

# Text-Guided Face Recognition using Multi-Granularity Cross-Modal Contrastive Learning

Md Mahedi Hasan, Shoaib Meraj Sami, and Nasser Nasrabadi  
West Virginia University, Morgantown, West Virginia, USA

mh00062@mix.wvu.edu, sms00052@mix.wvu.edu, nasser.nasrabadi@mail.wvu.edu

## Abstract

*State-of-the-art face recognition (FR) models often experience a significant performance drop when dealing with facial images in surveillance scenarios where images are in low quality and often corrupted with noise. Leveraging facial characteristics, such as freckles, scars, gender, and ethnicity, becomes highly beneficial in improving FR performance in such scenarios. In this paper, we introduce text-guided face recognition (TGFR) to analyze the impact of integrating facial attributes in the form of natural language descriptions. We hypothesize that adding semantic information into the loop can significantly improve the image understanding capability of an FR algorithm compared to other soft biometrics. However, learning a discriminative joint embedding within the multimodal space poses a considerable challenge due to the semantic gap in the unaligned image-text representations, along with the complexities arising from ambiguous and incoherent textual descriptions of the face. To address these challenges, we introduce a face-caption alignment module (FCAM), which incorporates cross-modal contrastive losses across multiple granularities to maximize the mutual information between local and global features of the face-caption pair. Within FCAM, we refine both facial and textual features for learning aligned and discriminative features. We also design a face-caption fusion module (FCFM) that applies fine-grained interactions and coarse-grained associations among cross-modal features. Through extensive experiments conducted on three face-caption datasets, proposed TGFR demonstrates remarkable improvements, particularly on low-quality images, over existing FR models and outperforms other related methods and benchmarks.*

## 1. Introduction

Current one-to-one face recognition (FR) algorithms such as ArcFace [5] and AdaFace [23] face challenges in achieving high verification rates (VR), particularly in

surveillance-based applications [13, 48], where a low-resolution face probe is captured under non-ideal conditions. Covariate factors like non-uniform lighting, occlusion, and non-frontal pose in such scenarios, severely degrade the image quality, leading to decreased matching performance [28, 48]. One way to improve the performance of an FR algorithm is to integrate auxiliary information, such as facial marks, gender, age, skin and hair color, facial expression, and other distinctive facial attributes to a face recognition model [11, 20, 50].

Natural language descriptions, which highlight the distinct characteristics of a face image, can also be employed as a soft biometric to improve the performance of a FR model [14]. In this research, we hypothesize that semantic information can substantially improve the image understanding capability of an FR algorithm compared to other soft biometrics. Therefore, the core idea of this research is to develop a model that effectively integrates textual descriptions with facial images, thereby maximizing performance. As both image and textual modalities are complementary to each other, the combination of textual and facial features will result in a significant performance leap. This text-guided face recognition (TGFR) has great potential for applications in systems where face images are captured in real-world settings, such as criminal investigations and intelligent video surveillance [14]. In criminal investigations, the testimony of the witnesses can be effectively utilized within the TGFR model to find or narrow down the potential suspects even from low-quality images.

However, integrating textual descriptions or captions within a FR loop poses several challenges. Firstly, natural language captions are inherently more abstract than facial images. A limited set of facial attributes may be insufficient to convey all the fine-grained details of a complex face image. Secondly, as describing a face image involves subjective judgment, captions written by different annotators for a particular face can be completely or partially different. Sometimes, an attribute, which one annotator considers important may be ignored by others. Thirdly, different captions of a particular face image may contain large word

variability. Fourthly, annotators might disagree on specific attributes, leading to incorrect representations of a person’s face. For example, annotators may differ on whether the person has a sharp nose. Lastly, annotators occasionally include a surplus of non-discriminative information, readily present on the face, thus making it challenging to differentiate the face-caption pair from others. Therefore, learning accurate embedding from such textual descriptions is the most challenging step of a text-guided FR model.

As the primary objective of this work is to improve the performance of SOTA FR algorithms [5, 23, 32], in our TGFR framework, we utilize them as fixed feature extractor. By keeping them in a frozen state, we can accurately measure the performance gain of the FR model resulting from textual supervision. However, the contextual embeddings generated from the foundation model are unaligned with image features and have limited distinctiveness for face recognition. To address these challenges, we propose a face caption alignment module (FCAM) that finetunes the text encoder with the specific aim of learning text embeddings that are both discriminative and well-aligned with visual features. The proposed FCAM incorporates a global caption-image contrastive loss (CICL) and a local word-region contrastive loss (WRCL) to learn alignments between the caption-image and word-regions, respectively. Additionally, it includes an intra-modal contrastive objective as well as an identity loss to produce discriminative features for both visual and semantic modalities.

Furthermore, in a conventional dual-encoder architecture, a fusion scheme is typically employed to learn a joint representation from image and text features. However, a simple feature-level fusion (FLF) scheme may be ineffective as the regional image features and word embeddings lack prior relationships [27]. Moreover, the FLF scheme treats textual and image features as separate entities, thus failing to capture their interactions and dependencies. Therefore, we introduce a novel and efficient attention-guided fusion scheme called face-caption fusion module (FCFM), aiming to explore cross-modal fine-grained interactions and coarse-grained associations for improved performance in multimodal fusion. In summary, our contributions are three-fold:

- We boost the performance of SOTA FR algorithms, including AdaFace [23] and ArcFace [5], by learning effective representations through cross-modal and intra-modal contrastive supervision.
- We introduce FCAM, a novel module to learn aligned global and local features within a shared semantic space by effectively utilizing the CICL to align caption-image embeddings and WRCL to align attention-weighted region with word embeddings.

