

LibreFace: An Open-Source Toolkit for Deep Facial Expression Analysis

Di Chang, Yufeng Yin, Zongjian Li, Minh Tran, Mohammad Soleymani
 Institute for Creative Technologies, University of Southern California

{dichang, yufengy, minhnttra}@usc.edu, {lizongjian, soleymani}@ict.usc.edu

Abstract

Facial expression analysis is an important tool for human-computer interaction. In this paper, we introduce LibreFace, an open-source toolkit for facial expression analysis. This open-source toolbox offers real-time and offline analysis of facial behavior through deep learning models, including facial action unit (AU) detection, AU intensity estimation, and facial expression recognition. To accomplish this, we employ several techniques, including the utilization of a large-scale pre-trained network, feature-wise knowledge distillation, and task-specific fine-tuning. These approaches are designed to effectively and accurately analyze facial expressions by leveraging visual information, thereby facilitating the implementation of real-time interactive applications. In terms of Action Unit (AU) intensity estimation, we achieve a Pearson Correlation Coefficient (PCC) of **0.63** on DISFA, which is **7%** higher than the performance of OpenFace 2.0 [4] while maintaining highly-efficient inference that runs **two times** faster than OpenFace 2.0 [4]. Despite being compact, our model also demonstrates competitive performance to state-of-the-art facial expression analysis methods on AffecNet, FFHQ, and RAF-DB. Our code will be released at <https://github.com/ihp-lab/LibreFace>

1. Introduction

Facial expression analysis is the process of automatic detection of subtle facial muscle movements and recognition of prototypical facial displays. Recognizing facial expressions can provide valuable information about users' social and affective states, with a wide range of applications in human-computer interaction (HCI) to mental health [10, 14], or human-agent negotiation [8].

Two widely used methods for facial expression analysis are facial Action Unit (AU) intensity estimation [12] and facial expression recognition. A facial action unit is an indicator of activation of an individual or a group of muscles, e.g., cheek raiser (AU6). AUs are formalized by Paul Ekman in Facial Action Coding System (FACS) [12]. Accurate and

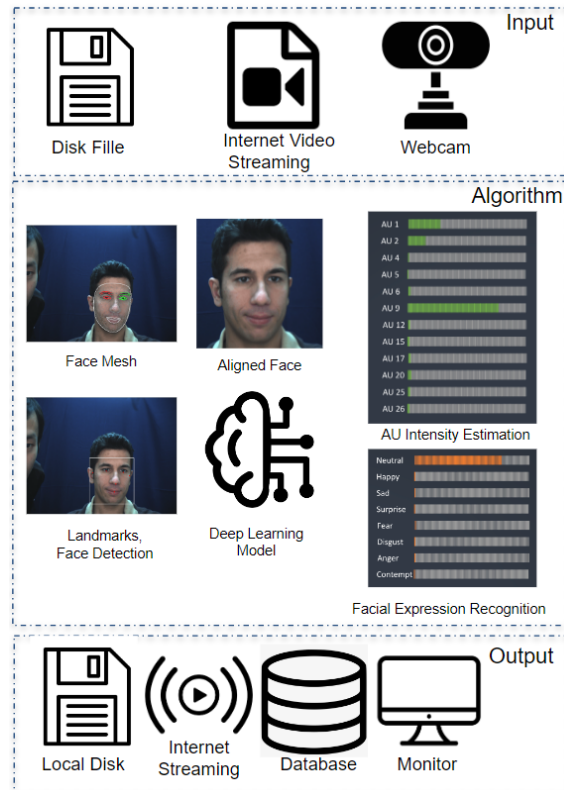


Figure 1. LibreFace is an open-source, accurate, and efficient framework for facial expression analysis, including: face mesh detection, landmark detection, AU intensity estimation, and facial expression recognition. LibreFace supports CPU-only environment as well as GPU acceleration when necessary.

efficient AU intensity estimation is crucial for recognizing complex facial expressions. Furthermore, facial expression recognition (FER) refers to identifying the prototypical facial displays or expressions, e.g., expressions of sadness. This is a challenging problem due to the high complexity and variability of facial expressions.

