

Reference Guided Face Component Editing

Qiyao Deng^{1,4}, Jie Cao^{1,4}, Yunfan Liu^{1,4}, Zhenhua Chai⁵, Qi Li^{1,2,4*}, Zhenan Sun^{1,3,4}

¹Center for Research on Intelligent Perception and Computing, NLPR, CASIA, Beijing, China

²Artificial Intelligence Research, CAS, Qingdao, China

³Center for Excellence in Brain Science and Intelligence Technology, CAS, Beijing, China

⁴School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

⁵Vision Intelligence Center, AI Platform, Meituandianping Group

{qiyao.deng, jie.cao, yunfan.liu}@cripac.ia.ac.cn, {qli, znsun}@nlpr.ia.ac.cn,
chaizhenhua@meituan.com}

Abstract

Face portrait editing has achieved great progress in recent years. However, previous methods either 1) operate on pre-defined face attributes, lacking the flexibility of controlling shapes of high-level semantic facial components (e.g., eyes, nose, mouth), or 2) take manually edited mask or sketch as an intermediate representation for observable changes, but such additional input usually requires extra efforts to obtain. To break the limitations (e.g. shape, mask or sketch) of the existing methods, we propose a novel framework termed r-FACE (Reference Guided FAce Component Editing) for diverse and controllable face component editing with geometric changes. Specifically, r-FACE takes an image inpainting model as the backbone, utilizing reference images as conditions for controlling the shape of face components. In order to encourage the framework to concentrate on the target face components, an example-guided attention module is designed to fuse attention features and the target face component features extracted from the reference image. Both qualitative and quantitative results demonstrate that our model is superior to existing literature.

1 Introduction

Face portrait editing is of great interest in the computer vision community due to its potential applications in movie industry, photo manipulation, and interactive entertainment, etc. With advances in Generative Adversarial Networks [Goodfellow *et al.*, 2014] in recent years, tremendous progress has been made in face portrait editing [Yang *et al.*, 2018; Choi *et al.*, 2018; Liu *et al.*, 2019]. These approaches generally fall into three main categories: label-conditioned methods, geometry-guided methods and reference-guided methods. Specifically, label-conditioned methods [He *et al.*, 2019; Choi *et al.*, 2018] only focus on several pre-defined conspicuous attributes thus lacking the flexibility of controlling shapes

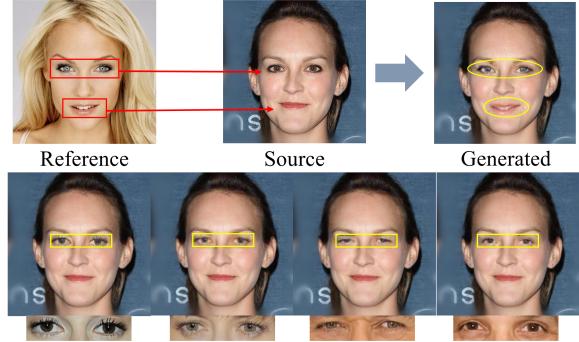


Figure 1: The illustration of reference guided face component editing. The first row is the definition diagram of the task, and the second row is the synthesized result based on the given different reference images.

of high-level semantic facial components (e.g., eyes, nose, mouth). This is because it is hard to produce results with observable geometric changes merely based on attribute labels. In order to tackle this, geometry-guided methods [Jo and Park, 2019; Gu *et al.*, 2019] propose to take manually edited mask or sketch as an intermediate representation for obvious face component editing with large geometric changes. However, directly taking such precise representations as a shape guide is inconvenient for users, which is laborious and requires painting skills. To solve this problem, reference-guided methods directly learns shape information from reference images without requiring precise auxiliary representation, relieving the dependence on face attribute annotation or precise sketch/color/mask information. As far as we know, reference-guided methods are less studied than the first two methods. ExGANs [Dolhansky and Canton Ferrer, 2018] utilizes exemplar information in the form of a reference image of the region for eye editing (in-painting). However, ExGANs can only edit eyes and requires reference images with the same identity, which is inconvenient to collect in practice. ELEGANT [Xiao *et al.*, 2018] transfers exactly the same type of attributes from a reference image to the source image by exchanging certain part of their encodings. However, ELEGANT is only used for editing obvious semantic attributes, and could not change abstract shapes.

