-quality generation capability of StyleGAN [18] have facilitated the development of facial attribute editing methods [4, 7, 11, 21, 36] that enable targeted modifications without affecting other attributes and identity. However, searching unified editing directions in the latent space for attribute editing typically requires substantial data from numerous subjects to disentangle the target attributes from others and identity. Limited number of subjects may lead to overfitting issues and poor generalization to new faces.

Considering these, it is challenging to search disentangled editing directions for manipulating intensities of multiple AUs based on the data from limited subjects. To address this, we propose a method to manipulate intensities of 12 AUs within the W^+ latent space [1] of StyleGAN [18] for high-resolution face image synthesis using only 18 subjects' data. Specifically, we introduce a novel pipeline designed to enforce disentanglement within the network, even with a dataset containing a limited number of subjects compared to the number of target facial attributes we aim to edit. This approach offers a potential solution to achieve desired facial attribute editing despite the dataset's limited subject count. To summarize, our contributions are as follows:

- Achieve accurate AU intensity manipulation in high-resolution synthesized face images conditioned by AU intensity values or target images without requiring net-

work retraining or extra AU estimators.

- Introduce an architecture designed to disentangle target attributes from others and identity, even when working with data containing very few subjects compared to the number of target facial attributes we aim to edit.
- Propose the encoding of labels to match the level-wise disentangled structure of latent vectors in W^+ to avoid entangled labels as conditions for editing.
- Demonstrate the ability to manipulate float or negative AU intensities while generating consistent results, despite the training set labels encompassing six levels.

2. Related Work

2.1. AU Intensity Manipulation

GANimation [29] is an early work that utilizes AU intensities as conditions for facial expression manipulation. However, it suffers from attention mechanism issues that could result in overlap artifacts in regions where facial deformations occur [30]. Ling *et al.* [20] propose using the relative AU intensities between the source and target images as conditions, avoiding the direct addition of new attributes onto the existing expression [40]. Alternatively, ICface [40] introduces a two-stage editing pipeline. The initial stage transforms the input image into a neutral one with all AU intensities set to zero, and the second stage maps this neutral status to the final output, depicting the desired driving attributes with two independent generators. However, the architecture of ICface is redundant and resource-intensive. FACEGAN [41] utilizes AU representations to construct facial landmarks for expression transfer, reducing the potential of identity leakage from the target image. These methods place greater emphasis on facial expressions compared to AUs, both in terms of their editing goals and evaluation criteria.

2.2. Image Editing in Latent Space

The latent space working with StyleGAN2 [18] is well-known of its meaningful and highly disentangled properties. Several unsupervised methods [11, 35, 42, 44] search editing directions in the latent space without the need for attributes labels. For instance, GANSpace [11] employs principal component analysis to identify semantic editing directions in the latent space. In contrast, supervised methods [4, 7, 14, 36] typically rely on pretrained attribute estimators or attribute labels. InterFaceGAN [36], for example, utilizes a binary support vector machine [25] to estimate hyperplanes for the corresponding attribute editing. Furthermore, some methods [21, 28, 48] use the CLIP loss [31] to enable text-driven image manipulation. These methods usually handle identity information effortlessly since commonly used datasets contain a much larger number of subjects compared to the attributes involved in the editing process. However, in certain cases with a limited number of subjects included, the

identity issue becomes significant. Therefore, in our work, we introduce a novel architecture designed to implicitly disentangle identity information from multiple attributes, even when dealing with a restricted number of subjects.

3. Proposed Method

3.1. Problem Setting

Our objective is to develop an intermediate module within the pretrained GAN inversion pipeline that enables the modification of specific facial attributes in input face images based on target conditions, while preserving the individual’s identity and leaving other attributes unaffected. A crucial aspect of achieving this lies in effectively disentangling the target facial attributes from others and from identity. Prior works [11, 36] focused on identifying global editing directions in latent space for desired facial attributes by analyzing data from thousands of subjects. The data includes a significantly larger number of subjects than the number of facial attributes aiming to edit. Consequently, it is common for different subjects in the dataset to share the same facial attributes. This characteristic naturally facilitates the disentanglement of identity-related influences from the identified global editing directions for the corresponding attribute editing.

However, while data availability from numerous subjects is abundant, obtaining fine-grained labels poses challenges. The significant tradeoff between data collection and annotation efforts, particularly when expert annotation is necessary, can hinder the inclusion of detailed labels. In specific face image editing tasks, such as AU intensity manipulation, multi-level intensity labels offer advantages over binary labels (activated or not). Yet, datasets with intensity labels covers more AUs often comprise fewer subjects. The limited subject pool in facial attribute editing may blend identity features with the target attributes, complicating disentanglement processes. To address this, we propose a novel framework named *AUEditNet*. This architecture enables seamless intensity adjustments across 12 AUs in face images, even when trained on a restricted dataset containing only 18 subjects.

3.2. Pipeline Overview

Consistent with previous works [4, 14, 21], we use a GAN inversion pair that consists of an encoder \mathcal{E} and a generator \mathcal{G} to achieve the transformation between the image space and the latent space. All editing occurs in the latent space. During training, we use a pair of images I_{src}, I_{tar} from one subject, while an additional face image I_{rnd} is randomly chosen from another subject’s data. The processes are visually depicted in Fig. 1 with detailed descriptions.

