Facial Attractiveness Prediction in Live Streaming: A New Benchmark and Multi-modal Method

Hui Li[1†], Xiaoyu Ren[2†], Hongjiu Yu[1], Huiyu Duan[2], Kai Li[1],
Ying Chen[1], Libo Wang[1], Xiongkuo Min[2], Guangtao Zhai[2], Xu Liu[1*]
[1]Alibaba Group [2]Shanghai Jiao Tong University
{lh401704,yuhongjiu.yhj,kaishi.lk,YingChen,zhuangshu.wlb,steven.liu}@taobao.com
{windkaiser,huiyuduan,minxiongkuo,zhaiguangtao}@sjtu.edu.cn

Abstract

Facial attractiveness prediction (FAP) has long been an important computer vision task, which could be widely applied in live streaming for facial retouching, content recommendation, etc. However, previous FAP datasets are either small, closed-source, or lack diversity. Moreover, the corresponding FAP models exhibit limited generalization and adaptation ability. To overcome these limitations, in this paper we present LiveBeauty, the first large-scale live-specific FAP dataset, in a more challenging application scenario, i.e., live streaming. 10,000 face images are collected from a live streaming platform directly, with 200,000 corresponding attractiveness annotations obtained from a well-devised subjective experiment, making LiveBeauty the largest open-access FAP dataset in the challenging live scenario. Furthermore, a multi-modal FAP method is proposed to measure the facial attractiveness in live streaming. Specifically, we first extract holistic facial prior knowledge and multi-modal aesthetic semantic features via a Personalized Attractiveness Prior Module (PAPM) and a Multi-modal Attractiveness Encoder Module (MAEM), respectively, then integrate the extracted features through a Cross-Modal Fusion Module (CMFM). Extensive experiments conducted on both LiveBeauty and other open-source FAP datasets demonstrate that our proposed method achieves state-of-the-art performance. Dataset will be available soon.

1. Introduction

Face image quality assessment plays an important role in face recognition, face analysis, *etc*. [1–3], which could be categorized into two sub-fields, including Biometric Face Image Quality Assessment (BFIQA) and Generic Face Image Quality Assessment (GFIQA) [4]. BFIQA assesses whether face images can keep identity information and are suitable for down-stream tasks such as face recognition [7–9], while GFIQA pays more attention to visual perceptual quality of face images [1, 2, 10]. Facial Attractiveness Prediction (FAP) is an important branch in GFIQA [5, 11, 12]. FAP aims to predict the facial attractiveness of face images, which has crucial implication in various fields [13]. Considering face images are widely present in live streaming and facial regions are generally the visual focus of users [14], predicting the facial attractiveness is also significant for live streaming applications, which can help to monitor and improve the quality and attractiveness of live videos [13, 15, 16], recommend live content [17, 18], *etc*. However, previous FAP studies mainly focus on the beauty industry, psychological research, *etc*. [13, 19]. Research on facial attractiveness prediction in live streaming is still lacking, despite its importance.

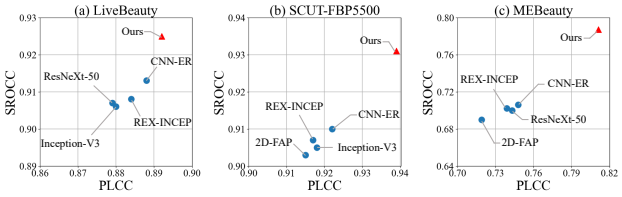

Figure 1. SROCC vs PLCC comparison on LiveBeauty, SCUT-FBP5500 [5] and MEBeauty [6]. Our proposed model achieves the best performance on all three datasets.

As shown in Tab. 1, numerous FAP datasets have been proposed in recent years. Most of these datasets are small or unavailable [21–25]. Some open-source datasets [5, 6, 20] primarily contains face portraits, with limited diversity and expressions. However, live scenarios are more challenging, due to the unpredictable conditions such as lighting, expressions, camera angles and background noise, *etc*., making the previous datasets unable to generalize to live applications. Considering the notable gap between face images in differ-

*Corresponding author.
†Equal contribution.

Table 1. Overview of common facial attractiveness prediction datasets.

Dataset	Year	Rating Scale	# Image	Gender	Ethnicity	# Annotator	Public
Gray *et al*. [20]	2010	[-3, 3]	2,056	F M	Caucasian	30	✓
MBW+FBW [21]	2014	[1, 5]	10,400	F M	Multi	30	-
De Vries *et al*. [22]	2015	Binary	9364	F M	Asian	20	-
LSFBD [23]	2016	[1, 5]	20,000	F M	-	-	-
SCUT-FBP5500 [5]	2018	[1, 5]	5,500	F M	Multi	60	✓
Tong *et al*. [24]	2020	Binary	4,512	F M	-	20	-
LSAFBD [25]	2020	[1, 5]	20,000	F	Asian	20	-
MEBeauty [6]	2022	[1, 10]	2,550	F M	Multi	300	✓
LiveBeauty (proposed)	2024	[1, 5]	10,000	F M	Asian	20	✓

ent contexts [6], the FAP dataset specifically collected in live streaming with diverse conditions is still lacking.

Also, there is an absence of a widely recognized FAP model suitable for live streaming applications. Specifically, early methods exhibit poor performance due to the limited feature representation capabilities, relying on *e.g*., handcrafted beauty descriptors [26, 27] or basic models like Support Vector Regression (SVR) [11]. The deep-learning methods based on Convolution Neural Networks (CNNs) or Transformers [12, 28–34] show limited generalization ability, partially due to the gaps of data distribution. More importantly, most of these methods [31–33, 35] utilize the simplistic feature fusion strategy and thus overlook the contextual and individual facial features, leading to poor adaption ability. The straightforward feature concatenation without context, including personalized facial priors or multi-scale features, fails to model individual facial features dynamically, thereby influencing the generalization ability.

To address the lacking of FAP dataset in live streaming applications, we first construct a large-scale live streaming facial attractiveness dataset **LiveBeauty**, which includes 10,000 face images collected from a popular live streaming platform, with the corresponding facial attractiveness annotations obtained from 20 annotators. To ensure the quality of the constructed dataset, a well-devised five-step image auto-sampling pipeline, well-designed subjective experiments and a data cleansing method [36] are employed.

Furthermore, based on the established dataset, we propose a no-reference FAP model termed Facial Prior Enhanced Multi-modal model (**FPEM**), to predict the facial attractiveness in live streaming by extracting and integrating holistic facial prior knowledge and multi-modal aesthetic semantic features. Specifically, we first devise a Personalized Attractiveness Prior Module (PAPM) to integrate extensive facial prior knowledge, in which multi-scale visual features extracted by Swin Transformer [37, 38] and facial prior features extracted by pretrained FaceNet [39] are integrated via a Cross-Attention block. To obtain face beauty attributes, we also introduce a Multi-modal Attractiveness Encoder Module (MAEM) based on CLIP [40] to extract multi-modal aesthetic semantic features. Furthermore, a Cross-Modal Fusion Module (CMFM) is proposed to integrate and refine the extracted features. Extensive experimental results confirm that our proposed method achieves the best performance compared to other state-of-the-art (SOTA) methods on both our LiveBeauty dataset and other FAP datasets, as shown in Fig. 1.