Over the past few years, previous studies on facial expression analysis have achieved promising performance with deep learning-based methods [23, 37, 38]. However, recent facial expression tools (see Table 1), e.g. OpenFace [3] and OpenFace 2.0 [4] are developed with traditional ma-




Figure 2. Overall pipeline of our proposed LibreFace, including face mesh and landmark detection, image alignment, action unit recognition, and facial expression recognition. All outputs from the pipeline can be saved, and training/inference code is open-sourced.

chine learning methods such as Support Vector Machine (SVM), which are less accurate than the recent deep models. On the other hand, despite achieving superior performance, state-of-the-art deep learning-based models are large networks with time-consuming training and inference requirements, which significantly prevent the deployment of such methods in real-world scenarios.

To address the aforementioned problems, we introduce LibreFace, an open-source, accurate, and efficient framework for facial expression analysis (see Figure 1). Specifically, we develop AU intensity estimation and facial expression recognition models with the SOTA pre-trained vision encoders, *e.g.*, ResNet [17], Swin-Transformer [44], and Masked Auto-encoder (MAE) [16]. To boost the performance, we further pre-train the vision encoders on various face datasets, including EmotioNet [13], AffectNet [24], and FFHQ [19]. Furthermore, to achieve real-time facial expression analysis, we utilize the feature-wise knowledge distillation based on [45, 49] to improve the inference efficiency. We deliver an executable cross-platform software that supports both CPU-only and GPU-accelerated solutions according to users’ hardware.

Extensive experiments show that LibreFace provides a more accurate, comprehensive, and efficient alternative to OpenFace 2.0 [4], the most commonly used facial expression analysis toolkit, and other open facial behavior analysis software. For AU intensity estimation, LibreFace achieves superior performance than OpenFace 2.0 [4] while running two times faster on a CPU-only environment. For FER, LibreFace can achieve competitive performance to the heavy state-of-the-art methods.

Our major contributions are summarized as follows.

- We present LibreFace, an open-source and comprehensive toolkit for accurate and real-time facial expression analysis with both CPU-only and GPU-acceleration versions. LibreFace eliminates the gap between cutting-edge research and an easy and free-to-use non-commercial toolbox.
- We propose to adaptively pre-train the vision encoders

with various face datasets and then distillate them to a lightweight ResNet-18 model in a feature-wise matching manner.

- We conduct extensive experiments of pre-training and distillation to demonstrate that our proposed pipeline achieves comparable results to state-of-the-art works while maintaining real-time efficiency.
- LibreFace system supports cross-platform running, and the code is open-sourced in C# (model inference and checkpoints) and Python (model training, inference, and checkpoints).

2. Related Work

2.1. Pre-trained Vision Encoders

Previous studies have proposed various pre-trained vision encoders to extract strong representations for computer vision tasks, including CNN-based models, *e.g.*, Residual Network (ResNet) [17], ResNeXt [43], and transformer-based approaches, *e.g.*, Vision Transformer (ViT) [11], Swin Transformer [27] and Masked Auto-encoder (MAE) [16].

He *et al.* propose Residual Network (ResNet) [17], a CNN-based model, to solve the gradient vanishing problem. He *et al.* introduce the residual block in which there is a direct connection skipping some layers in between, *i.e.*, skip connection. Swin Transformer [27] is a hierarchical vision transformer and can serve as a general-purpose backbone for computer vision. It has gained significant attention due to its strong performance in different recognition tasks. One of its key features is the use of shifted windows, which enables cross-window connections and leads to more efficient self-attention computation. Moreover, the hierarchical structure allows it to handle facial features of various scales. Another noteworthy work is the self-supervised Masked Auto-encoder (MAE) [16]. To eliminate the limitation of lacking labeled training data, the training process is conducted through an image reconstruction task. Subsets of image patches are masked randomly, and the remaining patches are passed to the encoder. A small decoder then processes the encoded patches and masked tokens to reconstruct the original image. After pre-training, the decoder is discarded, and the encoder can be applied to downstream classification tasks. In this paper, we extract the features with ResNet, Swin Transformer, and MAE for facial expression analysis. All of them are pre-trained on ImageNet [9], which is not a face dataset.