Feature Space for Target AUs. Initially, we encode the input source image I_{src} into the latent vectors $W_{src} = \mathcal{E}(I_{src})$. To achieve explicit control over facial attributes using conditions associated with physical interpretation, we perform

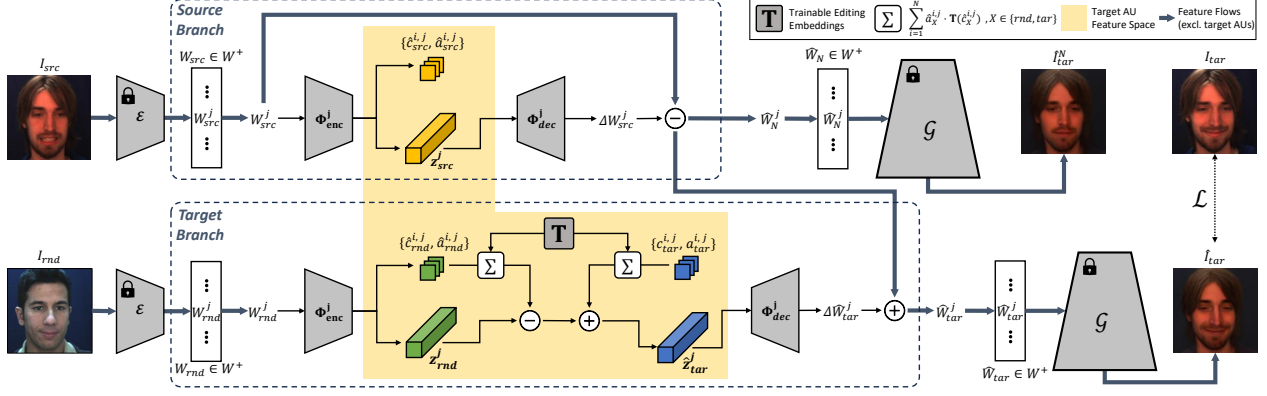


Figure 1. Overall scheme of the proposed AUEditNet. AUEditNet has a dual-branch structure that separately addresses source attribute removal (*Source Branch*) and target attribute addition (*Target Branch*). The *Source Branch* aims at removing the original status in I_{src} , maintaining other attributes and identity while keeping them distinct from the feature space of target facial attributes (highlighted in yellow). The *Target Branch* focuses on determining an edited direction $\Delta \hat{W}_{tar}^j$ for the new status of the target facial attribute, ensuring its independence from identity and other facial attributes. Instead of applying this branch directly to I_{src} , we randomly select another image I_{rnd} , facilitating implicit disentanglement of attributes and identity. The *blue bold arrows* present feature flows excluding the target facial attributes. In this configuration, AUEditNet guarantees that these flows remain outside the embedding space of the target facial attributes.

additional encoding of the latent vector W_{src}^j , one level from the multi-level vectors W_{src} , through a trainable encoder Φ_{enc}^j . Here, j corresponds to the level index, taking into account the disentangled level-wise structure of W_{src} , as outlined in Sec. 3.3. For the purpose of this subsection, we can disregard this index. The outcome of this encoding process is as follows:

$$\Phi_{enc}^j(W_{src}^j) = \{\hat{c}_{src}^{i,j}, \hat{a}_{src}^{i,j}, z_{src}^j\}, i \in [1, N], \quad (1)$$

where N represents the number of facial attributes included in the editing task. In this Eq. 1, $\hat{c}_{src}^{i,j}$ denotes whether the i -th facial attribute exists or not (AU is activated or not); $\hat{a}_{src}^{i,j}$ is the corresponding estimated detailed labels (AU intensities); z_{src}^j is the embedding which acts as a medium for delivering information pertaining to the target facial attributes in this newly encoded space. $\hat{c}_{src}^{i,j}$ would select an editing direction from a globally trainable matrix \mathbf{T} if the i -th facial attribute exists. \mathbf{T} contains N editing directions, each possessing the same dimension as the embeddings. When a specific editing direction $\mathbf{T}(\hat{c}_{src}^{i,j})$ is chosen, we scale it with the estimated labels $\hat{a}_{src}^{i,j}$ to serve as an intensity control. This yields a normalized embedding $z_N^j = z_{src}^j - \sum_{i=1}^N \hat{a}_{src}^{i,j} \cdot \mathbf{T}(\hat{c}_{src}^{i,j})$. Ideally, z_N^j exclusively represents a canonical status of the target facial attribute, free from any person-specific information. While it seems feasible to continue incorporating new target conditions into this normalized embedding for subsequent generation with edited attributes [14, 47], this approach has limitations.

- It cannot ensure the complete exclusion of other attributes or identity features from the normalized embedding.
- Achieving optimal disentanglement of identity from target

attributes requires training data that ideally encompasses as many subjects as possible to attain the desired normalized embedding.

- A loss function is necessary to enforce normalized embeddings identical within a batch, which heavily relies on the batch size and can be resource-intensive.

Given these limitations, instead of directly adding target conditions to the source embedding, our approach adopts a dual-branch structure to physically prevent irrelevant attribute or identity features (indicated by *blue bold arrows*) from infiltrating the feature space of target facial attributes (highlighted in yellow), as illustrated in Fig. 1.

We introduce I_{rnd} through the same processing steps with the shared-weights modules and build a normalized embedding instead of using the source one to compel the network to retain only the target-attribute related information within this encoded space during training. Finally, we introduce new conditions (the existence of the i -th attribute $\hat{c}_{tar}^{i,j}$ and the corresponding detailed labels $\hat{a}_{tar}^{i,j}$). This yields the edited embedding $\hat{z}_{tar}^j = z_{rnd}^j - \sum_{i=1}^N \hat{a}_{rnd}^{i,j} \cdot \mathbf{T}(\hat{c}_{rnd}^{i,j}) + \sum_{i=1}^N \hat{a}_{tar}^{i,j} \cdot \mathbf{T}(\hat{c}_{tar}^{i,j})$. During testing, we directly use source normalized embedding for efficiency considerations.