To overcome the aforementioned problems, we propose a

*Contact Author

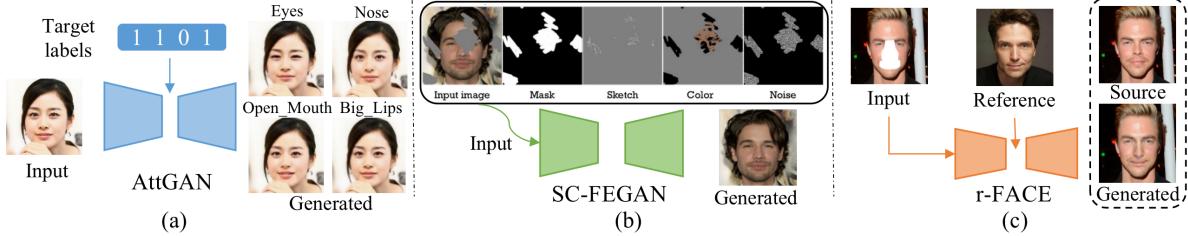


Figure 2: Different methods for face portrait editing. (a) AttGAN, (b) SC-FEGAN and (c) Our network.

new framework: **R**eference **G**uided **F**ace **C**omponent **E**diting (r-FACE for short), which can achieve diverse and controllable face semantic components editing (e.g., eyes, nose, mouth) without requiring paired images. The ideal editing is to transfer single or multiple face components from a reference image to the source image, while still preserving the consistency of pose, skin color and topological structure(see Figure 1). Our framework breaks the limitations of existing methods: 1) shape limitation. r-FACE can flexibly control diverse shapes of high-level semantic facial components by different reference images; 2) intermediate presentation limitation. There is no need to manually edit precise masks or sketches for observable geometric changes.

Our framework is based on an image inpainting model for editing face components by reference images even without paired images. r-FACE has two main streams including 1) an inpainting network \mathcal{G}_i and 2) an embedding network \mathcal{E}_r . As shown in Figure 3, \mathcal{G}_i takes the source image with target face components corrupted and the corresponding mask image as input, and outputs the generated image with semantic features extracted by \mathcal{E}_r . To encourage the framework to concentrate on the target face components, an example-guided attention module is introduced to combine features extracted by \mathcal{G}_i and \mathcal{E}_r . To supervise the proposed model, a contextual loss is adopted to constrain the similarity of shape between generated images and reference images, while a style loss and a perceptual loss are adopted to preserve the consistency of skin color and topological structure between generated images and source images. Both qualitative and quantitative results demonstrate that our model is superior to existing literature by generating high-quality and diverse faces with observable changes for face component editing.

In summary, the contributions of this paper are as follows:

- We propose a novel framework named reference guided face component editing for diverse and controllable face component editing with geometric changes, which breaks the shape and intermediate presentation (e.g., precise masks or sketches) limitation of existing methods.
- An example-guided attention module is designed to encourage the framework to concentrate on the target face components by combining attention features and the target face component features of the reference image, further boosting the performance of face portrait editing.
- Both qualitative and quantitative results demonstrate the superiority of our method compared with other benchmark methods.

2 Related Work

Face Portrait Editing. Face portrait editing aims at manipulating single or multiple attributes or components of a face image towards given conditions. Depending on different conditions, face portrait editing methods can be classified into three categories: label-conditioned methods, geometry-guided methods and reference-guided methods. Label-conditioned methods change predefined attributes, such as hair color [Choi *et al.*, 2018], age [Liu *et al.*, 2019] and pose [Cao *et al.*, 2019]. However, these methods focus on several conspicuous attributes [Liu *et al.*, 2015; Langner *et al.*, 2010], lacking the flexibility of controlling the shapes of different semantic facial parts. As shown in Figure 2(a), AttGAN [He *et al.*, 2019] attempts to edit attributes with shape changes, such as '*Narrow_Eyes*', '*Pointy_Nose*' and '*Mouth_Slightly_Open*', but it can only achieve subtle changes hard to be observed. Moreover, lacking of labeled data will extremely limit the performance of these methods. To tackle above problems, geometry-guided methods use an intermediate representation to guide observable shape changes. [Gu *et al.*, 2019] proposes a framework based on mask-guided conditional GANs which can change the shape of face components by manual editing precise masks. As shown in Figure 2(b), SC-FEGAN [Jo and Park, 2019] requires directly taking mask, precise sketch and color as input for editing the shape of face components. However, such precise input is difficult and inconvenient to obtain. Reference-guided methods can directly learn shape information from reference images without precise auxiliary representation, relieving the dependence on face attribute annotation or precise sketch/color/mask information for face portrait editing. Inspired by this, we propose a new framework (see Figure 2(c)), which can achieve diverse and controllable face semantic components editing (e.g., eyes, nose, mouth), which is shape free and precise landmark or sketch free.