To summarize, our contributions are as follow:

- We highlight the **significance of Facial Attractiveness Prediction (FAP) in live streaming**, and emphasize the lacking of corresponding research. To the best of our knowledge, this is the first study to analyze and model the facial attractiveness in live streaming.
- We present **LiveBeauty**, the first large-scale, open-source facial attractiveness prediction dataset for live streaming, which contains 10,000 live face images and the collected Mean Opinion Scores (MOSs).
- We propose a **multi-modal method termed FPEM** specifically designed for FAP, which predicts the perceptual facial attractiveness by leveraging facial prior knowledge and multi-modal aesthetic semantic features, and further employing cross-modal fusion to integrate and refine the extracted features. Comprehensive experiments demonstrate that our method achieves the **state-of-the-art performance** on common FAP datasets.

2. Related Works

2.1. Aesthetic Assessment of Generic Images

Image Aesthetic Assessment (IAA) aims to computationally assess the image quality from the aesthetic perspective. Early IAA methods employ carefully designed handcrafted features to model the photographic rules [41, 42], global image layout [15, 43] and typical objects [44] in the perceived images [45]. With the emergence of data-driven methods, the IAA models based on deep learning surpass the handcrafted methods notably [46–49]. Though IAA methods can evaluate the facial beauty to some extent, the existing methods mainly focus on generic images rather than face images, which limits their performance on FAP tasks.

2.2. Quality Assessment of Face Images

BFIQA. Biometric Face Image Quality Assessment (BFIQA) originated from biometric research focus on evaluating the suitability of facial images for recognition. Initially relying on hand-crafted features [50–53], recent advances in deep learning have improved BFIQA with regression and embedding-based solutions [9, 53–58]. Nonetheless, Chen *et al*. [4] noted these methods may not generalize well on other facial quality and aesthetics tasks due to the emphasis on facial biometric information.

GFIQA. Generic Face Image Quality Assessment (GFIQA) is a newly defined task that evaluates the conventional perceptual quality of face images. Facial prior knowledge like landmarks and parsing maps are integrated with visual features and degradation representations [1, 2, 4, 59, 60]. Furthermore, GFIQA methods have advanced through the use of transformer-based [4, 59] and multi-modal approaches [3]. However, emphasis on degradation of these methods results in poorer performance on FAP tasks.

2.3. Attractiveness Assessment of Face Images

To measure the facial attractiveness of face images, early handcrafted Facial Attractiveness Prediction (FAP) methods predominantly rely on manually designed features, *e.g*. geometric features [11, 26, 27], traditional image descriptors [5, 11, 26], holistic descriptors [11], *etc*. However, these methods exhibit poor performance due to difficulties in presenting comprehensive characteristics of facial attractiveness.

With the advancement in deep learning, FAP solutions based on Convolution Neural Networks (CNNs) show better performance. Inspired by psychological studies, specific architectures of CNNs are designed [28, 61]. The applications of Label Distribution Learning (LDL) [29, 62, 63] achieve better performance. Some methods [30, 64, 65] combine facial prior knowledge containing landmarks, parsing maps, *etc*., while the others [12, 66] integrate attractiveness-related attributes such as gender, race, *etc*. In addition, the ensemble solutions of different loss functions like regression loss and classification loss [67] and different models including ResNeXt50 and Inception-V3 [35] are further explored. However, these CNN-based methods primarily focus on local features, limiting their capacity for long-term modeling. To address this problem, transformer-based solutions [68, 69] are proposed, while Liu *et al*. [33] and Gan *et al*. [70] combine transformers with CNNs to leverage both global and local features related to facial attractiveness. Nevertheless, these approaches have limited generalization and adaptation ability in more challenging scenarios, due to the biased predictions from training data skewed towards certain demographics, and using straightforward single-modal modeling techniques which rely on simplistic concatenation of visual features.

Figure 2. Examples of the face images in our LiveBeauty dataset (from left to right are examples rated 1 to 5).

3. Dataset

3.1. Images Collection

To accurately reflect the distribution of facial attractiveness on live streaming platforms, we first collect large-scale raw face images from one of the most popular live streaming platforms. We retrieve replays for live broadcasts of all streamers conducted on the platform in March 2024 with the highest Page Views (PVs). Given that each broadcast lasts approximately 2 to 10 hours, we sample one image per hour for the first three hours of each broadcast, resulting in three images sampled from each broadcast.

To ensure the collection of high-quality face images suitable for facial attractiveness prediction, we propose a five-step auto-sampling pipeline: (i) Face region size measurement is accomplished using a SOTA face detection model [71], which generates the bounding box (bbox) of the face region, and we ensure the shorter boundary of the bbox is greater than 90. (ii) Blur detection is applied by ensuring the variance of the Laplacian operator in the Y channel of the face region is greater than 10. (iii) Face pose estimation is accomplished using a SOTA pose estimation model [72], and we ensure the pitch angle is no greater than 20, and the yaw angle is no greater than 15. (iv) Face proportion assessment is accomplished using a SOTA face segmentation model [72], and we ensure the face region proportion is greater than 60%. (v) Duplicate character removal is accomplished using a SOTA face recognition model.

We further conduct human evaluation to eliminate samples containing virtual characters, incomplete facial features, abnormal expressions that affect attractiveness measurement and repetitive characters. Finally, 10,000 images are collected, ensuring 10,000 different faces from over 9,500 broadcasts. Fig. 2 shows some examples of the face images in our LiveBeauty dataset.

3.2. Subjective Experiments

We recruit 20 subjects in our experiments, including 14 females and 6 males. This ratio is determined by the actual customer demographics of the live platform. Face stimulis are displayed on an iPhone 14 Pro Max with 6.7-inch Or-

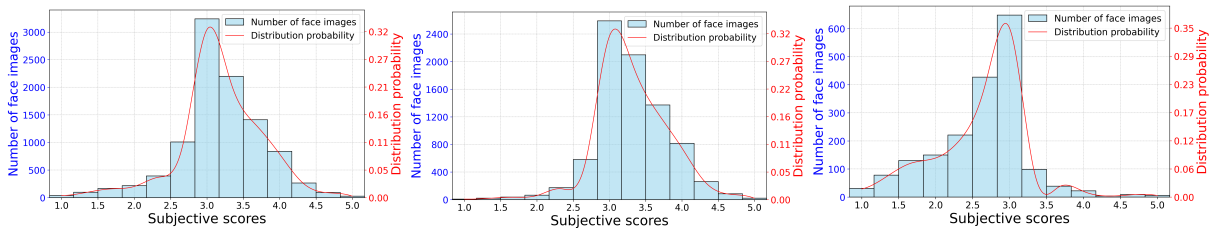

(a) MOS distribution of all samples.Skewness=-0.470 Kurtosis=1.498

(b) MOS distribution of female samples.Skewness=-0.083 Kurtosis=1.499

(c) MOS distribution of male samples.Skewness=-0.281 Kurtosis=1.123

Figure 3. Illustration of the proposed LiveBeauty MOS distributions from different perspectives.

ganic Light-Emitting Diode (OLED) screen, supporting a resolution up to 2,796 $\times$ 1,290. Subjective experiments are conducted in a well-controlled lab environment with fixed settings [36], *i.e.*, lighting, viewing distance, viewing angle, *etc*.

We split the experiment into 200 sessions with 50 images per session. Each face image is displayed for 10 seconds and followed by a 1-second gray screen mask. Before the experiment starting, we offer a brief explanation and 50 extra training samples for subjects to familiarize them with the task. During the experiment, subjects are asked to rate the facial attractiveness of the stimulis as distributed scores from 1 to 5. After each session, subjects are asked to take at least 5 minutes of break. All subjects participate in all sessions, and each subject is allowed to participate in no more than 5 sessions each day. Ultimately, all 10,000 images are evaluated by 20 subjects, resulting in a total of 200,000 (10,000 $\times$ 20) scores collected.