Source Latent Vectors Editing. For all other facial attributes and identity information, our goal is to preserve them within the original latent space [14]. We input the source embedding z_{src}^j and the edited target embedding \hat{z}_{tar}^j into the decoder Φ_{dec}^j to obtain the residuals ΔW_{src}^j and ΔW_{tar}^j respectively, which are used for editing W_{src}^j . The purpose of ΔW_{src}^j is to capture the source status of the target facial attributes in the input image, while ΔW_{tar}^j stores the new status. Rather than solely assessing the result with the new

status, we propose to supervise both outcomes through the following expressions:

$$\begin{cases} \hat{W}_N^j = W_{src}^j - \Delta W_{src}^j, \\ \hat{W}_{tar}^j = \hat{W}_N^j + \Delta \hat{W}_{tar}^j, \end{cases} \quad (2)$$

where \hat{W}_N^j represents the intermediate editing resulting from the removal of the source status of the aimed facial attributes, and \hat{W}_{tar}^j is the outcome achieved by incorporating the target conditions based on \hat{W}_N^j . After replacing the latent vector at the index j in W_{src} with \hat{W}_N^j (or \hat{W}_{tar}^j), we obtain the final edited latent vectors \hat{W}_N (or \hat{W}_{tar}) for image generation. $\hat{I}_{tar}^N = \mathcal{G}(\hat{W}_N)$ represents a synthesized face image with zero intensities (deactivation) for all AUs, while $\hat{I}_{tar} = \mathcal{G}(\hat{W}_{tar})$ is generated based on the target intensities.

3.3. Multi-level Architecture

The latent space used for editing is the W^+ space [1], compatible with StyleGAN [18]. Latent vectors in W^+ exhibit a multi-level structure, allowing them to control different semantic levels of images [43]. The level is indexed by j and $j \in [1, M]$, where $M \leq 18$ due to the dimension of W^+ . Instead of attempting to reintegrate the disentangled features by levels of the latent vectors in W^+ using a single editing module, we opt for multiple independent editing modules. Each of these modules is responsible for editing a specific level of the latent vectors. However, as we lack information regarding the relationship between the aimed facial attributes and the level index j , supervising the level-wise $\hat{c}_{src}^{i,j}, \hat{a}_{src}^{i,j}$ could be challenging. This is particularly true for the AU intensity editing task, where it's not reasonable to expect that a single level estimated results $\hat{c}_{src}^{i,j}, \hat{a}_{src}^{i,j}$ can accurately represent 12 AU intensities across various areas of the face. Incorrectly estimating the AU intensity of the input image could disrupt the normalization process and consequently affect the final manipulation accuracy. To address this, we draw inspiration from the concept of a latent space for images and propose the creation of another 'latent space' for labels. We use a fully connected network Ψ_{enc} to encode the labels of the target conditions (c_{tar}^i, a_{tar}^i) for the i -th facial attribute into corresponding level-wise pseudo-labels ($c_{tar}^{i,j}, a_{tar}^{i,j}$) to suit the editing needs of each level. Given each level's estimated ($\hat{c}_{src}^{i,j}, \hat{a}_{src}^{i,j}$), we decode them back to the label space, resulting in estimated source labels ($\hat{c}_{src}^i, \hat{a}_{src}^i$) through Ψ_{dec} . The entire process can be summarized as follows:

$$\begin{cases} \Psi_{enc}(c_{tar}^i, a_{tar}^i) = c_{tar}^{i,j}, a_{tar}^{i,j}, \\ \Psi_{dec}(\hat{c}_{src}^{i,j}, \hat{a}_{src}^{i,j}) = \hat{c}_{src}^i, \hat{a}_{src}^i, \end{cases} \quad (3)$$

where $i \in [1, N]$, $j \in [1, M]$, and the subscripts 'src' and 'tar' can be interchanged if we switch the roles of the source and target images during training.

3.4. Objectives

During training, AUEditNet requires source and target images from the same subject. And the random image can be randomly picked from other subjects. We train AUEditNet by minimizing the following loss:

$$\mathcal{L} = \lambda_R \mathcal{L}_R + \lambda_P \mathcal{L}_P + \lambda_F \mathcal{L}_F + \lambda_{ID} \mathcal{L}_{ID} + \lambda_L \mathcal{L}_L. \quad (4)$$

Pixel-wise and Perceptual Losses. We minimize both the pixel-wise loss \mathcal{L}_R and the perceptual loss \mathcal{L}_P [15] between the edited image \hat{I}_{tar} , which is generated based on the provided target conditions, and the actual target image I_{tar} .

$$\begin{aligned} \mathcal{L}_R &= \|\hat{I}_{tar} - I_{tar}\|_2, \\ \mathcal{L}_P &= \|F_{pcept}(\hat{I}_{tar}) - F_{pcept}(I_{tar})\|_2, \end{aligned} \quad (5)$$

where $F_{pcept}(\cdot)$ denotes the perceptual feature extractor.