Recently, Cai *et al.* propose MARLIN [7], an extension of MAE for facial video representation learning. MARLIN is pre-trained on the YouTube Face dataset [41] with video reconstruction, and it learns universal features which

Table 1. Comparison of facial behavior analysis tools. *AU* and *FER* stand for AU intensity estimation and facial expression recognition. Note that the checkmark for *Test* requires releasing of both checkpoints and codes for inference. *Executable* denotes if the corresponding toolbox releases a binary executable program. *: Toolbox without \checkmark does not support GPU for acceleration purposes, but can still be running on CPU-only environment.

Toolbox	Approach	Landmark	GPU support*	AU	FER	Train	Test	Executable	Real-time	Free
dlib [22]	[21]	\checkmark				\checkmark	\checkmark		\checkmark	\checkmark
FaceTracker	[35]	\checkmark					\checkmark	\checkmark	\checkmark	\checkmark
Mediapipe [29]	[20]	\checkmark	\checkmark		\checkmark		\checkmark	\checkmark	\checkmark	\checkmark
AFFDEX	unknown	\checkmark		\checkmark	\checkmark			\checkmark	\checkmark	
AFFDEX 2.0 [5]	CNN, [34]	\checkmark		\checkmark	\checkmark			\checkmark	\checkmark	
FACET	unknown	\checkmark		\checkmark				\checkmark	\checkmark	
OpenFace [3]	[1,2]	\checkmark		\checkmark		\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
OpenFace 2.0 [4]	[42,47,48]	\checkmark		\checkmark		\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
LibreFace (Ours)	[16,17,45]	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark

are transferable across various facial analysis tasks. Thus, motivated by MARLIN, we argue that further pre-training the vision encoders on various face datasets, *e.g.*, EmotionNet [13], AffectNet [24], and FFHQ [19] can boost the performance for facial expression analysis.

2.2. Facial Expression Analysis Toolkit

Over the past few years, several open-source toolkits for facial behavior analysis have been developed with traditional machine learning methods, *e.g.*, histogram of gradients (HOG) and support vector machine (SVM). The most popular tools among these are OpenFace [3] and OpenFace 2.0 [4].

OpenFace [3] is an open-source framework for facial behavior analysis. The framework is capable of various facial analysis tasks such as landmark detection, gaze tracking, and AU detection. OpenFace 2.0 [4] is an extension work of OpenFace with several improvements over the original one. The capabilities of OpenFace 2.0 include facial landmark detection, head pose estimation, eye-gaze estimation, and facial action unit recognition.

With this work, we aim to provide a more accurate and efficient facial expression analysis framework based on deep neural networks, *e.g.*, ResNet [17], Swin Transformer [27] and MAE [16], which have recently achieved resounding success in other computer vision tasks. We also deploy the idea of traditional knowledge distillation [18] and utilize feature-wise distillation based on [45,49] to further improve the inference efficiency to achieve real-time facial expression analysis better than the previous works.

3. Problem Formulation

3.1. AU Intensity Estimation

AU intensity estimation refers to the process of quantifying the intensity or strength of specific facial muscle movements known as Action Units (AUs). It involves analyzing facial images or video frames to detect and track the presence and magnitude of specific AUs. The intensity estimation is usually represented on a scale of from 0 to 5, with the bigger value indicating stronger strength of the AU activation. The approaches for AU intensity estimation can be categorized into traditional feature-based methods and deep learning-based methods. Traditional feature-based methods involve extracting hand-crafted features from facial images, such as texture, shape, or motion information. These features are then used as input to machine learning algorithms, such as support vector machines (SVMs) or random forests, to estimate AU intensities. OpenFace [3] and OpenFace 2.0 [4] are representative works of feature-based methods.