Face Completion/Inpainting. Face completion, also known as face inpainting, aims to complete a face image with a masked region or missing content. Early face completion works [Bertalmio *et al.*, 2000; Criminisi *et al.*, 2003; Bertalmio *et al.*, 2003] fill semantic contents based on the overall image and structural continuity between the masked and unmasked regions, which aims to reconstruct missing regions according to the ground-truth image. Recently, some learning-based method [Zheng *et al.*, 2019; Song *et al.*, 2019] are proposed for generating multiple and diverse plausible results. [Zheng *et al.*, 2019] proposes a probabilistically principled framework with two parallel paths, the VAE-based reconstructive path is used to impose smooth priors for the la-

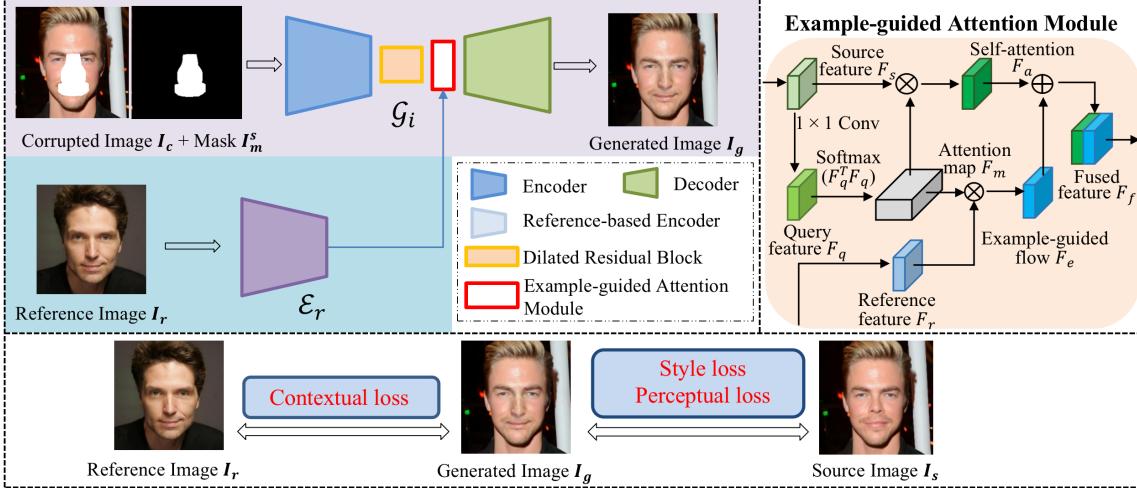


Figure 3: The overall structure of proposed framework. On the top left corner is the generator. The top-right figure shows detailed attention module. The constraints among the source image, the reference image and the generated image is shown on the bottom.

tent space of complement regions, the generative path is used to predict the latent prior distribution for missing regions, which can be sampled to generate diverse results. However, this method lacks controllability for diverse results. In light of this, [Song *et al.*, 2019] generate controllable diverse results from the same masked input by manual modifying facial landmark heatmaps and parsing maps.

3 Method

In this section, we first introduce the framework of reference-based face component editing. Then, the example-guided attention module are presented. Finally, objective functions of the proposed model are provided.

3.1 Reference Guided Face Component Editing

We propose a framework (See Figure 3), named reference guided face component editing, that transfers one or multiple face components of a reference image to corresponding components of the source image. The framework requires three inputs, a source image I_s , a reference image I_r and the target face component mask of the reference image I_m^s . The source mask I_m^s merely needs to roughly represent the target face components, which can be obtained by a face parsing or landmark detection network. The corrupted image can be obtained by Equation 1:

$$I_c = I_s * I_m^s, \quad (1)$$

where $*$ is an element-wise multiplication operator. The goal of this framework is to generate a photo-realistic image I_g , in which shape is semantically similar to corresponding face components of the reference image while the face color and topological structure are consistent with the source image.

In this work, we utilize an image inpainting generator G_i as backbone that can generate the completed image without constraints on shape, while a discriminator D is used for distinguishing face images from the real and synthesized domains. G_i is consist of an encoder, seven dilated residual blocks, an attention module and a decoder. To fill missing parts with semantically meaningful face components of a reference image,

a reference-guided encoder E_r is introduced to extract features of the reference image. The encoder of G_i and E_r have same structure but parameters are not shared. Attention module, to be described next, is effectively transferring semantic components from high-level features of the reference image to G_i , further improving the performance of our framework. The generated image I_g can be expressed as:

$$I_g = G_i(I_c, I_m^s, E_r(I_r)), \quad (2)$$

3.2 Example-guided Attention Module

Inspired by the short+long term attention of PICNet [Zheng *et al.*, 2019], we propose an example-guided attention. The short+long term attention uses the self-attention map to harness distant spatial context and the contextual flow to capture feature-feature context for finer-grained image inpainting. The example-guided attention replaces the contextual flow by the example-guided flow which combines the attention features and the reference features for clearly transforming the corresponding face component features of reference images to source images.