Pretrained Function Loss. Following the prior works [14, 47], the pretrained function loss \mathcal{L}_F focuses on task-relevant inconsistencies between \hat{I}_{tar} and I_{tar} . The inconsistencies include both intermediate activation feature maps $\{f_k, k \in [1, K]\}$ and estimation results derived from a network $F_{pre}(\cdot)$, which is pretrained on the specific task (e.g. AU intensity estimation).

$$\begin{aligned} \mathcal{L}_F &= \frac{1}{K} \sum_{k=1}^K \|f_k(\hat{I}_{tar}) - f_k(I_{tar})\|_2 \\ &+ \frac{1}{N} \|F_{pre}(\hat{I}_{tar}) - F_{pre}(I_{tar})\|_2, \end{aligned} \quad (6)$$

where K is the number of chosen layers from F_{pre} .

Identity Loss. We restrict the ID similarity between the real image I_{tar} and two generated images $\hat{I}_{tar}; \hat{I}_{tar}^N$ based on the pretrained ArcFace network [5], denoted as F_{id} .

$$\mathcal{L}_{ID} = \sum_{\hat{I} \in \{\hat{I}_{tar}, \hat{I}_{tar}^N\}} 1 - \langle F_{id}(\hat{I}), F_{id}(I_{tar}) \rangle \quad (7)$$

Label Loss. We propose the label loss to supervise the transformation process in the 'latent space' for labels through Ψ_{enc} and Ψ_{dec} as mentioned in Sec. 3.3. Let's assume we use the source image's labels (c_{src}^i and a_{src}^i) as the target conditions for the i -th facial attribute. In this scenario, the generated image should be identical to the source image. This implies that the estimated source embedding z_{src}^j should match the edited embedding \hat{z}_{tar}^j at each level. In other words, the removal of the source status and the addition of the target status should be entirely consistent. As a result, the level-wise conditions ($c_{src}^{i,j}, a_{src}^{i,j}$) in the label's latent space, encoded from the source image's labels (c_{src}^i, a_{src}^i), should align with the estimated conditions ($\hat{c}_{src}^{i,j}, \hat{a}_{src}^{i,j}$). This corresponds to the second part of Eq. 8, which supervises the learning of the encoder Ψ_{enc} . Supp. provides more details.

To further ensure that the level-wise estimation retains the information about the labels of the source image, we utilize the label decoder Ψ_{dec} to guarantee that the estimated results $(\hat{c}_{src}^i, \hat{a}_{src}^i)$ decoded from the level-wise conditions $(\hat{c}_{src}^{i,j}, \hat{a}_{src}^{i,j})$ are consistent with the source image’s labels. Thus, we build the loss as follows:

$$\mathcal{L}_L = \sum_{\xi \in \{c,a\}} \frac{1}{N} \sum_{i=1}^N \left(\|\hat{\xi}_{src}^i - \xi_{src}^i\|_2 + \frac{1}{M} \sum_{j=1}^M \|\hat{\xi}_{src}^{i,j} - \xi_{src}^{i,j}\|_2 \right). \quad (8)$$

In summary, the first three terms in Eq. 4 ensure the generated image’s similarity to the target image. The fourth term enforces identity consistency, and the final term supervises the learning of level-wise pseudo-labels for avoiding entangled labels as conditions and improving the attribute editing performance. The hyperparameters $\lambda_R, \lambda_P, \lambda_F, \lambda_{ID}, \lambda_L$ enable a balanced learning from these various losses.

4. Experiments

4.1. Implementation Details

We employed e4e [39] and StyleGAN2 [18] as the GAN inversion pair. We designed a Siamese network for the external AU intensity estimation for the pretrained function loss in Eq. 6. This network takes a pair of face images from the same subject as the input and estimates the intensity difference of AUs between these two images. This design reduces the impact of subject-specific facial attributes. During training, we used the convolutional part of VGG-16 [38] as the backbone to build the AU intensity estimation network F_{pre} . During test, we used a separate estimator H_{est} , which has the same architecture with ResNet-50 [12] as the backbone. Importantly, this estimator H_{est} was never exposed to the training phase but was trained on the same training dataset.

We trained AUEditNet using the DISFA training subset [23, 24]. DISFA comprises of 27 subjects and provides multi-level integral intensities for 12 AUs, offering annotations for the largest number of AUs among publicly available datasets for AU intensity estimation. We used 18 subjects for training and 9 subjects for testing, following the data split used in [19, 34]. To assess AUEditNet, we used the DISFA test subset to evaluate its accuracy in manipulating AU intensities while preserving other attributes. Furthermore, we expanded our evaluation to encompass facial expressions, beyond AUs alone, by using the BU-4DFE dataset [46]. Our evaluation involved tasks related to expression transfer and data augmentation for AU intensity estimation. For further assessment of out-of-domain editing performance, we incorporated CelebA-HQ [16] and FFHQ [17], which both are the benchmarks for the high-quality human face image datasets.

4.2. Evaluation Criteria

We assess the performance of AUEditNet by examining the comparison between the generated image \hat{I}_{tar} and the target image I_{tar} from four perspectives: the accuracy of intensity editing in AUs, identity preservation, image similarity, and smile expression manipulation (illustrated in Sec. 4.4).

Accuracy of AU Intensity Manipulation. We quantify the AU intensity manipulation performance in edited images by using the external pretrained ResNet-50 based estimator H_{est} , which is unseen during training. We report the Intra-Class Correlation (specifically ICC(3,1) [37]) and mean squared error (MSE), both calculated for 12 AUs, between the estimated values $H_{est}(\hat{I}_{tar})$ or $H_{est}(\hat{I}_{tar}^N)$ and their intended target values.