3.3. Data Analysis

According to [36], we employ subject post-screening using outlier ratio and Spearman Rank Correlation Coefficient (SROCC). A subject whose rated scores with SROCC $<$ 0.75 or outlier_ratio $>$ 2% will be marked as unreliable and will be rejected. Ultimately, 20 subjects are retained. We compute the Mean Opinion Score (MOS) as the attractiveness label of each face image:

$$\text{MOS}_j = \frac{1}{N}\sum_{i=1}^{N} r_{ij}, \tag{1}$$

where MOS_j represents the MOS for the j-th face image, N is the number of the valid subjects, and r_{ij} is the rated score of the i-th subject on the j-th face image.

MOS distribution of our LiveBeauty is shown in Fig. 3. It can be observed that the distributions of MOSs for all samples, as well as for female and male samples, exhibit a Gaussian-like shape. Skewness and kurtosis values of three MOS distributions are also shown in Fig. 3, which indicate that the distributions are characterized by thin tails and the MOS is concentrated around 3. This aligns with real-world facial attractiveness distributions, *i.e.*, most individuals possess average facial attractiveness. Additionally, Fig. 3b indicates both higher skewness and higher kurtosis than Fig. 3c, demonstrating high attractiveness in females is more prevalent in the collected live streaming videos.

4. Proposed Method

As illustrated in Fig. 4, our Facial Prior Enhanced Multimodal (FPEM) method utilizes a two-stage training strategy, and consists of four primary modules.

In the first stage, which is called Preliminary Training Phase, the **Personalized Attractiveness Prior Module** (PAPM in Sec. 4.1.1) receives the image I, in which a Swin Transformer E_{ms} extracts multi-scale feature f_Σ, and a frozen FaceNet E_{fp} extracts face-aware feature f_s. Subsequently, a Face-Aware Cross-Attention block incorporates f_Σ and f_s to form the personalized attractiveness feature f_p, which are then projected to $\widehat{Q}_1$. The **Multi-modal Attractiveness Encoder Module** (MAEM in Sec. 4.1.2) contains an image encoder E_i and a text encoder E_t, receives the image I and progressive texts, leverages multi-modal embeddings via CLIP [40], and obtains $\widehat{Q}_2$ by the weighted sum of five-level score and corresponding probability that generated by calculating the cosine similarities between textual and visual embeddings, which is called Similarity Regression.

In the second stage, which is called Hybrid Fusion Phase, the **Cross-Modal Fusion Module** (CMFM in Sec. 4.2.1) dynamically refines the extracted textual embeddings f_t to f_f based on the personalized attractiveness feature f_p, and employs the Similarity Regression strategy to predict $\widehat{Q}_3$. The **Decision Fusion Module** (DFM in Sec. 4.2.2) obtains the final attractiveness score $\widehat{Q}$, considering dynamically with three predicted scores $\widehat{Q}_1$, $\widehat{Q}_2$ and $\widehat{Q}_3$.

4.1. Preliminary Training Phase

4.1.1 Personalized Attractiveness Prior Module

Stem Swin. Considering that the visual perception progress is perceived from low level to high level, the hierarchical features from consecutive transformer stages are fully

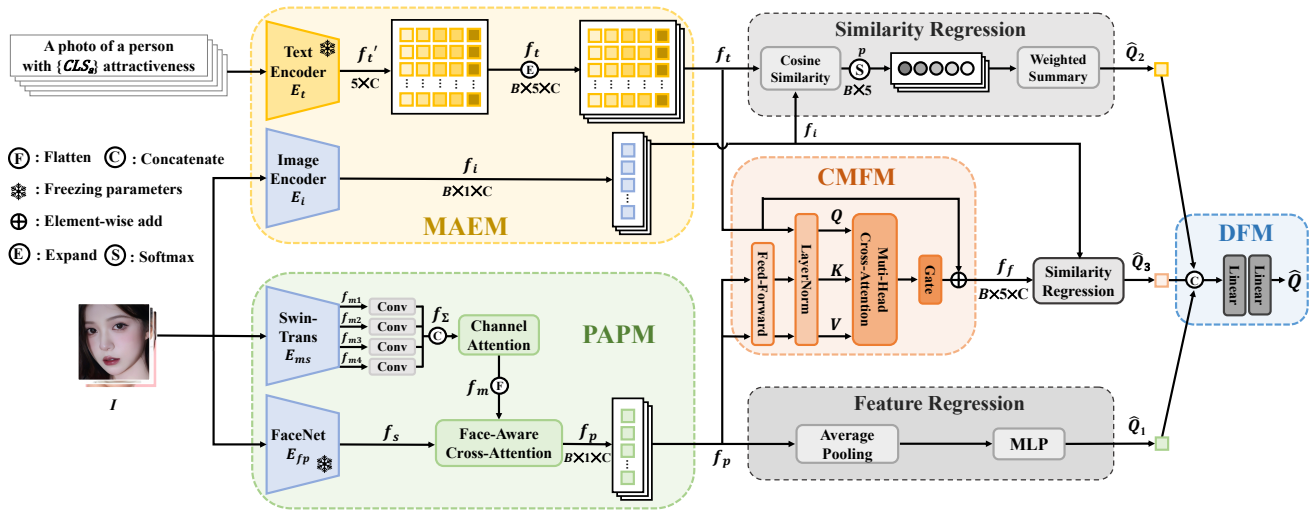

Figure 4. Architecture of our method, including the Personalized Attractiveness Prior Module (PAPM), the Multi-modal Attractiveness Encoder Module (MAEM), the Cross-Modal Fusion Module (CMFM) and the Decision Fusion Module (DFM).

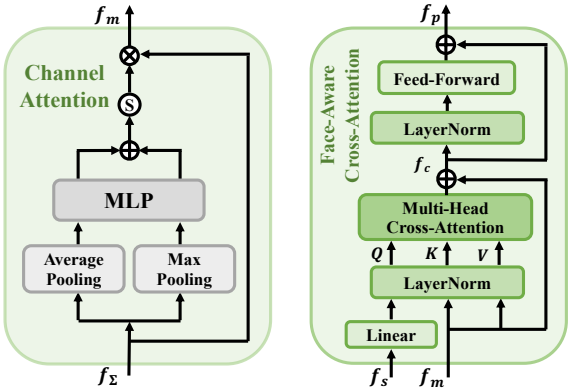

Figure 5. Structure of the Channel Attention and the Face-Aware Cross-Attention Block.

leveraged. Specifically, the features $f_{m_i} \in \mathbb{R}^{H_i \times W_i \times C_i}$ extracted from i-th stage are reshaped into same size, and are then concatenated as $f_\Sigma \in \mathbb{R}^{H_4 \times W_4 \times C_\Sigma}$. Meanwhile, f_Σ is refined by a Channel Attention block from channel perspective (which is shown in Fig. 5), and then flattened as $f_m \in \mathbb{R}^{N_4 \times C_\Sigma}$.