The proposed structure is shown in the upper right corner of Figure 3. Following [Zheng *et al.*, 2019], the self-attention map is calculated from the source feature, which can be expressed as $F_a = F_s \otimes F_m$. The attention map F_m is obtained by $F_m = \text{Softmax}(F_q^T F_q)$, in which $F_q = \text{Conv}(F_s)$ and Conv is a 1×1 convolution filter. To transfer the target face component features of reference images, the reference feature is embedded by multiplying the attention map with the source mask I_m^s . The example-guided flow is expressed as follows:

$$F_e = I_m^s * F'_r + (1 - I_m^s) * F_r, \quad (3)$$

where $F'_r = F_r \otimes F_m$. Finally, the fused feature $F_f = F_a \oplus F_e$ is sent to the decoder for generating results with the target components of the reference image.

3.3 Objective

We use the combination of a per-pixel loss, a style loss, a perceptual loss, a contextual loss, a total variation loss and an adversarial loss for training the framework.

To capture fine facial details we adopt the perceptual loss [Johnson *et al.*, 2016] and the style loss [Johnson *et al.*, 2016], which are widely adopted in style transfer, super resolution, and face synthesis. The perceptual loss aims to measure the similarity of the high dimensional features (e.g., overall spatial structure) between two images, while the style loss measure the similarity of styles (e.g., colors). The perceptual loss can be expressed as follows:

$$\mathcal{L}_{perc} = \sum_l \frac{1}{C_l H_l W_l} \|\phi_l(\mathbf{I}_g) - \phi_l(\mathbf{I}_s)\|_1, \quad (4)$$

The style loss compare the content of two images by using Gram matrix, which can be expressed as follows:

$$\mathcal{L}_{style} = \sum_l \frac{1}{C_l C_l} \left\| \frac{G_l(\mathbf{I}_g * \mathbf{I}_m^s) - G_l(\mathbf{I}_c)}{C_l H_l W_l} \right\|_1, \quad (5)$$

where $\|\cdot\|_1$ is the ℓ_1 norm. $\phi_l(\cdot) \in \mathbb{R}^{C_l \times H_l \times W_l}$ represents the feature map of the l -th layer of the VGG-19 network [Simonyan and Zisserman, 2014] pretrained on the ImageNet. $G_l(\cdot) = \phi_l(\cdot)^T \phi_l(\cdot)$ represents the Gram matrix corresponding to $\phi_l(\cdot)$.

The contextual loss [Mechrez *et al.*, 2018] measures the similarity between non-aligned images, which in our model effectively guarantees the consistent shape of the target face components in generated images and reference images. Given an image x and its target image y , each of which is represented as a collection of points(e.g. VGG-19 [Simonyan and Zisserman, 2014] features): $X = \{x_i\}$ and $Y = \{y_j\}$, $|X| = |Y| = N$. The similarity between the images can be calculated that for each feature y_j , finding the feature x_i that is most similar to it, and then sum the corresponding feature similarity values over all y_j . Formally, it is defined as:

$$CX(x, y) = CX(X, Y) = \frac{1}{N} \sum_j \max_i CX_{ij}, \quad (6)$$

where CX_{ij} is the similarity between features x_i and y_j . At training stage, we need the target components mask of reference images I_m^r for calculating the contextual loss. The contextual loss can be expressed as follows:

$$\mathcal{L}_{cx} = -\log(CX(\phi_l(\mathbf{I}_g * \mathbf{I}_m^s), \phi_l(\mathbf{I}_r * \mathbf{I}_m^r))), \quad (7)$$

The per-pixel loss can be expressed as follows:

$$\mathcal{L}_{pixel} = \|\phi_l(\mathbf{I}_g * \mathbf{I}_m^s) - \phi_l(\mathbf{I}_r * \mathbf{I}_m^r)\|_1, \quad (8)$$

Lastly, we use an adversarial loss for minimizing the distance between the distribution of the real image and the generated image. Here, LSGAN [Mao *et al.*, 2017] is adopted for stable training. The adversarial loss can be expressed as follows:

$$\mathcal{L}_{ad_G} = \mathbb{E}[(\mathcal{D}(\mathbf{I}_g) - 1)^2], \quad (9)$$

$$\mathcal{L}_{ad_D} = \mathbb{E}[\mathcal{D}(\mathbf{I}_g)^2] + \mathbb{E}[(\mathcal{D}(\mathbf{I}_s) - 1)^2], \quad (10)$$