Identity Preservation. A well-trained image editor should consistently maintain the identity given various provided conditions. To assess the similarity of identity, we measure the distance of embeddings between \hat{I}_{tar} and I_{tar} to assess the similarity of identity, where the embedding is extracted by a pretrained face recognition model [9].

Image Similarity. We employ two metrics: pixel-wise mean squared error and the Learned Perceptual Image Patch Similarity (LPIPS) [45] to measure the image similarity between \hat{I}_{tar} and I_{tar} .

4.3. Qualitative Evaluation

Within-Dataset Evaluation Fig. 2 illustrates a qualitative comparison of AU intensity manipulation based on provided target conditions. Both ReDirTrans [14] and our proposed AUEditNet employ a two-step editing process to prevent potential attribute status mixing. After the source status removal, the generated images should exhibit all AU intensities set to zero, serving as benchmarks when all AUs are deactivated. ReDirTrans and AUEditNet demonstrate the ability to learn the desired AU movements, under both cases when deactivating all AUs or assigning new target intensities. However, ReDirTrans fails to preserve identity information in intermediate and final generated images. Additionally, ReDirTrans attempts to address color discrepancy between real and inverted images during AU editing, resulting in undesired color distortion in images. In contrast, AUEditNet focuses only on editing the aimed AUs’ intensities, devoid of unrelated information, which is achieved through the dual-branch architecture. On the other hand, DeltaEdit [21] excels in maintaining identity information and other facial attributes. However, it is limited to learning noticeable AU movements and may ignore subtle motions such as eyebrow, cheek, and lip corner movements, potentially causing significant changes in the entire facial expression. AUEditNet successfully achieves accurate AU intensity editing under this two-phase editing process while maintaining identity.

Cross-Dataset Evaluation Fig. 3 presents the cross-dataset results involving single AU editing with multiple




Figure 2. Comparison of AU intensity manipulation using target AU intensities in DISFA. AUEditNet, ReDirTrans generate editing results that involve the removal (−) of source attributes and the addition (+) of target attributes. DeltaEdit uses intensity differences between source and target images for attribute addition (+ Δ). The removal (−) process yields ‘neutral-like’ face images with all AU intensities set to zero.




Figure 3. Cross-dataset evaluation of single AU intensity manipulation in CelebA-HQ. The descriptions of AUs (from top to bottom) are Outer Brow Raiser, Brow Lowerer, Upper Lid Raiser, Lip Corner Depressor, and Lips Part. a_{tar} represents the target intensity.

intensity levels on the CelebA-HQ dataset [16]. AUEditNet exhibits the capability to achieve consecutive AU intensity manipulation. Notably, even in the absence of negative intensities during training, AUEditNet produces reasonable

editing outcomes. For instance, applying negative intensity to AU 5 (Upper Lid Raiser) results in a generated image with partial eye closure. Regarding AU 25 (lips part), where intensity indicates mouth openness, providing negative intensity still maintains the closed configuration, aligning with the case of zero intensity instead of creating unrealistic results.

4.4. Quantitative Evaluation

Accuracy of AU Intensity Editing. Table 1 presents measurements of ICC and MSE for comparing the estimated AU intensities against ground truth. We categorize the methods under each evaluation metric based on their research directions, whether they focus on the estimation or editing of AU intensities in images. Among the editing methods, our proposed AUEditNet surpasses state-of-the-art facial attribute editing methods, especially in terms of the average performance across all 12 AUs. When it comes to the performance of deactivating all AUs, AUEditNet achieves a substantial 38.92% improvement in MSE compared to ReDirTrans [14]. This illustrates the complete and accurate attribute removal process, which, in turn, contributes to enhanced final performance since attribute removal and addition are entirely reversible processes with shared trainable parameters.

Furthermore, we expand our comparison to include both editing and estimation methods because the editing performance is also assessed using the same AU intensity estimation process. Moreover, the external estimator H_{est} is trained on the same data as AU intensity estimation methods. We still observe that the estimation performance, when evaluated with our edited face images, surpasses that of state-

	Method	AU1	AU2	AU4	AU5	AU6	AU9	AU12	AU15	AU17	AU20	AU25	AU26	Avg
ICC(3, 1) (\uparrow)	HR [26]	.56	.52	.75	.42	.51	.55	.82	.55	.37	.21	.93	.62	.57
	Aps [32]	.35	.19	.78	.73	.52	.65	.81	.49	.61	.28	.92	.67	.58
	MAE-Face [22]	.740	.688	.754	.666	.653	.584	.877	.527	.589	.331	.952	.721	.674
	DeltaEdit [21]	.091	.058	.114	.034	.383	.065	.694	.008	.004	.041	.581	.166	.179
	ReDirTrans [14]	.856	.631	.851	.436	.634	.278	.862	.364	.602	.481	.927	.480	.617
	AUEditNet	.848	.559	.874	.600	.577	.230	.890	.276	.669	.511	.950	.548	.628
MSE (\downarrow)	HR [26]	.41	.37	.70	.08	.44	.30	.29	.14	.26	.16	.24	.39	.32
	Aps [32]	.68	.59	.40	.03	.49	.15	.26	.13	.22	.20	.35	.17	.30
	MAE-Face [22]	.200	.186	.514	.032	.320	.222	.221	.093	.204	.146	.164	.260	.213
	DeltaEdit [21]	.605	.686	1.311	.031	.513	.485	.570	.080	.424	.454	1.157	.420	.561
	ReDirTrans [14]	.181	.397	.341	.034	.453	.552	.286	.070	.225	.333	.247	.367	.290
	AUEditNet	.191	.445	.309	.029	.492	.579	.228	.080	.188	.322	.169	.367	.283
	ReDirTrans (<i>N</i>)	.045	.117	.025	.019	.024	.009	.300	.032	.177	.032	.803	.427	.167
	AUEditNet (<i>N</i>)	.069	.101	.098	.024	.036	.006	.227	.004	.014	.063	.351	.228	.102