Face-Aware Cross-Attention. Facial prior knowledge, *i.e.*, face-aware features f_s extracted by FaceNet, are utilized in this block to guide the learning of facial attractiveness. Specifically, the refined features f_m extracted by the stem Swin are fused with the face-aware features f_s through cross-attention. As shown in Fig. 5, f_s is encoded as the query Q via a Linear layer and a Layer Normalization (LN) layer, while f_m is encoded as the key K and value V via a LN layer. Further, a Multi-Head Cross-Attention (MHCA) mechanism is applied, which is computed as:

$$\begin{cases} \text{head}_i = \psi(\frac{Q \cdot K^T}{\sqrt{d}}) \cdot V \cdot W_i, i \in \{1, ..., h\} \\ \text{MHCA}(Q, K, V) = \text{Concat}(\text{head}_i) W_o, \end{cases} \tag{2}$$

where $\psi(\cdot)$ denotes the softmax function and d denotes the query dimension. Let FFN indicate the Feed-Forward Network, the output feature f_p is computed as:

$$\begin{cases} f_c = f_m + \text{MHCA}(Q, K, V) \\ f_p = f_c + \text{FFN}(\text{LN}(f_c)). \end{cases} \tag{3}$$

Let MLP represent the Multi Layer Perceptron, then the attractiveness score $\widehat{Q}_1$ can be predicted by:

$$\widehat{Q}_1 = \text{MLP}(\text{AveragePooling}(f_p)). \tag{4}$$

4.1.2 Multi-modal Attractiveness Encoder Module

Inspired by [73], MAEM is devised to utilize multi-modal aesthetic semantic feature. We utilize five attractiveness levels $a \in \mathcal{A} = \{1, 2, 3, 4, 5\}$ ={"*bad*", "*poor*", "*fair*", "*good*", "*perfect*"}. Then the text template can be formed as "*a photo of a person with* $\{a\}$ *attractiveness*", resulting in five candidate text descriptions $\mathcal{T}$.

MAEM contains a frozen text encoder E_t and a trainable image encoder E_i, where E_t receives five candidate text descriptions $\mathcal{T}$ to form the textual embeddings $f_t \in \mathbb{R}^{5 \times 512}$, while E_i receives the image I to form the visual embeddings $f_i \in \mathbb{R}^{1 \times 512}$. Subsequently, the cosine similarity between textual embeddings f_t and visual embeddings f_i is calculated to obtain the attractiveness level probability.

Then the attractiveness scores $\widehat{Q}_2$ can be predicted by:

$$\widehat{Q}_2 = \sum_{i=1}^{A} \hat{p}(a|I) \cdot a, \tag{5}$$

where A denotes the number of attractiveness levels and $\hat{p}(a|\cdot)$ is the probability of the level a to be estimated.

4.2. Hybrid Fusion Phase

4.2.1 Cross-Modal Fusion Module

In this stage, we freeze MAEM and PAPM. Since our five candidate text descriptions $\mathcal{T}$ are general, resulting in general textual embeddings f_t, CMFM is devised to refine f_t according to the personalized attractiveness feature f_p extracted by PAPM. Therefore, personalized textual embeddings f_f are generated for similarity regression.

As depicted in Fig. 4, f_t is encoded as the query Q via a LN layer, while f_p is encoded as the key K and value V via a Feed-Forward Network and a LN layer. Further, the MHCA mechanism is applied same as Eq. (2). Subsequently, personalized textual embeddings f_f are acquired through a residual operation utilizing an empirical gate ratio $\alpha = 0.7$, the prediction $\widehat{Q}_3$ is computed same as Eq. (5).

4.2.2 Decision Fusion Module

To achieve a more robust common consensus, DFM is constructed to consider all predicted candidates. In particular, the three predicted scores $\widehat{Q}_1$, $\widehat{Q}_2$ and $\widehat{Q}_3$, are combined through two learnable linear layers to generate the final attractiveness score $\widehat{Q}$.

4.3. Training objectives

Loss for PAPM. The L1 loss between $\widehat{Q}_1$ and the ground truth label Q is calculated as the attractiveness prediction loss function $\mathcal{L}_1$, which can be formulated as:

$$\mathcal{L}_1 = \frac{1}{n} \sum_{i=1}^{n} \left| \widehat{Q}_{1_i} - Q_i \right|, \tag{6}$$

where n denotes number of training samples.

Loss for MAEM. To jointly optimize ranking and regression, we incorporate a merged ranking loss $\mathcal{L}_R$ along with the L1 scoring loss $\mathcal{L}_S$. Specifically, $\mathcal{L}_R$ is composed of a fidelity loss $\mathcal{L}_{R_1}$ and a two-direction ranking loss $\mathcal{L}_{R_2}$. For the fidelity loss $\mathcal{L}_{R_1}$, given an image pair (I_i, I_j), the binary label according to their ground truth (Q_i, Q_j) is defined as:

$$R(I_i, I_j) = \begin{cases} 1 & \text{if } Q_i \geq Q_j \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

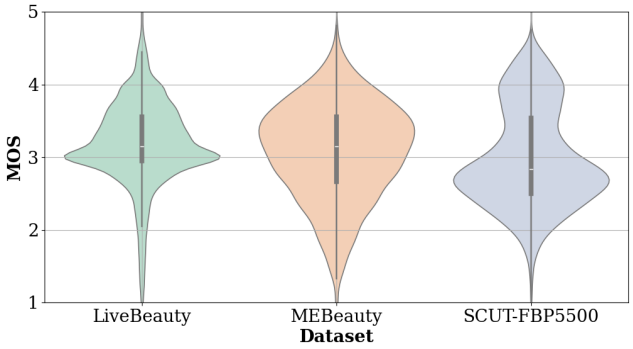

Figure 6. MOS distribution of three benchmark datasets.

According to the Thurstone's model [74], the probability of I_i perceiving more attractive than I_j is estimated as:

$$\widehat{R}(I_i, I_j) = \Psi\left(\frac{\widehat{Q}_{2_i} - \widehat{Q}_{2_j}}{\sqrt{2}}\right), \tag{8}$$

where $\Psi(\cdot)$ denotes the standard normal cumulative distribution function, and its variance is set to 1. Then the fidelity loss [75] can be calculated as:

$$\begin{aligned} \mathcal{L}_{R_1} = 1 &- \sqrt{R(I_i, I_j)\widehat{R}(I_i, I_j)} \\ &- \sqrt{(1 - R(I_i, I_j))(1 - \widehat{R}(I_i, I_j))}. \end{aligned} \tag{9}$$

For the two-direction ranking loss $\mathcal{L}_{R_2}$, given an image I_i, $\widehat{S}_{i,a}$ is the highest cosine similarity result of all five matchings. Since the human perception of facial attractiveness is gradual and rarely exhibits multiple peaks [76], we further obtain the ranking information by ensuring that the attractiveness level probability distribution only exhibits one peak and decreases in two directions (leftward and rightward). Then the two-direction ranking loss can be formulated as:

$$\begin{cases} \mathcal{L}_{\text{left}} = -\sum_{j=2}^{a-1} \log \frac{exp(\widehat{S}_{i,j})}{exp(\widehat{S}_{i,j}) + exp(\widehat{S}_{i,j-1})} \\ \mathcal{L}_{\text{right}} = -\sum_{j=a}^{A-1} \log \frac{exp(\widehat{S}_{i,j})}{exp(\widehat{S}_{i,j}) + exp(\widehat{S}_{i,j+1})} \\ \mathcal{L}_{R_2} = -(\mathcal{L}_{\text{left}} + \mathcal{L}_{\text{right}}), \end{cases} \tag{10}$$

where the labels of $\widehat{S}_{i,j-1}$, $\widehat{S}_{i,j}$ and $\widehat{S}_{i,j+1}$ are 0, 1, 0, respectively. In the end, the total loss of MAEM can be defined as $\mathcal{L}_2 = \mathcal{L}_S + \lambda_1 \mathcal{L}_{R_1} + \lambda_2 \mathcal{L}_{R_2}$, where λ_1 and λ_2 are empirically set as 1.