With total variational regularization loss \mathcal{L}_{tv} [Johnson *et al.*, 2016] added to encourage the spatial smoothness in the generated images, we obtain our full objective as:

$$\begin{aligned} \mathcal{L} = & \lambda_1 \mathcal{L}_{perc} + \lambda_2 \mathcal{L}_{style} + \lambda_3 \mathcal{L}_{cx} \\ & + \lambda_4 \mathcal{L}_{pixel} + \lambda_5 \mathcal{L}_{tv} + \lambda_6 \mathcal{L}_{ad_G} \end{aligned} \quad (11)$$

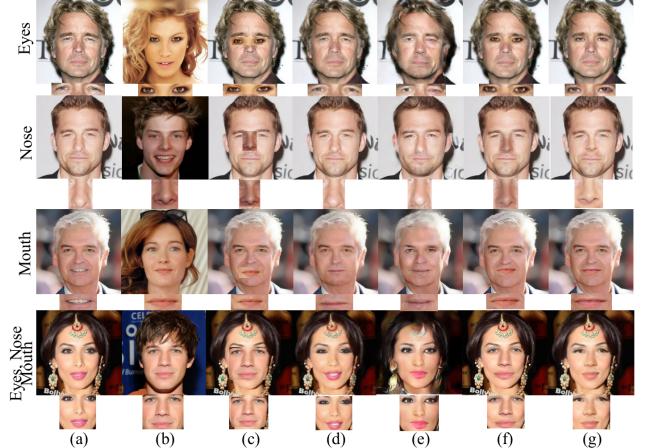


Figure 4: Comparisons among (c) copy-and-paste, (d) AttGAN, (e) ELEGANT, (f) the commercial state of the art algorithm in Adobe Photoshop and (g) the proposed rFACE technique. The source and reference images are shown in (a) and (b), respectively. AttGAN and ELEGANT edit attributes: 'Narrow_Eyes', 'Pointy_Nose', 'Mouth_Slightly_Open' and all of them respectively.

4 Experiments

4.1 Dataset and Preprocessing

The face attribute dataset CelebAMask-HQ [Lee *et al.*, 2019] contains 30000 aligned facial images with the size of 1024×1024 and corresponding 30000 semantic segmentation labels with the size of 512×512 . Each label in the dataset has 19 classes (e.g., "left eye", "nose", "mouth"). In our experiments, three face components are considered, i.e., eyes ("left eye & right eye"), nose ("nose"), and mouth ("mouth & u-lip & l-lip"). We obtain rough version of face components from semantic segmentation labels by an image dilation operation, which are defined as mask images. We take 2,000 images as the test set for performance evaluation, using rest images to train our model. All images are resized to 256×256 .

4.2 Implementation Details

Our end-to-end network is trained on four GeForce GTX TITAN X GPUs of 12GB memory. Adam optimizer is used in experiments with $\beta_1 = 0.0$ and $\beta_2 = 0.9$. The hyperparameters from λ_1 to λ_6 are assigned as 0.1, 250, 1, 0.5, 0.1, 0.01 respectively. For each source image, we remove two or three target face components for training. During testing, our model can change one or more face components.

4.3 Qualitative Evaluation

Comparison with Other Methods. In order to demonstrate the superiority of the proposed method, we compare the quality of sample synthesis results to several benchmark methods. In addition to face editing model AttGAN [He *et al.*, 2019] and ELEGANT [Xiao *et al.*, 2018], we also consider copy-and-paste as a naive baseline and Adobe Photoshop image editing as an interactive way to produce synthesized face images. According to the results presented in Figure 4, although margins of edited facial components in Adobe

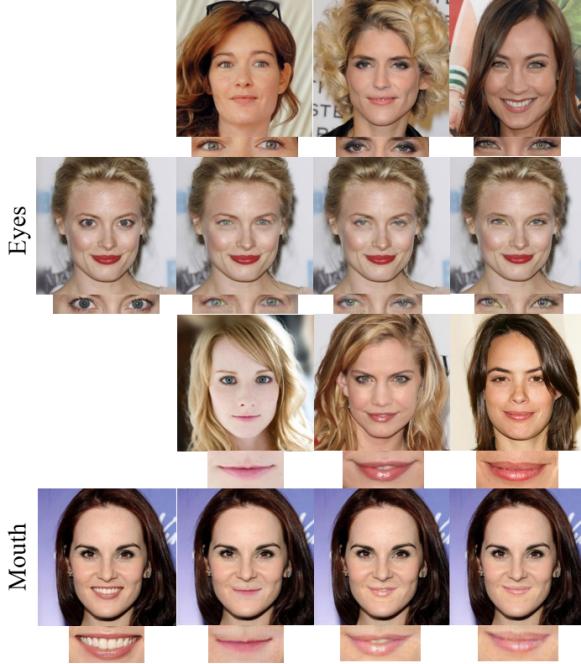


Figure 5: Illustrations of multi-modal face components editing, including ‘eyes’ and ‘mouth’. The first column represents source images, the first and third rows are reference images, and the second and fourth rows are synthesized images according to reference images above.