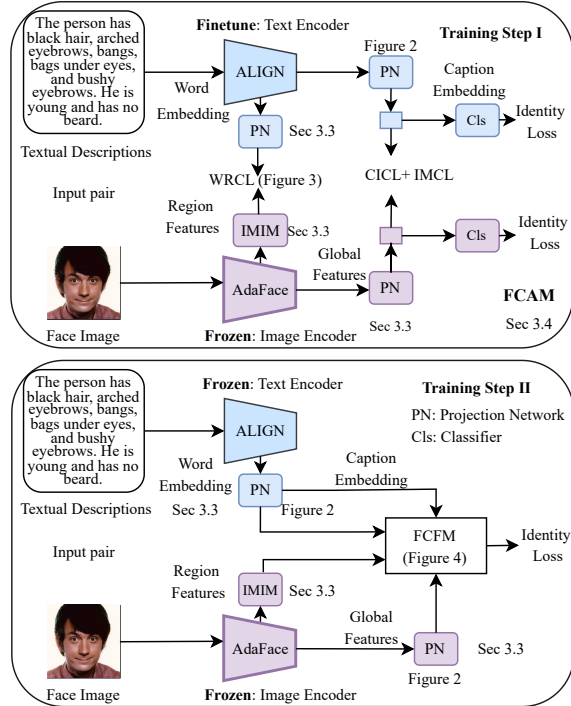


Figure 1. An overview of our proposed TGFR framework, comprising an image encoder, a text encoder, a face-caption alignment module (FCAM), and a face-caption fusion module (FCFM). We follow a two-step learning process. In the initial step, we train both the foundation model and the FCAM module while keeping the parameters of the FR model frozen. The goal of FCAM is to map the local and global embeddings of face-caption pairs into a shared space, ensuring a proper alignment between image-text modalities. In the final step, we train the FCFM module.

- Experiments conducted on three challenging face-caption datasets demonstrate significant improvements over SOTA methods and baseline approaches in both 1:1 verification and Rank-1 identification.

## 2. Related work

**Text-guided face recognition** Facial semantic attributes have been extensively exploited as auxiliary information for various tasks, such as face image retrieval [49], generation [38], and editing [17, 44]. Gonzalez *et al.* [11] presented an overview of soft biometrics for face recognition in unconstrained scenarios. The authors reported a maximum relative performance improvement of 40% over the FR models when utilizing 6 discriminative facial attributes.

Our work is closely related to a recent approach to face recognition using captions, namely CGFR, introduced by Hasan and Nasser [14]. In CGFR, the authors employed a pre-trained FR model and a pre-trained BERT model [8] for extracting features from the respective modalities. They further finetuned the BERT model using a refinement module.

Afterward, a cross-modal fusion scheme was employed to transform the uni-modal representations into a unified embedding space. A notable drawback of CGFR [14] is that its feature extractors were trained independently on objectives that were distinct and unrelated to the FR task. For example, language models such as BERT [8] and RoBERTa [30] were trained on vast corpora using masked language modeling. Therefore, these encoders generate irrelevant features that result in lower performance in downstream FR tasks.

In this research, rather than adopting BERT, we employ a CLIP-like foundation model [21,26,34] that has been pre-trained on a large-scale dataset containing image-text pairs for a wide range of multimodal tasks including image-text matching [4,10,21,27,43], cross-modal retrieval [27,31]. The rationale behind selecting a CLIP-like model, such as ALIGN [21], as our text encoder lies in its primary objective of aligning images and text. Furthermore, they trained CGFR using the SOTA CMPC [52] and DAMSM [46] losses on the low-quality MMCelebA dataset [44]. In contrast, we train our TGFR using a variety of cross-modal contrastive losses and evaluate our model on three challenging face-caption datasets. Additionally, we incorporate an intra-modal contrastive supervision and an identity loss to further enhance the discriminability of the learned features.

**Representation learning with text supervision** In recent years, with the availability of datasets containing large corpora of image-text pairs, representation learning with text supervision [4,7,21,27,31] has exhibited remarkable success across a multitude of downstream tasks such as visual question answering task [2,37], cross-modal retrieval [27,31], and zero-shot image classification [21]. The semantic context derived from textual supervision aids the model in capturing more comprehensive multimodal representations, resulting in improved performance. Desai and Johnson introduced the VirTex framework [7], where they trained a CNN model with textual descriptions for image captioning tasks. By fine-tuning the pre-trained model, they achieved improved performance in image classification and instance segmentation tasks. Additionally, in the iCAR framework [42], authors proposed a novel deep fusion approach for effective image-text alignment. This method demonstrated superior performance on several benchmark datasets for zero- and few-shot image classification.

**Contrastive learning** To tackle the challenges of computational complexities, noise-contrastive estimation (NCE) was introduced [12]. In this method, rather than computing the probability of the target word among all possible words in the vocabulary, a binary cross-entropy loss is employed to differentiate between the actual target word and noise samples. Another approach, information NCE [33], focuses on maximizing a lower bound on the mutual infor-

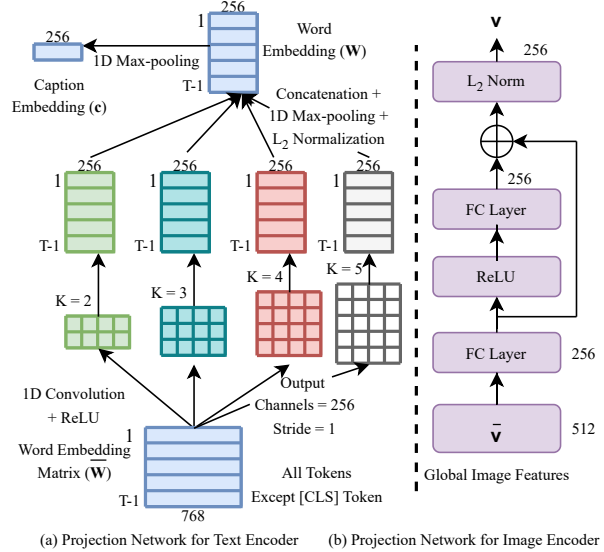


Figure 2. (a) The architecture of the proposed CNN-based projection for generating phrase-level contextual embeddings from the output of the foundation model. (b) The network for projecting global image features into a shared space. ( $K = 1D$  kernel)

mation existing between images and captions, particularly in the context of cross-modal retrieval tasks.

Another powerful technique for learning image-text representations is contrastive pretraining, which has recently demonstrated remarkable success in various multimodal tasks [21,27,34]. CLIP [34] introduces an approach for learning visual representations using caption by establishing semantic similarity between image-text pairs through contrastive loss. Jia *et al.* [21] designed a method called noisy student training, which utilized noisy text data to improve image-text representation learning by aligning image-text pairs using a contrastive loss. In another approach, authors [27] employed the momentum distillation technique for image-text representation learning. They proposed an image-text contrastive loss to align image-text features before fusion. Although these methods achieve remarkable performance on tasks related to multimodal representation learning, they often fail to capture complex interactions between image regions and words, resulting in lower performance on fine-grained image-text classification [24].