Deep learning-based methods leverage the power of deep neural networks to learn features from facial images or video frames automatically. Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Transformers structures can be used for this task. Affdex 2.0 [5] is a representative work of this category. In LibreFace, we use Masked Auto-encoder (MAE) and ResNet-18 for AU intensity estimation. In Table 2, we provide the AUs that LibreFace can predict.

3.2. Facial Expression Recognition

Facial Expression Recognition (FER) is the process of categorizing or recognizing the prototypical facial displays or expressions. It involves analyzing facial features, such as the position and movement of facial muscles, to iden-

Table 2. Available action units in LibreFace. **R** - LibreFace provides intensity estimation, **D** - LibreFace provides detection of occurrence.

AU	FACS name	Prediction
AU1	Inner brow raiser	R
AU2	Outer brow raiser	R
AU4	Brow lowerer	R
AU5	Upper lid raiser	R
AU6	Cheek raiser	R
AU7	Lid tightener	D
AU9	Nose wrinkler	R
AU10	Upper lip raiser	D
AU12	Lip corner puller	R
AU14	Dimpler	D
AU15	Lip corner depressor	R
AU17	Chin raiser	R
AU20	Lip stretcher	R
AU23	Lip tightener	D
AU24	Lip pressor	D
AU25	Lips part	R
AU26	Jaw drop	R

tify and label specific emotions or expressions. It is worth noting that FER is a complex task due to factors like inter-individual variability, context dependency, and cultural differences in expressions. Researchers continue to explore advanced techniques [26, 46, 53], such as deep learning architectures (e.g., CNNs, RNNs and Transformers), multi-modal approaches (combining facial and audio/text cues), and transfer learning, to improve the accuracy and robustness of FER systems. In Figure 3, we visualize the software of the LibreFace system deployed on a Personal Computer and show the example of feeding a local image file as input.

4. Method

4.1. Overview

The overall framework is illustrated in Figure 2. We first discuss the image pre-processing, which involves face mesh and landmark detection and image alignment in Section 4.2. Next, we feed the pre-processed images into a pre-trained MAE encoder, followed by a linear regression or classification layer that predicts the AU intensity values or facial expression labels, as detailed in Section 4.3. Once the MAE is fine-tuned, we employ feature-wise distillation to transfer the teacher model’s (MAE) knowledge to a lightweight student model (ResNet-18), as outlined in Section 4.4. Finally, we use the distilled ResNet-18 for efficient AU intensity estimation and FER.

4.2. Face Mesh and Landmark Detection

Prior to inputting an arbitrary image into deep neural networks, it is imperative to perform facial image alignment based on localized landmarks. Facial image alignment in-

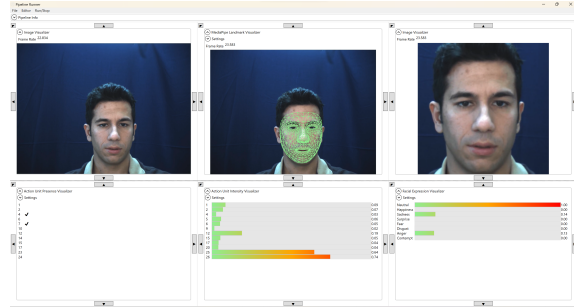


Figure 3. Overview of LibreFace software deployed in the Windows Operating System. We offer real-time facial action unit intensity estimation, detection, and facial expression recognition.


volves geometric transformations, such as translation, rotation, and scaling, to convert the input face image into a canonical or standardized form. This process ensures consistent positioning of facial features across various images, facilitating the learning of patterns by our model. In our methodology, we leverage an open-source tool, i.e., MediaPipe [29], to detect the face mesh within the input image. By utilizing the detected mesh, we determine the precise locations of the facial landmarks, including the eyes, eyebrows, nose, and mouth. Subsequently, these facial landmarks are utilized to crop and align the facial region, culminating in the production of the resulting aligned output.