Loss for CMFM and DFM. The loss used in the second stage between the ground truth label Q and the final attractiveness score $\widehat{Q}$ derived from three predicted scores is the L1 Loss:

$$\mathcal{L}_3 = \frac{1}{n} \sum_{i=1}^{n} \left| \widehat{Q}_i - Q_i \right|. \tag{11}$$

Table 2. Performance comparison on LiveBeauty, MEBeauty [6] and SCUT-FBP5500 [5]. Best in **red** and second-best in **blue**.

Aspects	Methods	LiveBeauty			MEBeauty [6]			SCUT-FBP5500 [5]		
		SROCC↑	PLCC↑	KROCC↑	SROCC↑	PLCC↑	KROCC↑	SROCC↑	PLCC↑	KROCC↑
Baseline	ViT-B [77]	0.897	0.847	0.733	0.620	0.659	0.443	0.869	0.880	0.686
	ResNeXt-50 [78]	0.907	0.879	0.746	0.700	0.743	0.516	0.899	0.911	0.727
	Inception-V3 [79]	0.906	0.880	0.745	0.676	0.727	0.492	0.905	0.918	0.736
IAA	AVA-MLSP [46]	0.833	0.815	0.650	0.602	0.643	0.431	0.764	0.794	0.566
	TANet [47]	0.891	0.858	0.728	0.651	0.690	0.467	0.870	0.882	0.686
	Dele-Trans [80]	0.888	0.841	0.721	0.611	0.651	0.440	0.870	0.860	0.682
	EAT [81]	0.893	0.867	0.728	0.640	0.697	0.462	0.853	0.872	0.665
FAP	ComboNet [31]	0.902	0.871	0.741	0.651	0.692	0.472	0.897	0.907	0.725
	2D-FAP [32]	0.896	0.854	0.733	0.690	0.719	0.506	0.903	0.915	0.734
	REX-INCEP [35]	0.908	0.884	0.750	0.702	0.739	0.514	0.907	0.917	0.739
	CNN-ER [35]	**0.913**	**0.888**	**0.757**	**0.706**	**0.748**	**0.519**	**0.910**	**0.922**	**0.745**
	MEBeauty [6]	0.826	0.796	0.643	0.649	0.677	0.467	0.799	0.802	0.599
	FPEM (Ours)	**0.925**	**0.892**	**0.773**	**0.787**	**0.811**	**0.595**	**0.931**	**0.939**	**0.778**
	Improvement	+1.2%	+0.4%	+1.6%	+8.1%	+6.3%	+7.6%	+2.1%	+1.7%	+3.3%

5. Experiments

5.1. Benchmark Datasets

The proposed LiveBeauty dataset and two open-source FAP datasets including SCUT-FBP5500 [5] and MEBeauty [6] are selected as benchmark datasets. The MOS distributions of them are shown in Fig. 6. We split SCUT-FBP5500 and MEBeauty with 60%-40% and 80%-20% train-test ratio separately, which is kept same with their own protocol, and LiveBeauty is split using the 90%-10% protocol.

5.2. Implementation Details

Model initialization. In MAEM, ViT-B/16 and GPT-2 are adopted as the image encoder and the text encoder separately, which are initialized by [40]. In PAPM, Swin-T is adopted as the trainable image encoder, where only four stages are reserved and initialized with parameters pre-trained on multiple face-related tasks [38].

Hyper parameters. During the whole training process, AdamW is selected as the optimizer, and the learning rate scheduler is set with linear warm-up and cosine annealing scheme. The base learning rate varies with different training phase. Each training phase lasts 50 epochs, and the batch size is set as 32. Details can be found in the supplementary.

5.3. Competitors & Metrics

To confirm the effectiveness of our proposed method, we select baseline models and several SOTA FAP and IAA methods, including ComboNet [31], 2D-FAP [32], REX-INCEP [35], CNN-ER [35], MEBeauty [6], AVA-MLSP [46], TANet [47], Dele-Trans [80] and EAT [81]. All the models are retrained ourselves with the default parameters defined in their papers. Three criteria are adopted to evaluate the performance of these methods, *i.e.*, Spearman Rank Order Correlation Coefficient (SROCC), Pearson Linear Correlation Coefficient (PLCC) and Kendall Rank Order Correlation Coefficient (KROCC).

Table 3. Ablation study on LiveBeauty test set. FACA denotes Face-Aware Cross-Attention, MS denotes Multi-Scale features, SR denotes Similarity Regression strategy, $\mathcal{L}_S$ denotes the L1 scoring loss, $\mathcal{L}_R$ denotes the total ranking loss $\mathcal{L}_R = \mathcal{L}_{R_1} + \mathcal{L}_{R_2}$.

Modules	Methods	Metrics		
		SROCC↑	PLCC↑	KROCC↑
PAPM	only Swin-T	0.9045	0.8710	0.7454
	w/o FACA	0.9078	0.8714	0.7475
	w/o MS	0.9089	0.8808	0.7487
	PAPM	0.9096	0.8810	0.7497
MAEM	zero-shot CLIP	0.3838	0.3582	0.2692
	w/o SR	0.8825	0.8226	0.7123
	w/o $\mathcal{L}_S$	0.9161	0.8433	0.7607
	w/o $\mathcal{L}_R$	0.9186	0.8911	0.7653
	MAEM	0.9198	0.8631	0.7653
ALL	w/o 2-stage	0.9187	**0.8925**	0.7624
	w/o DFM	0.9214	0.8833	0.7685
	w/o CMFM	0.9237	0.8834	0.7718
	all combined	**0.9247**	0.8916	**0.7734**

Table 4. Ablation study on four predicted attractiveness scores before and after DFM.

SROCC	$\widehat{Q}_1$	$\widehat{Q}_2$	$\widehat{Q}_3$	$\widehat{Q}$
LiveBeauty	0.9096	0.9198	0.9210	**0.9247**
MEBeauty	0.7390	0.7772	0.7797	**0.7868**
SCUT-FBP5500	0.9116	0.9203	0.9236	**0.9314**

5.4. Performance

The experimental results on three FAP datasets are shown in Tab. 2, from which we could draw the following conclusions: (i) Our proposed method achieves the first place

and surpasses the second place by about 0.012, 0.081, 0.021 in terms of SROCC values on LiveBeauty, MEBeauty and SCUT-FBP5500 respectively, which demonstrates the superiority of our proposed method. (ii) The IAA methods are inferior to the FAP methods, which manifests that the generic aesthetic assessment methods overlook the facial features involved in the subjective nature of facial attractiveness, leading to poor performance on FAP tasks. (iii) The performance of all methods drops significantly on MEBeauty. This is because the training samples are limited and the faces are ethnically diverse in MEBeauty, indicating that there is a large diversity in facial attractiveness. All these factors make the prediction of facial attractiveness in MEBeauty more challenging.

5.5. Ablation Study

We conduct elaborate ablation study to investigate the contributions of different modules in our proposed method, which is shown in Tab. 3. Conclusions can be drawn that: (i) In the ablation study of PAPM, we can observe that leveraging facial prior and multi-scale features are both helpful. (ii) The result of ablation study of MAEM indicates that the zero-shot CLIP has limited generalization ability in FAP tasks, which confirms its shortcoming in aesthetic evaluation [82]. Specifically, for MAEM w/o SR, we leverage concatenation of the visual features and the textual features with highest similarity, and utilize MLP for regression. We observe a significant performance drop without SR, which manifests that specifically devised regression strategy is important for CLIP to adapt to regression tasks. Meanwhile, the joint optimization of ranking and regression also leads to improvement of SROCC. (iii) In the ablation study of ALL modules, we can see that employing all modules yields the best performance, which confirms the effectiveness of hybridly incorporating multi-modal aesthetic semantic features with holistic facial prior knowledge via CMFM.