Moreover, in addition to image-text matching loss [21,27,31,34], numerous other contrastive learning-based cross-modality loss functions have been proposed in the literature. These encompass, triplet loss [29,47], word-region alignment [4,46], word-patch alignment [24], token-wise alignment [41] and more, all aimed at tackling a variety of multimodal tasks. For example, Zhang *et al.* [51] introduced ConVIRT to learn visual features from chest X-ray images through the analysis of paired reports. They employed a bidirectional contrastive objective between visual

and textual modalities, similar to our CICL, and achieved superior performance across four medical image classification and two zero-shot retrieval tasks. Moreover, Huang *et al.* [18] contrasted image sub-regions with the corresponding words in the textual report to learn both global and local representations in their proposed GLoRIA framework.

### 3. Our approach

#### 3.1. Model architecture

An overview of our method is illustrated in Figure 1. It consists of several components: an image encoder for extracting both local and global image features from an input face, a text encoder for learning embeddings from the associated captions, a face-caption alignment module, and a face-caption fusion module. The uni-modal features are then fed into the FCAM module, which serves three main purposes: 1) aligning cross-modal features, 2) applying intra- and inter-modal feature interactions, and 3) learning discriminative image-text features. Finally, the FCFM module is applied to implement fine-grained interactions between cross-modal features.

#### 3.2. Multimodal feature extraction

In this study, we employ a SOTA FR model [5,23] as our image encoder,  $f_v$ , and a foundation model [21,26,34,36], pre-trained on a large number of caption-image pairs, as our text encoder,  $f_c$ . The iRestNet18 [5,16] is adopted as the backbone network for the image encoder. In this network, the output of the final adaptive average pooling layer produces the global feature vector,  $\bar{\mathbf{v}} \in \mathbb{R}^{512}$ , for the input image which contains high-level semantic information of the face. Additionally, regional image features,  $\bar{\mathbf{R}} \in \mathbb{R}^{256 \times 14 \times 14}$ , are extracted from an intermediate convolutional stage. From the output of the foundation model, we obtain a contextualized word matrix,  $\bar{\mathbf{W}} \in \mathbb{R}^{D \times T}$ . Here,  $T$  represents the total token count, while  $D$  denotes the dimensionality of each token embedding.

#### 3.3. Projection networks (PNs)

The output of the text encoder comprises 768-dimensional contextualized embeddings for each token in the caption [21,26]. In this work, we design a 1D CNN-based projection network to map each embedding into a shared 256-dimensional multimodal space, Figure 2 (a), illustrates the proposed network which allows us to capture phrase-level contextual information. For convolutions, each utilizing a kernel size ranging from 2 to 5 and a stride of 1 are applied to the word matrix,  $\bar{\mathbf{W}} \in \mathbb{R}^{768 \times (T-1)}$ , to effectively capture local patterns within the text. Next, we concatenate all generated feature maps. Subsequently, we apply a 1D max-pooling and  $L_2$  normalization to the feature maps to generate word embeddings,  $\mathbf{W} \in \mathbb{R}^{256 \times (T-1)}$ .

Moreover, we incorporate an additional 1D max-pooling to compute the final caption embedding,  $\mathbf{c} \in \mathbb{R}^{256}$ . In addition, another projection network, as illustrated in Figure 2 (b), is designed to map the global image features,  $\bar{\mathbf{v}} \in \mathbb{R}^{512}$ , into a shared embedding space,  $\mathbf{v} \in \mathbb{R}^{256}$ .

To increase the inherent relationships embedded within the local image features,  $\bar{\mathbf{R}} \in \mathbb{R}^{256 \times 14 \times 14}$ , we introduce an intra-modal interaction module (IMIM). The proposed IMIM, as depicted in Figure 3, refines local features to generate more informative representations by capturing fine-grained local patterns. It also contributes in filtering out irrelevant information, thereby improving the model’s robustness to variations. In this module, we first normalize [19] the input features. Afterward, self-attention (SA) [40] mechanism is applied, where the keys, queries, and values are learned via 1x1 convolutions. The SA layer improves intra-modality relationships effectively, by capturing long-range dependencies of  $\bar{\mathbf{R}}$ . It is noteworthy that in our experiment we also applied the proposed IMIM module to the textual branch. However, due to the strong intra-modal interactions, that are already present within the contextual embeddings of transformer-based foundation models [8,21,26,30], the performance of our model remains the same. Hence, we skip it.

#### 3.4. Face caption alignment module (FCAM)

We design this module to learn well-aligned global and local features for both the visual and semantic modalities. The training objectives of this module are as follows: 1) CICL, designed to find relationships between the face image and its associated caption, 2) WRCL, designed to learn fine-grained alignments between image sub-regions and word embeddings. 3) IMCL, for improving feature representation within each modality and 4) an identity loss.

**Caption-image contrastive loss (CICL)** To learn aligned global image and caption embedding, similar to the approach in [18,41,51], we design a caption-image contrastive objective between the visual and textual domains. Our CICL aims to maximize the similarity score between true face-caption pairs against random pairs. We pass the input face image,  $\mathbf{x}_f$ , through an image encoder and then the proposed projection network to generate the global image features,  $\mathbf{v}$ . In parallel, we compute the contextualized caption embedding,  $\mathbf{c}$ , from the output of the text encoder and the 1D-CNN-based projection network. We encode a mini-batch of  $B$  input caption-image pairs ( $\mathbf{x}_f, \mathbf{x}_c$ ) to generate global caption-image embeddings ( $\mathbf{v}, \mathbf{c}$ ). Next, we compute the face-to-caption contrastive loss,  $L_i^{(f2c)}$ , for the  $i^{th}$  pair in the similar way to the InfoNCE loss [33]:

$$L_i^{(f2c)} = -\log \frac{\exp(\langle \mathbf{v}_i, \mathbf{c}_i \rangle / \tau)}{\sum_{k=1}^B \exp(\langle \mathbf{v}_i, \mathbf{c}_k \rangle / \tau)}. \quad (1)$$