4.3. Pre-training and Fine-tuning


Pre-training in facial expression analysis tasks refers to the process of training a deep learning model on a large dataset, typically using a supervised or self-supervised learning approach, before fine-tuning it on the specific tasks. The rationale behind pre-training is to leverage the general knowledge learned from a large-scale dataset to improve the performance and efficiency of the model on the target task, which allows the model to learn general visual representations, overcome data scarcity, facilitate transfer learning, improve feature extraction, and enhance the robustness and generalization capabilities.

In this section, we discuss the pre-training process for the MAE teacher model and the ResNet-18 student model. As illustrated in Figure 4, the ViT-base structural Masked Auto-encoder (MAE) is pre-trained on EmotioNet [13]. In this work, we use the model weights provided by [50]. Subsequently, we employ the weights of the ViT encoder as the backbone, add a linear classifier, and further pre-train the entire model on AffectNet [24] and FFHQ [19] datasets. After pre-training on these two datasets, we fine-tune the model on DISFA [31] for AU intensity estimation. We use a Mean Squared Error (MSE) loss as the task-specific loss for AU intensity estimation since it is a regression task, which can be expressed as

$$\mathcal{L}_{Task} = \| \hat{y} - y \|^2, \quad (1)$$



(a) The ViT-base backbone of the MAE is pre-trained on the EmotionNet [13] dataset with a self-supervised image reconstruction loss.



(b) We pre-train the whole pipeline with MAE or ResNet-18 encoder on the AffectNet [32] and FFHQ [19] dataset with cross-entropy loss.

Figure 4. Pre-training strategies of LibreFace for facial expression analysis.




Figure 5. Overview of feature-wise distillation. We train the student model encoder E_S with feature-wise distillation, which transfers the knowledge from the teacher model encoder E_T . C_T and C_S represent the classifier of the teacher and student model, respectively.

where \hat{y} and y represent the prediction and ground-truth for AU intensities.

Similarly, we employ ResNet-18 as the encoder and train it on AffectNet [24] and FFHQ [19] along with a linear classifier. Subsequently, we fine-tune the whole model on DISFA [31]. More details can be found in Section 5.4. We use a Cross-Entropy (CE) loss for the classification of FER. The task-specific loss is defined as

$$\mathcal{L}_{Task} = - \sum_{c=1}^M y \log(\hat{y}), \quad (2)$$

where M is the number of categories of facial expressions, while \hat{y} and y represent one-hot encoded prediction and ground-truth labels.

4.4. Feature-wise Distillation

Given that MAE is constructed on a ViT-base backbone, which demands significant computational resources, we aim to enhance the pipeline’s efficiency and speed during inference. Although ResNet-18 already demonstrates strong performance for AU intensity estimation, we can fur-

ther distill knowledge from the MAE into ResNet-18 to simultaneously boost performance and reduce computational cost. As illustrated in Figure 5, we propose to distill in a feature-matching manner, following the idea in [45]. During distillation, all the weights from the teacher model, including an encoder E_T and a classifier C_T , are frozen, and we fine-tune the student model encoder E_S and classifier C_S with a feature-matching loss \mathcal{L}_{FM} , a task-specific regression or classification loss \mathcal{L}_{Task} (Equations 1 and 2) and a Kullback–Leibler (KL) divergence loss L_{KL} .

We minimize the difference between the features from teacher encoder f_T and the features from student encoder f_S by a simple MSE loss

$$\mathcal{L}_{FM} = \| f_T - \mathbf{I}(f_S) \|^2, \quad (3)$$

where $\mathbf{I}(\cdot)$ denotes the linear interpolation function to match the feature dimension.

\mathcal{L}_{Task} is the same as any model trained in a non-distillation fashion (Equations 1 and 2), which corrects the wrong prediction from C_S with ground-truth labels.