Moreover, the comparison experiments with four predicted attractiveness scores before and after DFM are conducted as well, whose SROCC results are shown in Tab. 4. It can be observed that the SROCC of $\widehat{Q}_2$ is noticeably higher than that of $\widehat{Q}_1$, attributed to the complementary information provided by two modalities. The performance of $\widehat{Q}_3$ shows further improvement compared to $\widehat{Q}_2$, confirming that the personalized facial prior integrated in text embeddings for modeling contextual and individual facial feature helps improve the performance. The performance of the ensembled $\widehat{Q}$ outperforms the other individual predictions, which verifies the effectiveness of the DFM.

5.6. Cross-Dataset Study

Considering the diversity of facial attractiveness perception, the cross-dataset validation is carried out to validate the generalization ability of these methods. Specifically, all the mentioned models are trained on the proposed LiveBeauty dataset and subsequently tested on the entire SCUT-FBP5500 and the entire MEBeauty.

Table 5. Cross-Dataset study trained on LiveBeauty and tested on the entire MEBeauty dataset and the entire SCUT-FBP5500 dataset. Best in **red** and second-best in **blue**.

Methods	MEBeauty [6]			SCUT-FBP5500 [5]		
	SROCC↑	PLCC↑	KROCC↑	SROCC↑	PLCC↑	KROCC↑
ViT-B [77]	0.573	0.583	0.402	**0.553**	**0.548**	**0.394**
ResNeXt-50 [78]	0.535	0.546	0.372	0.478	0.467	0.335
Inception-V3 [79]	0.575	0.581	0.404	0.484	0.480	0.339
AVA-MLSP [46]	0.587	0.597	0.414	0.400	0.412	0.275
TANet [47]	0.551	0.529	0.384	0.401	0.411	0.277
Dele-Trans[80]	0.576	0.572	0.402	0.480	0.476	0.335
EAT[81]	**0.603**	**0.624**	**0.427**	0.484	0.466	0.338
ComboNet [31]	0.540	0.548	0.375	0.384	0.375	0.264
2D-FAP [32]	0.556	0.489	0.390	0.444	0.445	0.306
REX-INCEP [35]	0.588	0.592	0.413	0.473	0.471	0.332
CNN-ER [35]	0.589	0.596	0.414	0.502	0.499	0.352
MEBeauty [6]	0.520	0.516	0.361	0.366	0.339	0.251
FPEM (Ours)	**0.640**	**0.647**	**0.457**	**0.615**	**0.579**	**0.447**
Improvement	+3.7%	+2.3%	+3.0%	+6.2%	+3.1%	+5.3%

As shown in Tab. 5, our proposed method outperforms all other methods in both two datasets indicating that inserting personalized facial prior containing rich contextual and individual information is not only effective for improving performance but also for enhancing the generalizability. Compared to tested results on LiveBeauty, a significant performance drop can be observed, due to the domain gap between LiveBeauty and the other two datasets, *i.e.*, data from LiveBeauty are obtained from live streaming scenarios. In addition, LiveBeauty consisting of Asian face images, while SCUT-FBP5500 and MEBeauty contains multi-ethnic face images, which are more comprehensive. These unignorable gaps between data distribution lead to the performance drop on cross-dataset experiments.

6. Conclusion

In this paper, we tackle the problem of facial attractiveness prediction in live streaming. To address the existing gap in datasets, we introduce LiveBeauty, the first large-scale FAP dataset for live streaming. This dataset is high-quality and authentic, collected from real-world environments, and is currently the largest open-access facial attractiveness dataset. Additionally, we propose a multi-modal facial attractiveness prediction method, termed FPEM. This method captures holistic facial prior knowledge and multi-modal aesthetic semantic features, integrating them through a cross-modal fusion module. Extensive experiments indicate the superiority of our proposed method over other advanced methods. Meanwhile, the LiveBeauty and FPEM can be effectively leveraged for facial attractiveness prediction in other video applications.

References

[1] Shaolin Su, Hanhe Lin, Vlad Hosu, Oliver Wiedemann, Jinqiu Sun, Yu Zhu, Hantao Liu, Yanning Zhang, and Dietmar Saupe. Going the extra mile in face image quality assessment: A novel database and model. *IEEE Transactions on Multimedia*, 2023. 1, 3

[2] Nicolas Chahine, Stefania Calarasanu, Davide Garcia-Civiero, Theo Cayla, Sira Ferradans, and Jean Ponce. An image quality assessment dataset for portraits. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9968–9978, 2023. 1, 3

[3] Wei Sun, Weixia Zhang, Yanwei Jiang, Haoning Wu, Zicheng Zhang, Jun Jia, Yingjie Zhou, Zhongpeng Ji, Xiongkuo Min, Weisi Lin, and Zhai Guangtao. Dual-branch network for portrait image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2024. 1, 3

[4] Wei-Ting Chen, Gurunandan Krishnan, Qiang Gao, Sy-Yen Kuo, Sizhou Ma, and Jian Wang. Dsl-fiqa: Assessing facial image quality via dual-set degradation learning and landmark-guided transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2931–2941, 2024. 1, 3

[5] Lingyu Liang, Luojun Lin, Lianwen Jin, Duorui Xie, and Mengru Li. Scut-fbp5500: A diverse benchmark dataset for multi-paradigm facial beauty prediction. In *Proceedings of the IEEE International Conference on Pattern Recognition*, pages 1598–1603, 2018. 1, 2, 3, 7, 8

[6] Irina Lebedeva, Yi Guo, and Fangli Ying. Mebeauty: a multi-ethnic facial beauty dataset in-the-wild. *Neural Computing and Applications*, pages 1–15, 2022. 1, 2, 7, 8

[7] P Jonathon Phillips, J Ross Beveridge, David S Bolme, Bruce A Draper, Geof H Givens, Yui Man Lui, Su Cheng, Mohammad Nayeem Teli, and Hao Zhang. On the existence of face quality measures. In *Proceedings of the IEEE International Conference on Biometrics: Theory, Applications and Systems*, pages 1–8, 2013. 1

[8] Lacey Best-Rowden and Anil K Jain. Learning face image quality from human assessments. *IEEE Transactions on Information Forensics and Security*, 13(12):3064–3077, 2018.