Photoshop are much smoother than those in results of copy-and-paste, obvious ghosting artifacts and color distortions still exist and more fine-grained manual labour is required to improve the quality. In contrast, AttGAN could generate realistic synthetic images which are indistinguishable to generic ones in an end-to-end manner. However, pre-defined facial attribute labels are used to guide AttGAN to transform facial components thus the diversity of generated images is limited. Moreover, as can be seen in Figure 4, AttGAN only produces subtle changes that could hardly reflect the desired translation. ELEGANT, as a reference-guided face attribute editing method, can learn obvious semantic attributes (e.g., open eyes or close mouth), but could not learn abstract shape information (e.g., nose editing). Moreover, ELEGANT produces large deformation and unexpected artifacts at other attribute-unrelated areas, especially the editing of multiple components. On the contrary, our method takes arbitrary face images as reference, which significantly increases the diversity of generated images.

Multi-Modal Face Components Editing. Reference-guided face components editing improves the diversity and controllability of generated faces, as the stylistic information is designated by arbitrary reference images. As shown in Figure 5, target face components of interest (e.g. eyes and mouths) are transformed to be of the same style as the corresponding reference image. For example, synthesized mouths in faces of the last row accurately simulate the counterpart in reference images, in terms of both overall shape (e.g. raised corners of pursed mouth) and local details (e.g. partially covered teeth). Meanwhile, they are naturally blended in

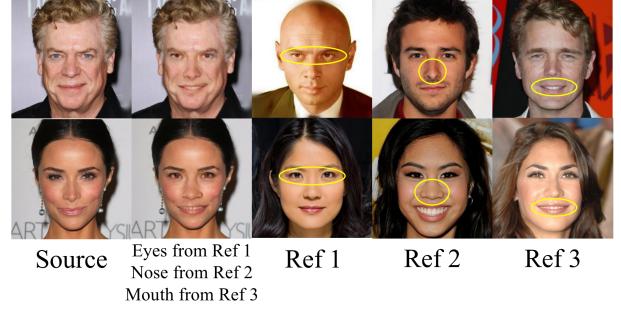


Figure 6: Illustrations of hybrid face components editing. Given the source image (2nd column), the synthesized image (3rd column) are obtained by transfer the eyes from reference 1 (R1), nose from reference 2 (R2) and mouth from reference 3 (R3).

to the source face without observable color distortions and ghosting artifacts, demonstrating the effectiveness of proposed method.

Hybrid Face Components Editing. With the aid of mask images, the proposed method could integrate multiple references images with different facial components into the source face, as shown in Figure 6. It could be observed that only the target facial component of interest is transformed to present the same style as the corresponding reference and blended into the background seamlessly, while leaving the rest part of the image unchanged. This suggests that our framework could synthesize face images with high diversity (arbitrary reference faces) and strong controllability (ROI present by masks), while maintaining the visual fidelity of the source image.

4.4 Quantitative Evaluation

Following most of face portrait methods [He *et al.*, 2019; Wu *et al.*, 2019], we leverage Fréchet Inception Distance (FID, lower value indicates better quality) [Heusel *et al.*, 2017] and Multi-Scale Structural SIMilarity (MS-SSIM, higher value indicates better quality) [Wang *et al.*, 2003] to evaluate the performance of our model. FID is used to measure the quality and diversity of generated images. MS-SSIM is used to evaluate the similarity of two images from luminance, contrast, and structure.

We compare our method with AttGAN[He *et al.*, 2019], one of label-conditioned methods, and ELEGANT [Xiao *et al.*, 2018], one of reference-guided methods. For AttGAN and ELEGANT, the value of FID or MS-SSIM are the result of averaging three pre-defined attributes for shape changing, including ‘Narrow_Eyes’, ‘Pointy_Nose’ and ‘Mouth_Slightly_Open’. The backbone of r-FACE is similar to image inpainting task, so we also compare our method with GLCIC[Iizuka *et al.*, 2017], one of popular face inpainting methods. As is shown in TABLE 1, comparing with these methods, the FID of our method is much lower. With the observation that the MS-SSIM of our method is lower than AttGAN and GLCIC, we analyze the reasons: MS-SSIM is sensitive to luminance, contrast, and structure, however, 1) GLCIC does not have any constraints on the structure or shape of components, just requiring the missing regions to be completed; 2) AttGAN edits shape-changing attributes with