In the original knowledge distillation approach, as proposed by Hinton *et al.* [18], the KL divergence loss minimizes the discrepancy between the teacher and student classifiers. In our work, however, we input the student features into the teacher classifier, utilizing the frozen, pre-trained teacher classifier for both the teacher and student models. This approach grants more flexibility to the optimization algorithm, allowing it to adjust the weights of both student encoder E_S and classifier C_S to minimize the loss more effectively.

$$L_{KL} = -\text{Softmax}(\hat{y}_T) \log \text{Softmax}(\hat{y}_S) \quad (4)$$

where \hat{y}_T and \hat{y}_S denote the predictions from C_T with teacher features and student features respectively.

Finally, we train the whole pipeline with the following overall loss

$$\mathcal{L} = \mathcal{L}_{FM} + \alpha \mathcal{L}_{Task} + \beta \mathcal{L}_{KL}, \quad (5)$$

where α and β are the loss weights.

5. Experiments

5.1. Dataset

EmotionNet [13], **FFHQ** [19] and **AffectNet** [32] are used to pre-train our models. EmotionNet dataset [13] is a large-scale collection of facial expressions in the wild, containing approximately 975,000 annotated images. The annotations include Action Units (AUs) and 421 emotion keywords, which are automatically assigned using the EmotionNet algorithm [13]. This dataset is suitable for a variety of AU applications, including recognition and intensity estimation, as well as basic and compound emotion recognition. Mixed results may occur when facial expressions are not well composed or ambiguous. For instance, some annotations may indicate a combination of multiple emotions or an uncertain degree of intensity, such as 68.66% disgust and 27.98% fear.

AffectNet [13] is a vast facial expression dataset comprising over 1,000,000 facial images. The images are gathered from the internet by searching for 1250 keywords related to emotions in six different languages. To support both discrete emotion classification and continuous affective computing, the dataset is manually labeled for eight different facial expressions and the intensity of valence and arousal. It is currently the largest database of facial expression, valence, and arousal, providing a valuable resource for research in affective computing. FFHQ [19] is a dataset of high-quality human facial images. Its initial purpose was to serve as a benchmark for the evaluation of generative adversarial networks (GANs). This dataset provides 70,000 images in the PNG format with a resolution of 1024×1024 and a substantial range of variation with respect to age, ethnicity, and image background. Furthermore, FFHQ boasts extensive coverage of facial accessories such as eyeglasses, sunglasses, and hats.

DISFA [31] is utilized for training and evaluation of AU intensity estimation. This dataset is a non-posed large-scale facial expression database captured by stereo videos of 27 adult subjects (12 females and 15 males) from different ethnicities. These videos are annotated with facial action unit intensity values ranging between [0, 5] manually labeled by two human FACS experts. Following the setting from OpenFace 2.0 [4], we train our proposed pipeline with five-fold cross-validation so that we fairly compare our results. We apply the same model we used for AU intensity estimation to **BP4D** [51] dataset to provide AU detection results for those AUs that are not provided with labels in DISFA [31].

For FER, we report the result of our proposed LibreFace on the test set of AffectNet [32] and **RAF-DB** [25]. We apply the same pre-training strategy and feature-wise distillation we used for AU intensity estimation to FER.

RAF-DB [25] is a large-scale facial expression dataset

consisting of approximately 30,000 diverse facial images which feature a wide range of variations in age, gender, ethnicity, head pose, and lighting conditions. This work uses the single-label subset, including seven classes of basic emotions, for evaluation purposes.

5.2. Implementation and Training Details

All methods are implemented in PyTorch [33]. Training code and model weights are available for the sake of reproducibility. All tasks are completed with one NVIDIA RTX 8000 GPU for training, and our pipeline can be tested with either a CPU-only machine or a machine with a GPU.

The input image is first resized to 256×256 . For data augmentation, each face is randomly cropped into 224×224 and randomly flipped in the horizontal direction. We train the model with the AdamW optimizer [28] with a batch size of 128. The initial learning rate is $3e-5$ with a weight decay of $1e-4$. We train the model for a maximum of 20 epochs with early stopping. The model is trained with MSE loss for AU intensity estimation and cross-entropy loss for FER while being evaluated on the validation set at the end of every epoch. For the loss weights in the feature-wise distillation, we set $\alpha = 1.0$ and $\beta = 1.0$ for all experiments after hyper-parameter tuning.