Table 1. Comparison to the state-of-the-art action unit (AU) intensity estimation and editing methods on DISFA [24]. The ‘Method’ column under each metric is categorized into two parts: 1. Upper part: AU intensity estimation methods; 2. Lower part: AU intensity editing methods. In the estimation task, we evaluate the performance by comparing the estimated intensities of the input image to the ground truth. For the editing task, the procedure begins with the editing of the input image based on the target conditions. Then we acquire the estimated AU intensities from the edited image via the external pretrained estimator H_{est} . Finally, we compare these estimated intensities with the provided target conditions. ‘(*N*)’ denotes the results obtained after the source attribute removal, where all AU intensities are set to zero. Each group’s best result is highlighted in bold. Without extra facial data, MAE-Face becomes MAE-IN1k [22], leading to ICC dropping to .599.

Method	Target Image	Identity Preservation	Image Similarity	
		Distance (\downarrow)	L2 (\downarrow)	LPIPS (\downarrow)
GAN Inversion [39]	Real	.368	.025	.173
	Inverted	.278	.011	.065
DeltaEdit [21]	Real	.396	.022	.165
	Inverted	.309	.011	.074
ReDirTrans [14]	Real	.505	[.024]	.175
	Inverted	.479	.018	.153
AUEditNet	Real	[.468]	.026	[.174]
	Inverted	[.435]	[.016]	[.126]

Table 2. Comparison of identity preservation and image similarity in facial attribute editing methods. ‘GAN Inversion’ as a baseline illustrates that the accuracy of action unit intensity editing cannot be reflected in the performance of the Image Similarity criteria. The best performance is indicated in bold, while the second best is highlighted within brackets.

of-the-art AU intensity estimation methods on the DISFA test subset [23, 24]. This finding further solidifies the high level of consistency between the provided target intensities and the edited images generated by AUEditNet.

Identity Preservation and Image Similarity. Table 2 summarizes the performance of identity preservation and image similarity given image editing results. In addition to comparing the edited images with the real target images, we also conduct a comprehensive comparison using GAN-

inverted images as the target. All three editing methods focus on the latent code editing, without adjusting the image encoder and generator. From the identity perspective, DeltaEdit [21] achieves the best performance, nearly matching the GAN inversion performance. However, this is at the cost of AU intensity manipulation accuracy, resulting in a decline of 71.50% in ICC and 98.23% in MSE compared to AUEditNet. Comparing our AUEditNet with ReDirTrans [14], we observe the identity preservation improvements of 7.33% and 9.19% considering real and inverted images, respectively. These results further validate the effectiveness of our method’s ability to achieve disentanglement and preserve identity during intensity manipulation.

Regarding image similarity, DeltaEdit [21] continues to outperform the other two editing methods. However, when using the GAN inversion as the baseline to compare the inverted source image with the real or inverted target images separately, we find that the image similarity criteria still maintain good performance, even when dealing with different AU intensities between source and target images. In other words, the difference in AU intensities is not reflected over the image similarity. When compared to ReDirTrans [14], AUEditNet achieves comparable performance with the real target image and achieves better performance with the inverted one. These results further demonstrate AUEditNet’s disentanglement ability when achieving AU intensity editing.

Smile Manipulation. We evaluate smile attribute manipulation using metrics proposed in [7] on the FFHQ dataset

Method	Smile Attribute	
	E_d (\downarrow)	ρ (\uparrow)
Talk-to-Edit [13]	0.212	40.9
StyleFlow [2]	0.099	88.9
Do <i>et al.</i> [7] (W/StyleGAN2)	0.103	96.9
AUEditNet	0.099	121.3

Table 3. Comparison of smile intensity manipulation performance on the FFHQ test dataset. AUEditNet achieves the best performance given identity preservation (E_d) and manipulation efficiency (ρ).

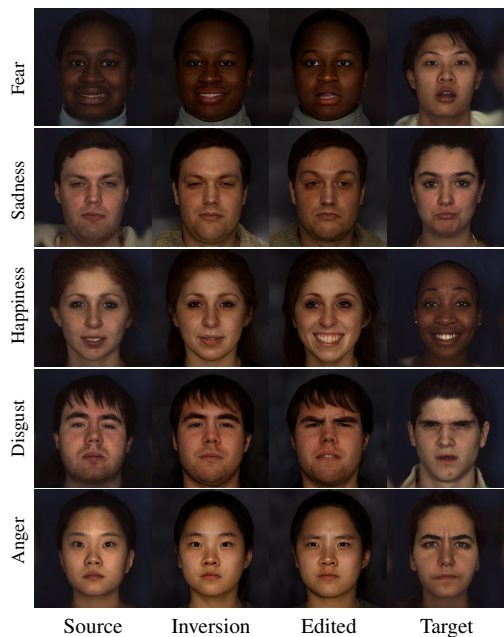


Figure 4. AU intensity manipulation conditioned on target images to achieve facial expression transfer on the BU-4DFE dataset. The fine-grained facial expressions, such as AU 17 (Chin Raiser) in ‘Sadness’ and AU 25 (Lips Part) in ‘Disgust’, are transferred accurately.