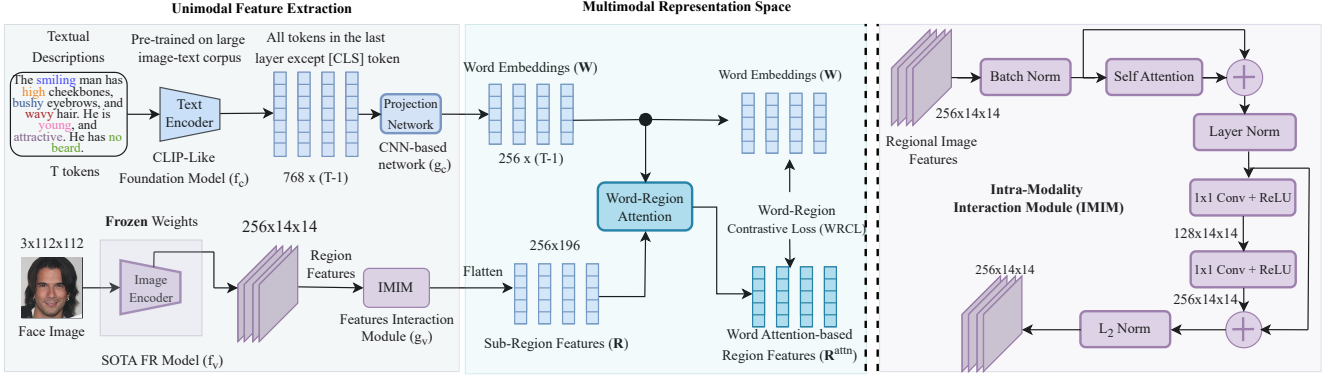


Figure 3. The block diagram of the proposed face-caption alignment module where we learn context-aware sub-regional features by emphasizing the distinctive attributes of a face image, corresponding to a specific word. Initially, the extracted local image features and word embeddings are fed into the projection networks to map them in a shared multimodal semantic space. Next, a word-region contrastive loss is applied to the context-aware regional features and their respective word representations.

where  $\langle \mathbf{v}_i, \mathbf{c}_i \rangle = \mathbf{v}_i^\top \mathbf{c}_i / \|\mathbf{v}_i\| \|\mathbf{c}_i\|$  represents the cosine similarity. The loss drives the TGFR model to align positive face-caption pairs while effectively preserve the mutual information. In addition, we define a similar caption-to-face contrastive loss,  $L^{(c2f)}$ , as:

$$L_i^{(c2f)} = -\log \frac{\exp(\langle \mathbf{c}_i, \mathbf{v}_i \rangle / \tau)}{\sum_{k=1}^B \exp(\langle \mathbf{c}_i, \mathbf{v}_k \rangle / \tau)}. \quad (2)$$

Here,  $\tau$  represents a learnable parameter. We calculate the final CICL as the summation of these two losses.

**Word-region contrastive loss (WRCL)** To align sub-regional image features,  $\mathbf{R} \in \mathbb{R}^{256 \times 196}$ , with word embeddings,  $\mathbf{W} \in \mathbb{R}^{256 \times (T-1)}$ , we begin by mapping the  $\mathbf{R}$  and  $\mathbf{W}$  into a multimodal semantic space. Then, we calculate a dot-product similarity matrix denoted as  $\mathbf{S} \in \mathbb{R}^{(T-1) \times 196}$ , where  $(T-1)$  represents the number of all tokens except [CLS] token, and 196 corresponds to the image sub-regions. Then, the similarities of each sub-region are normalized.

$$\bar{s}_{i,j} = \frac{\exp(s_{i,j})}{\sum_{i=0}^{T-2} \exp(s_{k,j})}, \text{ where } \mathbf{S} = \mathbf{W}^T \mathbf{R}. \quad (3)$$

Next, we learn an attention weight vector,  $\alpha$ , which assigns varying weights to different image sub-regions according to their association with a specific word in the caption. To calculate this, we compute the attention weight for the  $i^{\text{th}}$  word by normalizing  $\bar{s}_{i,j}$  across all image sub-regions. Finally, we compute attention-weighted sub-region features,  $\mathbf{r}^{\text{attn}}$ , using the attention weighted vector,  $\alpha$ .

$$\mathbf{r}_i^{\text{attn}} = \sum_{j=0}^{195} \alpha_j \times \mathbf{r}_j \text{ where } \alpha_j = \frac{\exp(\bar{s}_{i,j} / \tau_1)}{\sum_{k=0}^{195} \exp(\bar{s}_{i,k} / \tau_1)}. \quad (4)$$

Here,  $\tau_1 \in \mathbb{R}$  is the temperature parameter. Now, to find the attention-guided matching score between all words

and attention-weighted sub-regions features, we define the following matching function,  $f$ .

$$f(\mathbf{x}_f, \mathbf{x}_c) = \log \left( \sum_{i=1}^{T-1} \exp(\langle \mathbf{r}_i^{\text{attn}}, \mathbf{w}_i \rangle / \tau_2)^{\tau_2} \right). \quad (5)$$

Here,  $\tau_2$  is another hyperparameter and  $(\mathbf{x}_f, \mathbf{x}_c)$  is the input face-caption pair. Similar to approaches in [9, 46], we define the word-region contrastive loss as the negative log posterior probability that aims to maximize the posterior probability of the attention-weighted region features. The contrastive loss for image sub-regions being matched with their corresponding words are given below:

$$L_{WRCL}^{\tau|w} = \sum_{i=1}^B -\log \frac{\exp(f(\mathbf{x}_{f_i}, \mathbf{x}_{c_i}) / \tau_3)}{\sum_{k=1}^B \exp(f(\mathbf{x}_{f_i}, \mathbf{x}_{c_k}) / \tau_3)}. \quad (6)$$

In addition, we also minimize  $L_{WRCL}^{w|r}$ . The final loss,  $L_{WRCL}$ , is the summation of these two loss.