5.3. Software Implementation

The software is implemented in C# and consists of three major components, MediaPipe [29], image aligner, and ONNX models. MediaPipe is used for detecting faces and providing facial landmarks. The image aligner crops and aligns faces from input images according to detected facial landmarks. Aligned images will be provided as inputs for our models. Models run concurrently and can be enabled or disabled according to use cases. All components are developed as .NET libraries and are multi-platform. For easy deployment, we have also developed a graphical user interface and turned them into Microsoft’s Platform for Situated Intelligence components [6] that can be easily used in OpenSense [40].

5.4. Experimental Setup

Pre-training and Fine-tuning Strategy. For the final model deployed in the LibreFace pipeline, we apply knowledge distillation and pre-training to boost performance. We first pre-train a Masked Auto-encoder on EmotionNet, AffectNet, and FFHQ training set. The training strategy is shown in Figure 4. We then take the checkpoint of the MAE encoder and fine-tune it on DISFA [31] for AU intensity estimation. After we have finished the five-fold cross-validation for MAE, we adopt the MAE checkpoint (including the encoder and linear classifier) from these five folds as the teacher model and apply feature-wise knowledge distillation to train a ResNet-18 encoder and a linear classifier.

Table 3. Performance of our proposed LibreFace pipeline on DISFA [31]. Following [3, 4], we evaluate our pipeline with five-fold cross-validation. We use Pearson Correlation Coefficient (PCC) for comparison to other methods.

Method	AU1	AU2	AU4	AU5	AU6	AU9	AU12	AU15	AU17	AU20	AU25	AU26	Avg. ↑
CNN [15]	0.60	0.53	0.64	0.38	0.55	0.59	0.85	0.22	0.37	0.15	0.88	0.60	0.53
D-CNN [52]	0.49	0.39	0.62	0.44	0.53	0.55	0.85	0.25	0.41	0.19	0.87	0.59	0.51
OpenFace 2.0 [4]	0.64	0.50	0.70	0.67	0.59	0.54	0.85	0.39	0.49	0.22	0.85	0.67	0.59
LibreFace (Ours)	0.63	0.72	0.78	0.59	0.59	0.62	0.85	0.36	0.50	0.34	0.94	0.68	0.63

Table 4. Performance of our pipeline with different encoders and distillation on BP4D [51]. Res-18 is our baseline model without any pre-training or distillation. MAE_ViT_base denotes the encoder is replaced by MAE encoder with ViT-base structure being pre-trained. We report the cross-validated F1 score as the evaluation metric for all experiments. Shao *et al.* [36] and Luo *et al.* [30] are two sota methods that share the same evaluation setting.

Method	AU1	AU2	AU4	AU6	AU7	AU10	AU12	AU14	AU15	AU17	AU23	AU24	Avg. ↑
Shao <i>et al.</i> [36]	47.2	44.0	54.9	77.5	74.6	84.0	86.9	61.9	43.6	60.3	42.7	41.9	60.0
Luo <i>et al.</i> [30]	53.7	46.9	59.0	78.5	80.0	84.4	87.8	67.3	52.5	63.2	50.6	52.4	64.7
Res-18	19.6	13.9	33.8	75.8	71.1	80.5	86.4	50.3	23.4	59.0	21.1	31.2	47.2
MAE_ViT_base	57.0	51.1	61.1	76.4	76.1	79.1	86.3	48.3	52.5	64.7	51.3	54.0	63.2
LibreFace (Ours)	49.9	47.8	56.5	77.9	79.6	84.0	87.0	59.0	46.0	63.1	43.2	49.9	62.0

Table 5. Performance of facial expression recognition leveraging our distilled and pre-trained ResNet-18 model evaluated on AffectNet [32] and RAF-DB [25] with overall accuracy and comparison to a few state-of-the-art methods. We report the performance of sota methods.