Method	FID ↓	MS-SSIM ↑
GLCIC[Iizuka <i>et al.</i> , 2017]	8.09	0.95
AttGAN[He <i>et al.</i> , 2019]	6.28	0.96
ELEGANT [Xiao <i>et al.</i> , 2018]	15.97	0.82
r-FACE (Ours)	5.81	0.92
w / o attention	6.25	0.90
w / o contextual loss	5.27	0.90
w / o style loss	8.28	0.90
w / o perceptual loss	8.49	0.89

Table 1: Comparisons of FID and MS-SSIM on the CelebA-HQ dataset.

subtle changes that are hard to be observed as shown in Figure 4, so the change of luminance, contrast, and structure is limited. In contrast, r-FACE imposes a geometric similarity constraint on the components of source images and reference images, which changes the shape or structure drastically and even affects the identity of the face.

4.5 Ablation Study

Example-guided Attention Module. To investigate the effectiveness of the example-guided attention module, we conduct a variant of our model, denoted as “r-FACE w / o attention”. In “r-FACE w / o attention”, we train r-FACE without example-guided attention module. To learn face components from a reference image, we directly concatenate the features of reference images with that of source images, which introduced the features of the whole reference image. TABLE 1 shows the comparison on FID and MS-SSIM between “r-FACE w / o attention” and “r-FACE”. Our method outperforms “r-FACE w / o attention” on both FID and MS-SSIM, which indicates that example-guided attention module can explicitly transfer the corresponding face components of a reference image and reduce the impact of other information of the reference image.

Loss Configurations. In this section, we conduct ablation studies on different loss, such as the contextual loss, the style loss and the perceptual loss, for evaluating their individual contributions on our framework. As shown in TABLE 1, the quantitative comparison among “r-FACE w / o contextual loss”, “r-FACE w / o style loss”, “r-FACE w / o perceptual loss” and “r-FACE” are presented. “r-FACE” outperforms other loss configurations on MS-SSIM, which indicates the effectiveness of the three losses. It is observed that the FID of “r-FACE” is also better than other loss configurations except “r-FACE w/o contextual loss”. We argue that contextual loss plays an important role in restricting the shape of face components. The framework has no constraints on the shape after removing the contextual loss, so “r-FACE w / o contextual loss” is easier to generate missing components, merely requiring generated images to look real. Figure 7 compares the visual effects of different loss configurations. We find “r-FACE w / o contextual loss” could not synthesize components with the corresponding shape of reference images, which indicates the contextual loss is significantly important in shape constraints. In the results of “r-FACE w / o style loss” and “r-FACE w / o perceptual loss”, color distortions and obvious

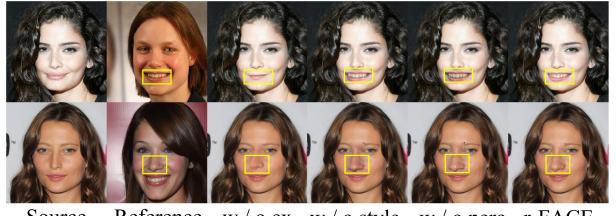


Figure 7: Visual comparisons for different loss configurations, including “r-FACE w / o contextual loss”, “r-FACE w / o style loss”, “r-FACE w / o perceptual loss” and “r-FACE”.

ghosting artifacts are observed, which show the style loss and the perceptual loss are able to preserve skin color of source image. Above all, we prove that three losses contribute to the performance of our framework and the combination of all losses achieves the best results.

4.6 Discussion and Limitation

Reference guided face component editing has wide real-life applications in interactive entertainment, security and data augmentation. Given whole reference images of any identity, r-FACE performs various face component editing, which can achieve the effect of “plastic surgery”. Besides, when all face components of a reference face are transformed to the source face, r-FACE is further extended to face swapping task. Moreover, As the face component is a part of face, editing face components may change the identity of the person. Therefore, r-FACE can be used for data augmentation to generate face images of different identities. Although r-FACE obtains diverse and controllable face component editing results, reference images with significant differences in pose are still challenging for our model. We will continue to explore solving extreme pose problems in further work.

5 Conclusion

In this work, we have proposed a novel framework, reference guided face components editing (r-FACE), for high-level face components editing, such as eyes, nose and mouth. r-FACE can achieve diverse, high-quality, and controllable component changes in shape from given references, which breaks the shape and precise intermediate presentation limitation of existing methods. For embedding the target face components of reference images to source images specifically, an example-guided attention module is designed to fuse the features of source images and reference images, further boosting the performance of face component editing. The extensive experiments demonstrate that our framework can achieve state-of-art face editing results with observable geometric changes.