[17]. Specifically, we modify the intensities of AU 6 (Cheek Raiser) and AU 12 (Lip Corner Puller) across eight levels simultaneously to enable smile intensity editing [10]. Table 3 provides comparisons based on identity preservation (E_d) and manipulation efficiency (ρ). The result indicates that AUEditNet better preserves identity when an attribute undergoes the same quantity of change than others.

4.5. Expression Transfer

Setting individual AU values is a cumbersome process and requires expertise for achieving desired expression synthesis [27]. In contrast, our proposed AUEditNet demonstrates the capability to directly transfer facial expressions from target images without the need for retraining the network. The process involves inputting the target image with the desired expression into the target branch in Fig. 1. Instead of employing removal and addition processes, we directly

Model	Target (Removal & Addition)			Neutral (Removal)		
	MSE \downarrow	ICC \uparrow	ID \downarrow	MSE \downarrow	ID \downarrow	
Training	w/o \mathcal{L}_R & \mathcal{L}_P	0.388	0.584	0.502	0.253	0.440
	w/o \mathcal{L}_F	0.507	0.356	0.480	0.467	0.454
	w/o \mathcal{L}_{ID}	0.288	0.619	0.533	0.115	0.515
Design	Sngl.	0.317	0.598	0.545	0.621	0.724
	+ Encoder, Decoder	0.290	0.617	0.505	0.167	0.600
	+ Dual.	0.288	0.617	0.471	0.111	0.439
	+ Label Mapping	0.283	0.628	0.468	0.102	0.426

Table 4. Ablation Study for AUEditNet.

feed the estimated embeddings of the target image into the decoder Φ_{dec}^j , similar to the procedure in the source branch, to acquire editing residuals with target facial expressions. Fig. 4 shows expression transfer results on the BU-4DFE dataset [46]. The edited images demonstrate the contributions of AU intensity manipulation to the facial expression reenactment.

4.6. Ablation Study

Table 4 shows the results of ablation studies for AUEditNet. In **module design**, the integration of dual branch (+ Dual.) leads to improvements in both AU manipulation accuracy (MSE, ICC) and ID preservation (ID). Notably, in the ‘Removal’ case for neutral face generation, MSE and ID get 33.5% and 26.8% improvements, respectively. The evaluation in this removal-only process is valuable for assessing whether unrelated information is introduced into the target AU space during editing, which is often invisible when ‘Removal and Addition’ processes are implemented. Level-wise label mapping can further improve manipulation accuracy. Regarding **training loss**, a well-trained AU intensity estimator (\mathcal{L}_F) plays a more crucial role than a paired target image (\mathcal{L}_R & \mathcal{L}_P). This observation aligns with the fact that pixel-wise MSE and perceptual loss may not effectively capture AU motions. The absence of ID loss leads to a performance drop in ID. However, it also loosens constraints on latent code editing, resulting in more accurate AU manipulation.

5. Conclusion

In this work, we achieved accurate AU intensity manipulation in high-resolution synthetic face images. Our method allows conditioning manipulation on intensity values or target images without retraining the network or requiring extra estimators. This pipeline presents a promising solution for editing facial attributes despite the dataset’s limited subject count. We validated our method both qualitatively and quantitatively through extensive experiments. The performance boost with synthetic augmented data confirms the quality of generated samples, mitigating the challenge of data scarcity. In the future, we aim to explore weakly- (or self-) supervised methods to further advance AU intensity manipulation.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4432–4441, 2019. 1, 4
- [2] Rameen Abdal, Peihao Zhu, Niloy J Mitra, and Peter Wonka. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (ToG)*, 40(3):1–21, 2021. 8
- [3] Tadas Baltrusaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 59–66. IEEE, 2018. 1
- [4] Yusuf Dalva, Said Fahri Altındış, and Aysegül Dundar. Vecgan: Image-to-image translation with interpretable latent directions. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVI*, pages 153–169. Springer, 2022. 1, 2
- [5] 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. 4
- [6] Hui Ding, Kumar Sricharan, and Rama Chellappa. Exprgan: Facial expression editing with controllable expression intensity. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 1
- [7] Hoseok Do, EunKyung Yoo, Taehyeong Kim, Chul Lee, and Jin Young Choi. Quantitative manipulation of custom attributes on 3d-aware image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8529–8538, 2023. 1, 2, 7, 8
- [8] Paul Ekman and Wallace V Friesen. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*, 1978. 1
- [9] Adam Geitgey. Github – face recognition. https://github.com/ageitgey/face_recognition/, 2021. 5
- [10] Jeffrey M Girard, Gayatri Shandar, Zhun Liu, Jeffrey F Cohn, Lijun Yin, and Louis-Philippe Morency. Reconsidering the duchenne smile: indicator of positive emotion or artifact of smile intensity? In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 594–599. IEEE, 2019. 8
- [11] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems*, 33:9841–9850, 2020. 1, 2
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [13] Yuming Jiang, Ziqi Huang, Xingang Pan, Chen Change Loy, and Ziwei Liu. Talk-to-edit: Fine-grained facial editing via dialog. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13799–13808, 2021. 8
- [14] Shiwei Jin, Zhen Wang, Lei Wang, Ning Bi, and Truong Nguyen. Redirtrans: Latent-to-latent translation for gaze and head redirection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5547–5556, 2023. 2, 3, 4, 5, 6, 7
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 694–711. Springer, 2016. 4
- [16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 5, 6
- [17] 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. 5, 8
- [18] 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. 1, 2, 4, 5
- [19] Wei Li, Farnaz Abtahi, Zhigang Zhu, and Lijun Yin. Eac-net: Deep nets with enhancing and cropping for facial action unit detection. *IEEE transactions on pattern analysis and machine intelligence*, 40(11):2583–2596, 2018. 5
- [20] Jun Ling, Han Xue, Li Song, Shuhui Yang, Rong Xie, and Xiao Gu. Toward fine-grained facial expression manipulation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*, pages 37–53. Springer, 2020. 1, 2
- [21] Yueming Lyu, Tianwei Lin, Fu Li, Dongliang He, Jing Dong, and Tieniu Tan. Deltaedit: Exploring text-free training for text-driven image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6894–6903, 2023. 1, 2, 5, 7
- [22] Bowen Ma, Rudong An, Wei Zhang, Yu Ding, Zeng Zhao, Rongsheng Zhang, Tangjie Lv, Changjie Fan, and Zhipeng Hu. Facial action unit detection and intensity estimation from self-supervised representation. *IEEE Transactions on Affective Computing*, pages 1–15, 2024. 7
- [23] S Mohammad Mavadati, Mohammad H Mahoor, Kevin Bartlett, and Philip Trinh. Automatic detection of non-posed facial action units. In *2012 19th IEEE International Conference on Image Processing*, pages 1817–1820. IEEE, 2012. 5, 7
- [24] S Mohammad Mavadati, Mohammad H Mahoor, Kevin Bartlett, Philip Trinh, and Jeffrey F Cohn. Disfa: A spontaneous facial action intensity database. *IEEE Transactions on Affective Computing*, 4(2):151–160, 2013. 5, 7
- [25] William S Noble. What is a support vector machine? *Nature biotechnology*, 24(12):1565–1567, 2006. 2
- [26] Ioanna Ntinou, Enrique Sanchez, Adrian Bulat, Michel Valstar, and Yorgos Tzimiropoulos. A transfer learning approach to heatmap regression for action unit intensity estimation. *IEEE Transactions on Affective Computing*, 2021. 7