**Intra-modal contrastive loss (IMCL)** While CICL and WRCL effectively align caption-image and word-region embeddings, respectively, they failed to incorporate self-supervision within each modality. Therefore, we introduce an additional contrastive objective, called intra-modal contrastive learning (IMCL), that introduces intra-modal self-supervision, leading to the semantic alignment between true pairs within each modality. The application of the IMCL, in our work, significantly improves image-text feature representations and also makes them more discriminative. For input captions, we pair two different captions belonging to a particular subject as true pairs  $(\mathbf{x}_{c1}, \mathbf{x}_{c2})$ . Likewise, for the visual input, we consider two different face images belonging to a subject, to form a true pair  $(\mathbf{x}_{f1}, \mathbf{x}_{f2})$ . However, in cases where a subject has only one view in the dataset,

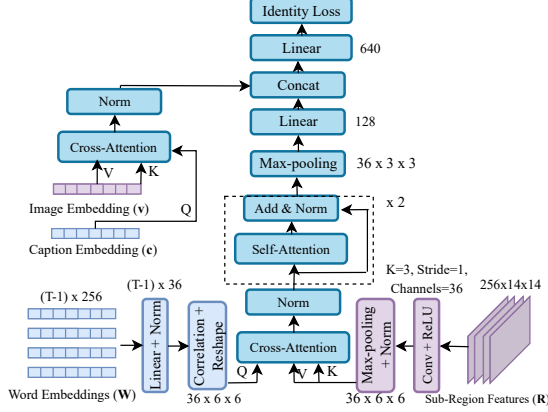


Figure 4. Block diagram of the proposed face-caption fusion module (FCFM) to capture complex fine-grained relationships and dependencies between regional features and word embeddings.

we generate two random views of the subject through random data augmentation. The following objective function represents  $L_{IMCL}$  loss.

$$L_{IMCL} = \frac{1}{2}(L_{IMCL}^v(\mathbf{x}_{f1}, \mathbf{x}_{f2}) + L_{IMCL}^c(\mathbf{x}_{c1}, \mathbf{x}_{c2})). \quad (7)$$

Here,  $L_{IMCL}^v$  represents the InfoNCE loss [33] between face image pairs  $(\mathbf{x}_{f1}, \mathbf{x}_{f2})$  where similarity is calculated using the function,  $s(\mathbf{x}_{f1}, \mathbf{x}_{f2}) = \langle f_v(g_v(\mathbf{x}_{f1})), f_v(g_v(\mathbf{x}_{f2})) \rangle$ . Here,  $f_v$  is the image encoder and  $g_v$  is the projection network. Likewise,  $L_{IMCL}^c$  represents the InfoNCE loss between caption pairs.

**Identity loss** To increase the discriminative power of both textual and facial features, we further employ a normalized cross-entropy loss [5], as our identity loss,  $L_{IDL}$ . This loss improves intra-class compactness while promoting inter-class dispersion. The total loss is the summation of the identity loss of both the textual and visual modalities.

**Overall loss** The following Eq. 8 delineates the objective function for the initial step of training process.

$$L_{total} = L_{WRCL} + \lambda_1 L_{IDL} + \lambda_2 L_{CICL} + \lambda_3 L_{IMCL}. \quad (8)$$

Here,  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are the hyperparameters.

### 3.5. Face caption fusion module (FCFM)

Figure 4 depicts the block diagram of the proposed FCFM module. The module incorporates a cross-attention between the word embeddings,  $\mathbf{W}$ , and regional image features,  $\mathbf{R}$ . This cross-attention mechanism [40] allows the TGFR to put attention on relevant image regions while processing a specific word. A self-attention [40] block is applied twice to the cross-modal feature which helps the

Table 1. A comprehensive list of face-to-image datasets.

Dataset	Public	Images	Annotations
Face2Text [39]	Yes	10,177	1~
MM CelebA-HQ [44]	Yes	30,000	10
FFHQ-Text [53]	Yes	760	9
CelebA-Dialog [22]	Yes	202,599	~5
SCU-Text2face [3]	No	1,000	5

TGFR to capture inherent dependencies and interactions within the input features. Afterward, a max-pooling layer is applied to the cross-modal features before they are fed into a linear layer. To capture the coarse-grained associations between global image-caption features, we employed an additional cross-attention mechanism [40], which is then followed by a normalization layer [1]. Lastly, the output of the linear layer is concatenated with the global cross-modal features before being fed into the final linear layer.

## 4. Experiments

### 4.1. Datasets and baselines

Table 1 presents a list of all available face-caption datasets. It is worth noting that the SCU-Text2face [3] dataset is not publicly available and FFHQ-Text [53] has only 760 pairs. Other face-caption datasets also have a limited number of face-caption pairs. This limitation arises due to the complexity of annotating face images, as mention in Section 1. In this work, we conduct experiments on three publicly available face-caption datasets, namely Multi-Modal CelebA-HQ [44], Face2Text [39], and CelebA-Dialog [22]. We consider the following baselines:

**BiLSTM-FLF** The bidirectional LSTM [35] is utilized as a text encoder and trained with an identity loss. The feature-level fusion scheme is applied using a linear layer.

**BERT-FLF** The pre-trained BERT [8] is employed as a text encoder and fine-tuned with an identity loss. Also, a linear layer is used for the FLF scheme.

### 4.2. Implementation details

Our TGFR model has two steps training process. In the initial step, we finetune the ALIGN encoder [21] and train the FCAM module, utilizing the objective function outlined in Eq. 8. We empirically set the hyperparameters  $\lambda_1 = 100$ ,  $\lambda_2 = 2$ , and  $\lambda_3 = 1$ . The ALIGN encoder [21] is fine-tuned for 20 epochs using an Adam optimizer [25] with a weight decay of 0.01, and a batch size of 16. In parallel, for the projection network of both branches, we employ another Adam optimizer [25], initialized with a learning rate of 0.001. Furthermore, for the ArcFace [5] image encoder, the backbone iResNet18 network was pre-trained on the MS1MV3 dataset [6] whereas, for the AdaFace [23] image encoder, the same backbone was pre-trained on the WebFace4M dataset [54]. In the subsequent step, the FCFM

Table 2. Comparison between the proposed TGFR framework and the baselines for 1:1 verification rate and Rank-1 Identification.