[9] Fu-Zhao Ou, Xingyu Chen, Ruixin Zhang, Yuge Huang, Shaoxin Li, Jilin Li, Yong Li, Liujuan Cao, and Yuan-Gen Wang. Sdd-fiqa: Unsupervised face image quality assessment with similarity distribution distance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7670–7679, 2021. 1, 3

[10] Yu Tian, Zhangkai Ni, Baoliang Chen, Shiqi Wang, Hanli Wang, and Sam Kwong. Generalized visual quality assessment of gan-generated face images. *arXiv preprint arXiv:2201.11975*, 2022. 1

[11] Yael Eisenthal, Gideon Dror, and Eytan Ruppin. Facial attractiveness: Beauty and the machine. *Neural Computation*, 18(1):119–142, 2006. 1, 2, 3

[12] Luojun Lin, Lingyu Liang, Lianwen Jin, and Weijie Chen. Attribute-aware convolutional neural networks for facial beauty prediction. In *Proceedings of the Information Joint Conference on Artificial Intelligence*, pages 847–853, 2019. 1, 2, 3

[13] David Zhang, Fangmei Chen, Yong Xu, et al. A new hypothesis on facial beauty perception. In *Computer Models for Facial Beauty Analysis*, pages 143–163, 2016. 1

[14] Xiaoyu Ren, Huiyu Duan, Xiongkuo Min, Yucheng Zhu, Wei Shen, Linlin Wang, Fangyu Shi, Lei Fan, Xiaokang Yang, and Guangtao Zhai. Where are the children with autism looking in reality? In *CAAI International Conference on Artificial Intelligence*, pages 588–600, 2022. 1

[15] Naila Murray, Luca Marchesotti, and Florent Perronnin. Ava: A large-scale database for aesthetic visual analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2408–2415, 2012. 1, 2

[16] Luoqi Liu, Junliang Xing, Si Liu, Hui Xu, Xi Zhou, and Shuicheng Yan. Wow! you are so beautiful today! *ACM Transactions on Multimedia Computing, Communications, and Applications*, 11(1s):1–22, 2014. 1

[17] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Some like it hot-visual guidance for preference prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5553–5561, 2016. 1

[18] Luojun Lin, Zhifeng Shen, Jia-Li Yin, Qipeng Liu, Yuanlong Yu, and Weijie Chen. Metafbp: Learning to learn high-order predictor for personalized facial beauty prediction. In *Proceedings of the ACM International Conference on Multimedia*, pages 6072–6080, 2023. 1

[19] Wang Chen, Peizhen Chen, Weijie Chen, and Luojun Lin. Customized automatic face beautification. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–5, 2023. 1

[20] Douglas Gray, Kai Yu, Wei Xu, and Yihong Gong. Predicting facial beauty without landmarks. In *Proceedings of the European Conference on Computer Vision*, pages 434–447, 2010. 1, 2

[21] Haibin Yan. Cost-sensitive ordinal regression for fully automatic facial beauty assessment. *Neurocomputing*, 129:334–342, 2014. 1, 2

[22] Harm De Vries and Jason Yosinski. Can deep learning help you find the perfect match? *arXiv preprint arXiv:1505.00359*, 2015. 2

[23] Yikui Zhai, Yu Huang, Ying Xu, Junying Zeng, Fei Yu, and Junying Gan. Benchmark of a large scale database for facial beauty prediction. In *Proceedings of the International Conference on Intelligent Information Processing*, pages 1–5, 2016. 2

[24] Song Tong, Xuefeng Liang, Takatsune Kumada, and Sunao Iwaki. Putative ratios of facial attractiveness in a deep neural network. *Vision Research*, 178:86–99, 2021. 2

[25] Yikui Zhai, Yu Huang, Ying Xu, Junying Gan, He Cao, Wenbo Deng, Ruggero Donida Labati, Vincenzo Piuri, and Fabio Scotti. Asian female facial beauty prediction using deep neural networks via transfer learning and multi-channel feature fusion. *IEEE Access*, 8:56892–56907, 2020. 1, 2

[26] Tharun J. Iyer, Rahul K, Ruban Nersisson, Zhemin Zhuang, Alex Noel Joseph Raj, and Imthiaz Refayee. Machine

learning-based facial beauty prediction and analysis of frontal facial images using facial landmarks and traditional image descriptors. *Computational Intelligence and Neuroscience*, 2021(1):4423407, 2021. 2, 3

[27] Shu Liu, Yang-Yu Fan, Zhe Guo, Ashok Samal, and Afan Ali. A landmark-based data-driven approach on 2.5 d facial attractiveness computation. *Neurocomputing*, 238:168–178, 2017. 2, 3

[28] Jie Xu, Lianwen Jin, Lingyu Liang, Ziyong Feng, Duorui Xie, and Huiyun Mao. Facial attractiveness prediction using psychologically inspired convolutional neural network (pi-cnn). In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1657–1661, 2017. 2, 3

[29] Yang-Yu Fan, Shu Liu, Bo Li, Zhe Guo, Ashok Samal, Jun Wan, and Stan Z Li. Label distribution-based facial attractiveness computation by deep residual learning. *IEEE Transactions on Multimedia*, 20(8):2196–2208, 2017. 3

[30] Shengjie Shi, Fei Gao, Xuantong Meng, Xingxin Xu, and Jingjie Zhu. Improving facial attractiveness prediction via co-attention learning. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4045–4049, 2019. 3

[31] Lu Xu and Jinhai Xiang. Comboloss for facial attractiveness analysis with squeeze-and-excitation networks. *arXiv preprint arXiv:2010.10721*, 2020. 2, 7, 8

[32] Shu Liu, Enquan Huang, Yan Xu, Kexuan Wang, Xiaoyan Kui, Tao Lei, and Hongying Meng. Lightweight facial attractiveness prediction using dual label distribution. *arXiv preprint arXiv:2212.01742*, 2022. 7, 8

[33] Qipeng Liu, Luojun Lin, Zhifeng Shen, and Yuanlong Yu. Fbpformer: Dynamic convolutional transformer for global-local-contexual facial beauty prediction. In *Proceedings of the International Conference on Artificial Neural Networks*, pages 223–235, 2023. 2, 3

[34] Tianhao Peng, Mu Li, Fangmei Chen, Yong Xu, and David Zhang. Geometric prior guided hybrid deep neural network for facial beauty analysis. *CAAI Transactions on Intelligence Technology*, 9(2):467–480, 2024. 2

[35] Fares Bougourzi, Fadi Dornaika, and Abdelmalik Taleb-Ahmed. Deep learning based face beauty prediction via dynamic robust losses and ensemble regression. *Knowledge-Based Systems*, 242:108246, 2022. 2, 3, 7, 8

[36] B Series. Methodology for the subjective assessment of the quality of television pictures. *Recommendation ITU-R BT*, 500(13), 2012. 2, 4

[37] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 2

[38] Lixiong Qin, Mei Wang, Chao Deng, Ke Wang, Xi Chen, Jiani Hu, and Weihong Deng. Swinface: a multi-task transformer for face recognition, expression recognition, age estimation and attribute estimation. *IEEE Transactions on Circuits and Systems for Video Technology*, 2023. 2, 7

[39] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015. 2

[40] 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 *Proceedings of the International Conference on Machine Learning*, pages 8748–8763, 2021. 2, 4, 7

[41] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z Wang. Studying aesthetics in photographic images using a computational approach. In *Proceedings of the European Conference on Computer Vision*, pages 288–301, 2006. 2

[42] Yan Ke, Xiaoou Tang, and Feng Jing. The design of high-level features for photo quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 419–426, 2006. 2

[43] Luca Marchesotti, Florent Perronnin, Diane Larlus, and Gabriela Csurka. Assessing the aesthetic quality of photographs using generic image descriptors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1784–1791, 2011. 2

[44] Congcong Li, Andrew Gallagher, Alexander C. Loui, and Tsuhan Chen. Aesthetic quality assessment of consumer photos with faces. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3221–3224, 2010. 2

[45] Yubin Deng, Chen Change Loy, and Xiaoou Tang. Image aesthetic assessment: An experimental survey. *IEEE Signal Processing Magazine*, 34(4):80–106, 2017. 2

[46] Vlad Hosu, Bastian Goldlucke, and Dietmar Saupe. Effective aesthetics prediction with multi-level spatially pooled features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9375–9383, 2019. 2, 7, 8

[47] Shuai He, Yongchang Zhang, Rui Xie, Dongxiang Jiang, and Anlong Ming. Rethinking image aesthetics assessment: Models, datasets and benchmarks. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 942–948, 2022. 7, 8

[48] Shu Kong, Xiaohui Shen, Zhe Lin, Radomir Mech, and Charless Fowlkes. Photo aesthetics ranking network with attributes and content adaptation. In *Proceedings of the European Conference on Computer Vision*, pages 662–679, 2016.