- [27] Foivos Paraperas Papantoniou, Panagiotis P Filntisis, Petros Maragos, and Anastasios Roussos. Neural emotion director: Speech-preserving semantic control of facial expressions in “in-the-wild” videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18781–18790, 2022. 8
- [28] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. 2
- [29] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *Proceedings of the European conference on computer vision (ECCV)*, pages 818–833, 2018. 1, 2
- [30] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: One-shot anatomically consistent facial animation. *International Journal of Computer Vision*, 128:698–713, 2020. 2
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [32] Enrique Sanchez, Mani Kumar Tellamekala, Michel Valstar, and Georgios Tzimiropoulos. Affective processes: stochastic modelling of temporal context for emotion and facial expression recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9074–9084, 2021. 7
- [33] Klaus R Scherer. Emotion as a process: Function, origin and regulation, 1982. 1
- [34] Zhiwen Shao, Zhilei Liu, Jianfei Cai, and Lizhuang Ma. Jaanet: joint facial action unit detection and face alignment via adaptive attention. *International Journal of Computer Vision*, 129:321–340, 2021. 5
- [35] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1532–1540, 2021. 2
- [36] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 1, 2
- [37] Patrick E Shrout and Joseph L Fleiss. Intraclass correlations: uses in assessing rater reliability. *Psychological bulletin*, 86(2):420, 1979. 5
- [38] K Simonyan and A Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations (ICLR 2015)*. Computational and Biological Learning Society, 2015. 5
- [39] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 5, 7
- [40] Soumya Tripathy, Juho Kannala, and Esa Rahtu. Icfac: Interpretable and controllable face reenactment using gans. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3385–3394, 2020. 1, 2
- [41] Soumya Tripathy, Juho Kannala, and Esa Rahtu. Facegan: Facial attribute controllable reenactment gan. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1329–1338, 2021. 2
- [42] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *International conference on machine learning*, pages 9786–9796. PMLR, 2020. 2
- [43] Ceyuan Yang, Yujun Shen, and Bolei Zhou. Semantic hierarchy emerges in deep generative representations for scene synthesis. *International Journal of Computer Vision*, 129:1451–1466, 2021. 4
- [44] Oğuz Kaan Yüksel, Enis Simsar, Ezgi Gülperi Er, and Pinar Yanardag. Latentclr: A contrastive learning approach for unsupervised discovery of interpretable directions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14263–14272, 2021. 2
- [45] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 5
- [46] Xing Zhang, Lijun Yin, Jeffrey F Cohn, Shaun Canavan, Michael Reale, Andy Horowitz, and Peng Liu. A high-resolution spontaneous 3d dynamic facial expression database. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*, pages 1–6. IEEE, 2013. 5, 8
- [47] Yufeng Zheng, Seonwook Park, Xucong Zhang, Shalini De Mello, and Otmar Hilliges. Self-learning transformations for improving gaze and head redirection. *Advances in Neural Information Processing Systems*, 33:13127–13138, 2020. 3, 4
- [48] Yiming Zhu, Hongyu Liu, Yibing Song, Ziyang Yuan, Xintong Han, Chun Yuan, Qifeng Chen, and Jue Wang. One model to edit them all: Free-form text-driven image manipulation with semantic modulations. *Advances in Neural Information Processing Systems*, 35:25146–25159, 2022. 2