Acknowledgments

This work was partially supported by the Natural Science Foundation of China (Grant No. U1836217, Grant No. 61702513, Grant No. 61721004, and Grant No. 61427811). This research was also supported by Meituan-Dianping Group, CAS-AIR and Shandong Provincial Key Research and Development Program (Major Scientific and Technological Innovation Project) (NO.2019JZZY010119).

References

- [Bertalmio *et al.*, 2000] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *SIGGRAPH*, pages 417–424, 2000.
- [Bertalmio *et al.*, 2003] Marcelo Bertalmio, Luminita Vese, Guillermo Sapiro, and Stanley Osher. Simultaneous structure and texture image inpainting. *IEEE TIP*, 12(8):882–889, 2003.
- [Cao *et al.*, 2019] Jie Cao, Yibo Hu, Bing Yu, Ran He, and Zhenan Sun. 3D aided duet GANs for multi-view face image synthesis. *IEEE TIFS*, 14(8):2028–2042, 2019.
- [Choi *et al.*, 2018] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, pages 8789–8797, 2018.
- [Criminisi *et al.*, 2003] Antonio Criminisi, Patrick Perez, and Kentaro Toyama. Object removal by exemplar-based inpainting. In *CVPR*, volume 2, pages II–II, 2003.
- [Dolhansky and Canton Ferrer, 2018] Brian Dolhansky and Cristian Canton Ferrer. Eye in-painting with exemplar generative adversarial networks. In *CVPR*, pages 7902–7911, 2018.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014.
- [Gu *et al.*, 2019] Shuyang Gu, Jianmin Bao, Hao Yang, Dong Chen, Fang Wen, and Lu Yuan. Mask-guided portrait editing with conditional gans. In *CVPR*, pages 3436–3445, 2019.
- [He *et al.*, 2019] Zhenliang He, Wangmeng Zuo, Meina Kan, Shiguang Shan, and Xilin Chen. AttGAN: Facial attribute editing by only changing what you want. *IEEE TIP*, 2019.
- [Heusel *et al.*, 2017] 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 *NeurIPS*, pages 6626–6637, 2017.
- [Iizuka *et al.*, 2017] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM ToG*, 36(4):107, 2017.
- [Jo and Park, 2019] Youngjoo Jo and Jongyoul Park. SC-FEGAN: Face editing generative adversarial network with user’s sketch and color. In *ICCV*, October 2019.
- [Johnson *et al.*, 2016] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711, 2016.
- [Langner *et al.*, 2010] Oliver Langner, Ron Dotsch, Gijsbert Bijlstra, Daniel HJ Wigboldus, Skyler T Hawk, and AD Van Knippenberg. Presentation and validation of the Radboud Faces Database. *Cognition and Emotion*, 24(8):1377–1388, 2010.
- [Lee *et al.*, 2019] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. MaskGAN: Towards diverse and interactive facial image manipulation. *arXiv preprint arXiv:1907.11922*, 2019.
- [Liu *et al.*, 2015] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738, 2015.
- [Liu *et al.*, 2019] Yunfan Liu, Qi Li, and Zhenan Sun. Attribute-aware face aging with wavelet-based generative adversarial networks. In *CVPR*, pages 11877–11886, 2019.
- [Mao *et al.*, 2017] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, pages 2794–2802, 2017.
- [Mechrez *et al.*, 2018] Roey Mechrez, Itamar Talmi, and Lihai Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *ECCV*, pages 768–783, 2018.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Song *et al.*, 2019] Linsen Song, Jie Cao, Linxiao Song, Yibo Hu, and Ran He. Geometry-aware face completion and editing. In *AAAI*, pages 2506–2513, 2019.
- [Wang *et al.*, 2003] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *ACSSC*, volume 2, pages 1398–1402, 2003.
- [Wu *et al.*, 2019] Po-Wei Wu, Yu-Jing Lin, Che-Han Chang, Edward Y Chang, and Shih-Wei Liao. RelGAN: Multi-domain image-to-image translation via relative attributes. In *ICCV*, pages 5914–5922, 2019.
- [Xiao *et al.*, 2018] Taihong Xiao, Jiapeng Hong, and Jinwen Ma. Elegant: Exchanging latent encodings with gan for transferring multiple face attributes. In *ECCV*, pages 168–184, 2018.
- [Yang *et al.*, 2018] Hongyu Yang, Di Huang, Yunhong Wang, and Anil K Jain. Learning face age progression: A pyramid architecture of gans. In *CVPR*, pages 31–39, 2018.
- [Zheng *et al.*, 2019] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. Pluralistic image completion. In *CVPR*, pages 1438–1447, 2019.