Methods	MMCelebA [44] Ver. (%)				Face2Text [39] Ver. (%)				CelebA-Dialog [22] Ver. (%)			
	1e-5	1e-4	1e-3	Id.(%)	1e-5	1e-4	1e-3	Id.(%)	1e-5	1e-4	1e-3	Id.(%)
Only AdaFace [23]	59.90	70.75	81.35	16.43	54.65	57.90	63.30	8.38	4.29	10.32	19.38	1.01
BiLSTM-FLF	66.20	78.60	85.0	16.43	57.40	62.40	69.0	8.38	17.47	31.66	48.60	4.77
BERT-FLF	65.50	78.0	85.80	16.43	60.60	63.0	70.40	8.38	22.78	33.51	49.31	5.96
CGFR [14]	65.80	81.0	87.60	18.90	61.60	65.0	71.80	9.97	22.36	34.88	50.39	7.39
Ours	<b>68.20</b>	<b>81.0</b>	<b>88.20</b>	<b>21.86</b>	<b>64.29</b>	<b>67.81</b>	<b>74.52</b>	<b>14.33</b>	<b>25.50</b>	<b>37.20</b>	<b>53.0</b>	<b>10.49</b>
Only ArcFace [5]	49.47	55.71	66.56	24.65	46.10	52.14	59.09	8.38	4.90	8.11	14.13	5.96
BiLSTM-FLF	62.50	71.96	78.73	16.43	61.02	64.21	72.0	8.38	14.13	21.77	38.65	4.17
BERT-FLF	63.89	72.83	79.89	16.43	61.20	63.96	72.67	8.38	15.77	25.23	40.82	12.70
CGFR [14]	66.32	78.05	82.23	28.59	63.50	65.92	74.48	16.64	16.90	26.83	42.45	18.06
Ours	<b>67.72</b>	<b>78.73</b>	<b>84.47</b>	<b>33.94</b>	<b>64.28</b>	<b>67.06</b>	<b>76.85</b>	<b>21.21</b>	<b>19.02</b>	<b>27.84</b>	<b>44.64</b>	<b>22.0</b>




Figure 5. ROC curves depict the performance comparison of the proposed TGFR with other methods in terms of 1:1 verification rates.

module is trained for an additional 36 epochs. We employed an SGD optimizer with a weight decay of  $1e-4$  and a momentum of 0.9. The initial learning rate is 0.1 and is divided by 10 after the 6th and 24th epochs.

### 4.3. Evaluation on face recognition

In Table 2, we conduct experiments to evaluate the performance of our TGFR framework for both face verification and identification. As depicted in Figure 5, our framework substantially improves the verification rate (VR) by utilizing textual descriptions as compared to the baseline approaches and the pre-trained AdaFace model [23] across all three experimental datasets. For the MMCelebA dataset [44], our framework achieves a VR of 68.20% at  $TAR@FAR=1e-5$ , marking a 3.52% improvement over CGFR [14] and a remarkable 12.17% improvement over the AdaFace model. Moreover, for the identification task, our proposed TGFR attains a Rank-1 accuracy of 21.86%, surpassing CGFR [14] by 13.54% and the baselines by 24.84%. Similarly, on the Face2Text dataset [39], our TGFR achieves a 4.18% higher VR at  $TAR@FAR=1e-5$  and a 30.43% higher Rank-1 accuracy compared to CGFR [14]. The CelebA-Dialog dataset [22] is more challenging than the other two datasets. In this dataset, the proposed TGFR remarkably improves the Rank-1 accuracy of pre-trained AdaFace from 1.01% to 10.49%.

We also perform experiments on the TGFR using the pre-trained ArcFace encoder [5] across all three datasets.

Table 3. Ablation experiments on the objective function of the proposed FCAM module using the MMCelebA dataset.

Identity	CICL	WRCL	IMCL	1e-5	1e-4	1e-5
✓	-	-	-	64.67	78	86.35
✓	✓	-	-	65.84	80.50	87.17
✓	-	✓	-	66.64	81.0	87.70
✓	-	-	✓	65.23	78.60	87.00
✓	✓	✓	✓	<b>68.20</b>	<b>81.0</b>	<b>88.20</b>

In the MMCelebA dataset [44], TGFR achieves a VR of 67.72% at  $TAR@FAR=1e-5$ , achieving a 2.07% improvement over CGFR and a 26.95% improvement over the pre-trained ArcFace model. In addition, for the identification task, TGFR achieves a Rank-1 accuracy of 33.94%, outperforming CGFR by 15.76% and the baseline models by 51.59%. Similarly, on the Face2Text dataset [39] and the CelebA-Dialog dataset [22], the proposed TGFR, using ArcFace encoder, demonstrates superior performance compared to other approaches. These results validate the generalizability and robustness of our TGFR framework.

### 4.4. Ablation study

**Analysis of FCAM** We analyze the role of each loss function in the proposed FCAM module, as shown in Table 3. In all the experiments conducted on the MMCelebA dataset [44] using the AdaFace [23] encoder, the identity loss is included, which is pivotal for generating discrimina-



Figure 6. Face images of a subject across different image qualities.

Table 4. Ablation experiments on the text encoders

Encoders	1e-5	1e-4	1e-3
BERT-base [8]	66.0	81.0	87.60
ALIGN [21]	68.20	81.0	88.20
CLIP [34]	67.65	81.0	87.96
BLIP [26]	68.20	81.0	88.20
GroupViT [45]	68.0	81.0	88.06
FLAVA [36]	67.90	81.0	88.0

Table 5. Comparison of 1:1 VR on TGFR framework across various image qualities.

Score	Level	Method	EER	1e-5	1e-4	1e-3
4	normal	AdaFace [23]	5.88	59.90	70.75	81.35
		CGFRR [14]	4.11	65.80	81.0	87.60
		TGFR	4.00	68.20	81.0	88.20
3	poor	AdaFace [23]	10.59	35.66	41.41	53.41
		CGFRR [14]	9.88	41.50	44.14	61.08
		TGFR	9.11	44.80	46.61	63.12
2	bad	AdaFace [23]	16.30	23.75	31.55	40.10
		CGFRR [14]	15.10	31.23	35.85	48.35
		TGFR	14.68	32.31	37.20	50.92
1	very bad	AdaFace [23]	19.06	11.67	18.98	27.77
		CGFRR [14]	18.78	21.39	29.95	38.39
		TGFR	17.65	24.40	31.0	39.60

tive features. From Table 3, it is observed that the FCAM module without any contrastive loss experiences a decrease in VR. This proves the necessity for both intra- and inter-modal contrastive losses. Furthermore, under the evaluation metric of  $\text{TAR@FAR}=1e-5$ , WRCL demonstrates an improvement from 64.67% to 66.64% over the performance of the identity loss alone. In contrast to the 2.96% improvement achieved by WRCL, we also note a 1.78% improvement achieved by CICL and a 0.86% improvement achieved by IMCL. These results signify the idea that each contrastive loss has a positive impact on VR performance. By combining all four components, we observe the highest VR, surpassing the identity loss alone by 5.16%.