[49] Hui Zeng, Zisheng Cao, Lei Zhang, and Alan C Bovik. A unified probabilistic formulation of image aesthetic assessment. *IEEE Transactions on Image Processing*, 29:1548–1561, 2019. 2

[50] Ayman Abaza, Mary Ann Harrison, and Thirimachos Bourlai. Quality metrics for practical face recognition. In *Proceedings of the International Conference on Pattern Recognition*, pages 3103–3107, 2012. 3

[51] Ayman Abaza, Mary Ann Harrison, Thirimachos Bourlai, and Arun Ross. Design and evaluation of photometric image quality measures for effective face recognition. *IET Biometrics*, 3(4):314–324, 2014.

[52] Abhishek Dutta, Raymond Veldhuis, and Luuk Spreeuwers. A bayesian model for predicting face recognition performance using image quality. In *Proceeding of the International Joint Conference on Biometrics*, pages 1–8, 2014.

[53] Philipp Terhorst, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Ser-fiq: Unsupervised estimation of face image quality based on stochastic embedding robustness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5651–5660, 2020. 3

[54] Omkar Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *Proceedings of the British Machine Vision Conference*, page 6, 2015.

[55] Fadi Boutros, Meiling Fang, Marcel Klemt, Biying Fu, and Naser Damer. Cr-fiqa: face image quality assessment by learning sample relative classifiability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5836–5845, 2023.

[56] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14225–14234, 2021.

[57] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18750–18759, 2022.

[58] Fu-Zhao Ou, Chongyi Li, Shiqi Wang, and Sam Kwong. Clib-fiqa: Face image quality assessment with confidence calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1694–1704, 2024. 3

[59] Tie Liu, Shengxi Li, Mai Xu, Li Yang, and Xiaofei Wang. Assessing face image quality: A large-scale database and a transformer method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 3

[60] Nicolas Chahine, Sira Ferradans, Javier Vazquez-Corral, and Jean Ponce. Generalized portrait quality assessment. *arXiv preprint arXiv:2402.09178*, 2024. 3

[61] Kerang Cao, Kwang-nam Choi, Hoekyung Jung, and Lini Duan. Deep learning for facial beauty prediction. *Information*, 11(8):391, 2020. 3

[62] Shu Liu, Bo Li, Yang-Yu Fan, Zhe Quo, and Ashok Samal. Facial attractiveness computation by label distribution learning with deep cnn and geometric features. In *Proceedings of the IEEE International Conference on Multimedia and Expo*, pages 1344–1349, 2017. 3

[63] Shu Liu, Enquan Huang, Yan Xu, Kexuan Wang, Xiaoyan Kui, Tao Lei, and Hongying Meng. Lightweight facial attractiveness prediction using dual label distribution. *arXiv preprint arXiv:2212.01742*, 2022. 3

[64] Lian Gao, Weixin Li, Zehua Huang, Di Huang, and Yunhong Wang. Automatic facial attractiveness prediction by deep multi-task learning. In *Proceedings of the IEEE International Conference on Pattern Recognition*, pages 3592–3597, 2018. 3

[65] Qinjie Xiao, You Wu, Dinghong Wang, Yong-Liang Yang, and Xiaogang Jin. Beauty3dfacenet: deep geometry and texture fusion for 3d facial attractiveness prediction. *Computers & Graphics*, 98:11–18, 2021. 3

[66] Lu Xu, Heng Fan, and Jinhai Xiang. Hierarchical multi-task network for race, gender and facial attractiveness recognition. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3861–3865, 2019. 3

[67] Lu Xu and Jinhai Xiang. Comboloss for facial attractiveness analysis with squeeze-and-excitation networks. *arXiv preprint arXiv:2010.10721*, 2020. 3

[68] Junying Gan, Xiaoshan Xie, Guohui He, and Heng Luo. Transbls: transformer combined with broad learning system for facial beauty prediction. *Applied Intelligence*, 53(21): 26110–26125, 2023. 3

[69] Djamel Eddine Boukhari, Ali Chemsa, and Riadh Ajgou. Facial beauty prediction based on vision transformer. *International Journal of Electrical and Electronic Engineering & Telecommunications*, pages 2319–2518, 2023. 3

[70] Junying Gan, Xiaoshan Xie, Yikui Zhai, Guohui He, Chaoyun Mai, and Heng Luo. Facial beauty prediction fusing transfer learning and broad learning system. *Soft Computing*, 27(18):13391–13404, 2023. 3

[71] Shifeng Zhang, Xiangyu Zhu, Zhen Lei, Hailin Shi, Xiaobo Wang, and Stan Z Li. Faceboxes: A cpu real-time face detector with high accuracy. In *Proceedings of the IEEE International Joint Conference on Biometrics*, pages 1–9, 2017. 3

[72] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3d dense face alignment. In *Proceedings of the European Conference on Computer Vision*, 2020. 3

[73] Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14071–14081, 2023. 5

[74] Louis L Thurstone. A law of comparative judgment. In *Scaling*, pages 81–92. Routledge, 2017. 6

[75] Ming-Feng Tsai, Tie-Yan Liu, Tao Qin, Hsin-Hsi Chen, and Wei-Ying Ma. Frank: a ranking method with fidelity loss. In *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 383–390, 2007. 6

[76] Qinkai Yu, Jianyang Xie, Anh Nguyen, He Zhao, Jiong Zhang, Huazhu Fu, Yitian Zhao, Yalin Zheng, and Yanda Meng. Clip-dr: Textual knowledge-guided diabetic retinopathy grading with ranking-aware prompting. *arXiv preprint arXiv:2407.04068*, 2024. 6

[77] Alexander Kolesnikov, Alexey Dosovitskiy, Dirk Weissenborn, Georg Heigold, Jakob Uszkoreit, Lucas Beyer, Matthias Minderer, Mostafa Dehghani, Neil Houlsby, Sylvain Gelly, Thomas Unterthiner, and Xiaohua Zhai. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2021. 7, 8

[78] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1492–1500, 2017. 7, 8

[79] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016. 7, 8

[80] Shuai He, Anlong Ming, Yaqi Li, Jinyuan Sun, ShunTian Zheng, and Huadong Ma. Thinking image color aesthetics assessment: Models, datasets and benchmarks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21838–21847, 2023. 7, 8

[81] Shuai He, Anlong Ming, Shuntian Zheng, Haobin Zhong, and Huadong Ma. Eat: An enhancer for aesthetics-oriented transformers. In *Proceedings of the ACM International Conference on Multimedia*, pages 1023–1032, 2023. 7, 8

[82] Xiangfei Sheng, Leida Li, Pengfei Chen, Jinjian Wu, Weisheng Dong, Yuzhe Yang, Liwu Xu, Yaqian Li, and Guangming Shi. Aesclip: Multi-attribute contrastive learning for image aesthetics assessment. In *Proceedings of the ACM International Conference on Multimedia*, page 1117–1126, 2023. 8