**Analysis of text encoder** We investigate the impact of different foundation models, such as CLIP [34], and ALIGN [21] on the proposed TGFR model in Table 4. While the VR of all the CLIP-like foundation models is similar, they outperform the BERT [8]. Therefore, in TGFR, we adopt the ALIGN model [21] as our text encoder.

**Effect of image quality** We investigate the impact of image qualities on the proposed TGFR model, utilizing the AdaFace model [23] as the image encoder, which is trained on the “normal” images of the MMCelebA dataset [44]. Figure 6 depicts a face image of a subject across various image qualities. In this study, we report the 1:1 VR using two evaluation metrics: EER, and  $\text{TAR@FAR}$ . Our analysis, presented in Table 5, leads to several key observations. First, we observe a drastic drop in the performance of the AdaFace model as image quality degrades. For instance, when transitioning from an image quality of “normal” (score 4) to “very bad” (score 1), the VR of the AdaFace substantially degraded from 59.90% to 11.67% at  $\text{TAR@FAR}=1e-5$ .

Secondly, across all levels of image qualities, our TGFR consistently achieves significantly higher VR compared to both CGFR [14] and the pre-trained AdaFace model. For instance, in the “very bad” image quality, the proposed TGFR achieves an EER score of 17.16%, which is 6.40% lower than CGFR [14] and 7.99% lower than the AdaFace model. Additionally, we get a 12.34% improvement compared to CGFR at  $\text{TAR@FAR}=1e-5$ , and a 3.39% improvement at  $\text{TAR@FAR}=1e-4$ . Thirdly, as the quality of the input image degrades, the performance gain in the proposed TGFR becomes higher. For example, in the case of “very bad” quality images, the improvement in VR of our TGFR reaches 52.17% at  $\text{TAR@FAR}=1e-5$  which is higher than the 20.40% improvement observed in “poor” quality images. These results show the efficacy of textual descriptions in boosting the performance of the FR model, specifically in low-quality face images that are corrupted with noise.

## 5. Conclusion

This paper presents TGFR, a framework designed to enhance the relative performance of existing FR algorithms using textual descriptions. In this framework, we present a face-caption alignment module that effectively handles the inherent heterogeneity between the visual and the semantic domain by increasing the mutual information between the local and the global features of the face-caption pair. Our FCAM employs multiple contrastive losses across different granularities to implement inter-modal alignment between cross-modal features and intra-modal alignment within each modality. Aligning image features with textual features not only ensures the capture of complete shared semantics but also helps to focus on the distinctive content of the facial image, leading to improved performance in FR tasks.

## 6. Acknowledgment

This material is based upon a work supported by the Center for Identification Technology Research and the National Science Foundation under Grant 1650474.



## References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [6](#)
- [2] Hangbo Bao, Wenhui Wang, Li Dong, Qiang Liu, Owais Khan Mohammed, Kriti Aggarwal, and et al. VLMo: Unified vision-language pre-training with mixture-of-modality-experts. In *Adv. Neural Inform. Process. Syst.*, 2022. [3](#)
- [3] Xiang Chen, Lingbo Qing, Xiaohai He, Xiaodong Luo, and Yining Xu. FTGAN: A fully-trained generative adversarial networks for text to face generation. *arXiv preprint arXiv:1904.05729*, 2019. [6](#)
- [4] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, and et al. UNITER: Universal image-text representation learning. In *Eur. Conf. Comput. Vis.*, pages 104–120. Springer, 2020. [3](#)
- [5] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: additive angular margin loss for deep face recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4685–4694, 2019. [1](#), [2](#), [4](#), [6](#), [7](#)
- [6] Jiankang Deng, Jia Guo, Debing Zhang, Yafeng Deng, Xiangju Lu, and Song Shi. Lightweight face recognition challenge. In *Int. Conf. Comput. Vis. Worksh.*, pages 2638–2646, 2019. [6](#)
- [7] Karan Desai and Justin Johnson. VirTex: Learning visual representations from textual annotations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11162–11173, 2021. [3](#)
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, pages 4171–4186, 2019. [2](#), [3](#), [4](#), [6](#), [8](#)
- [9] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, and et al. From captions to visual concepts and back. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1473–1482, 2015. [5](#)
- [10] Zheren Fu, Zhendong Mao, Yan Song, and Yongdong Zhang. Learning semantic relationship among instances for image-text matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15159–15168, 2023. [3](#)
- [11] Ester Gonzalez-Sosa, Julian Fierrez, Ruben Vera-Rodriguez, and Fernando Alonso-Fernandez. Facial soft biometrics for recognition in the wild: Recent works, annotation, and COTS evaluation. *IEEE Trans. Inf. Forensics Secur.*, 13(8):2001–2014, 2018. [1](#), [2](#)
- [12] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304, 2010. [3](#)
- [13] Mohammad Haghighat and Mohamed Abdel-Mottaleb. Low resolution face recognition in surveillance systems using discriminant correlation analysis. In *IEEE International Conference on Automatic Face & Gesture Recognition*, pages 912–917, 2017. [1](#)
- [14] Md Mahedi Hasan and Nasser Nasrabadi. Improving face recognition from caption supervision with multi-granular contextual feature aggregation. *arXiv preprint arXiv:2308.06866*, 2023. [1](#), [2](#), [3](#), [7](#), [8](#)
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9729–9738, 2020.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. [4](#)
- [17] Zhenliang He, Wangmeng Zuo, Meina Kan, Shiguang Shan, and Xilin Chen. AttGAN: Facial attribute editing by only changing what you want. *IEEE Trans. Image Process.*, 28(11):5464–5478, 2019. [2](#)
- [18] Shih-Cheng Huang, Liyue Shen, Matthew P Lungren, and Serena Yeung. GLoRIA: A multimodal global-local representation learning framework for label-efficient medical image recognition. In *Int. Conf. Comput. Vis.*, pages 3942–3951, 2021. [4](#)
- [19] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Int. Conf. Mach. Learn.*, pages 448–456, 2015. [4](#)
- [20] Anil K. Jain and Unsang Park. Facial marks: Soft biometric for face recognition. In *IEEE Int. Conf. Image Process.*, pages 37–40, 2009. [1](#)
- [21] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, and et al. Scaling up visual and vision-language representation learning with noisy text supervision. In *Int. Conf. Mach. Learn.*, pages 4904–4916, 2021. [3](#), [4](#), [6](#), [8](#)
- [22] Yuming Jiang, Ziqi Huang, Xingang Pan, Chen Change Loy, and Ziwei Liu. Talk-to-edit: Fine-grained facial editing via dialog. In *Int. Conf. Comput. Vis.*, pages 13799–13808, 2021. [6](#), [7](#)
- [23] Minchul Kim, Anil K Jain, and Xiaoming Liu. AdaFace: Quality adaptive margin for face recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 18750–18759, 2022. [1](#), [2](#), [4](#), [6](#), [7](#), [8](#)
- [24] Wonjae Kim, Bokyung Son, and Ildoo Kim. ViLT: Vision-and-language transformer without convolution or region supervision. In *Int. Conf. Mach. Learn.*, pages 5583–5594, 2021. [3](#)
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [26] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Int. Conf. Mach. Learn.*, pages 12888–12900, 2022. [3](#), [4](#), [8](#)
- [27] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Adv. Neural Inform. Process. Syst.*, 34:9694–9705, 2021. [2](#), [3](#)
- [28] Pei Li, Loreto Prieto, Domingo Mery, and Patrick J Flynn. On low-resolution face recognition in the wild: Compar-

- isons and new techniques. *IEEE Trans. Inf. Forensics Secur.*, 14(8):2000–2012, 2019. [1](#)
- [29] Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, and et al. UNIMO: Towards unified-modal understanding and generation via cross-modal contrastive learning. *59th Annual Meeting of the Association for Computational Linguistics*, pages 2592–2607, 2021. [3](#)
- [30] Y. Liu, M. Ott, Naman G., J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyano. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. [3](#), [4](#)
- [31] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Adv. Neural Inform. Process. Syst.*, 32, 2019. [3](#)
- [32] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. MagFace: A universal representation for face recognition and quality assessment. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 14225–14234, June 2021. [2](#)
- [33] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [3](#), [4](#), [6](#)
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, and et al. Learning transferable visual models from natural language supervision. In *Int. Conf. Mach. Learn.*, pages 8748–8763, 2021. [3](#), [4](#), [8](#)
- [35] Mike Schuster and Kuldeep K Paliwal. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.*, 45(11):2673–2681, 1997. [6](#)
- [36] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. FLAVA: A foundational language and vision alignment model. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15638–15650, 2022. [4](#), [8](#)
- [37] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: pre-training of generic visual-linguistic representations. In *Int. Conf. Learn. Represent.*, 2020. [3](#)
- [38] Jianxin Sun, Qiyao Deng, Qi Li, Muye Sun, Min Ren, and Zhenan Sun. AnyFace: Free-style text-to-face synthesis and manipulation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 18687–18696, 2022. [2](#)
- [39] Marc Tanti, Shaun Abdilla, Adrian Muscat, Claudia Borg, Reuben A Farrugia, and Albert Gatt. Face2Text revisited: Improved data set and baseline results. *Proceedings of the workshop on People in Vision, Language, and the Mind*, pages 41–47, 2022. [6](#), [7](#)
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, and et al. Attention is all you need. *Adv. Neural Inform. Process. Syst.*, 30, 2017. [4](#), [6](#)
- [41] Fuying Wang, Yuyin Zhou, Shujun Wang, Varut Vardhanabhuti, and Lequan Yu. Multi-granularity cross-modal alignment for generalized medical visual representation learning. *Adv. Neural Inform. Process. Syst.*, 35:33536–33549, 2022. [3](#), [4](#)
- [42] Yixuan Wei, Yue Cao, Zheng Zhang, Zhuliang Yao, Zhenda Xie, Han Hu, and Baining Guo. iCAR: Bridging image classification and image-text alignment for visual recognition. *arXiv preprint arXiv:2204.10760*, 2022. [3](#)
- [43] Keyu Wen, Xiaodong Gu, and Qingrong Cheng. Learning dual semantic relations with graph attention for image-text matching. *IEEE Trans. Circuit Syst. Video Technol.*, 31(7):2866–2879, 2020. [3](#)
- [44] Weihao Xia, Yujiu Yang, Jing-Hao Xue, and Baoyuan Wu. TediGAN: Text-guided diverse face image generation and manipulation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2256–2265, 2021. [2](#), [3](#), [6](#), [7](#), [8](#)
- [45] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. GroupViT: Semantic segmentation emerges from text supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 18134–18144, 2022. [8](#)
- [46] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1316–1324, 2018. [3](#), [5](#)
- [47] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Chen, and et al. Vision-language pre-training with triple contrastive learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15671–15680, 2022. [3](#)
- [48] Xi Yin, Ying Tai, Yuge Huang, and Xiaoming Liu. FAN: Feature adaptation network for surveillance face recognition and normalization. In *As. Conf. Comput. Vis.*, 2020. [1](#)
- [49] Alireza Zaeemzadeh, Shabnam Ghadar, Baldo Faieta, Zhe Lin, Nazanin Rahnavard, Mubarak Shah, and et al. Face image retrieval with attribute manipulation. In *Int. Conf. Comput. Vis.*, pages 12116–12125, 2021. [2](#)
- [50] Hao Zhang, J Ross Beveridge, Bruce A Draper, and P Jonathon Phillips. On the effectiveness of soft biometrics for increasing face verification rates. *Computer Vision and Image Understanding*, 137:50–62, 2015. [1](#)
- [51] Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. Contrastive learning of medical visual representations from paired images and text. In *Machine Learning for Healthcare Conference*, pages 2–25, 2022. [3](#), [4](#)
- [52] Ying Zhang and Huchuan Lu. Deep cross-modal projection learning for image-text matching. In *Eur. Conf. Comput. Vis.*, pages 686–701, 2018. [3](#)
- [53] Yutong Zhou. Generative adversarial network for text-to-face synthesis and manipulation. In *ACM Int. Conf. Multimedia*, pages 2940–2944, 2021. [6](#)
- [54] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, and et al. WebFace260M: A benchmark unveiling the power of million-scale deep face recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10492–10502, 2021. [6](#)