Method	AffectNet↑	RAF-DB↑
VGG-16 [39]	51.11	80.96
DLP-CNN [53]	54.47	80.89
GAN-Inpainting [46]	52.97	81.87
gaCNN [26]	58.78	85.07
LibreFace (Ours)	49.71	82.79

For AU detection on BP4D [51], we repeat the above pre-training and fine-tuning process and evaluate with three-fold cross-validation as in [30, 36].

We apply the same pre-training strategy for MAE on FER. We take the checkpoint of both the encoder and linear classifier from the MAE trained on the training set of AffectNet and FFHQ, then apply feature-wise knowledge distillation to the ResNet-18 encoder and classifier with the previous MAE checkpoint as the teacher model. We train the pipeline on each training set of AffectNet, FFHQ, and RAF-DB, and report our performance on the test set of each dataset, respectively.

5.5. Results

AU Intensity Estimation. We show the AU intensity estimation results of the proposed LibreFace on DISFA [31] in Table 3. Since OpenFace 2.0 [4] evaluates their model with five-fold cross-validation and reports Pearson Correlation Coefficient (PCC) of the results, we follow their setting to compare fairly. We also compare LibreFace to other methods following this setting. LibreFace achieves 0.63 av-

Table 6. Ablation of our proposed pipeline with different pre-training and encoder settings on DISFA [31]. We also evaluate our pipeline with five-fold cross-validation and use Pearson Correlation Coefficient (PCC), Mean Absolute Error (MAE), and Mean Squared Error (MSE) as the evaluation metric for ablation analysis. **Pre-train** denotes the corresponding encoder is pre-trained on AffectNet [32] and FFHQ [19]. **Distill** denotes the model is trained with the original distillation manner, and **FM_distill** denotes the model is trained with feature-matching distillation. We also use Swin-Transformer tiny structure and Masked Auto-encoder with Vision-Transformer base structure as the encoder, represented as Swin-Tiny and MAE_ViT_base. **Res-18+FM_distill** is the final model we deployed in LibreFace.

Method	PCC↑	MAE↓	MSE↓
Res-18	0.518	0.278	0.352
Res-18+Pre-train	0.614	0.236	0.260
Res-18+Distill	0.620	0.896	0.550
Res-18+FM_distill	0.628	0.244	0.260
Swin-Tiny	0.606	0.252	0.286
Swin-Tiny+Pre-train	0.644	0.236	0.246
MAE_ViT_base	0.668	0.202	0.272
MAE_ViT_base+Pre-train	0.674	0.202	0.270

erage PCC score across all AUs, while OpenFace 2.0 [4] achieves an average PCC of 0.59. In addition to achieving the best average PCC, our method also performs better than other methods on 9 out of total 12 AUs included in DISFA [31]. Since the amount of labeled AUs is limited in DISFA [31], we also fine-tune our model with different encoders on BP4D [51], as shown in Table 4, in order to provide AU detection (presence or absence of AUs) result for AUs not labeled in DISFA [31].

Table 7. Efficiency Comparison between LibreFace and OpenFace 2.0 [4]. We report the average (Avg.) and standard deviation (Std.) of durations of each round of tests and averaged frame rate per second (FPS).

Method	Avg.↓	Std.	FPS↑
OpenFace 2.0 [4] (AU only)	50.11	0.22	19.96
LibreFace (AU only)	25.11	0.60	39.82
LibreFace (Ours)	37.20	0.66	26.88

Table 8. Efficiency and exported executable model size comparison among different encoders in our proposed pipeline. We report the average (Avg.) and standard deviation (Std.) of durations of each round of test, averaged frame rate per second (FPS), and model size in MB(Size).

Method	Avg.↓	Std.	FPS↑	Size↓
Res-18	37.20	0.66	26.88	43
Swin-Tiny	55.94	1.45	17.88	185
MAE_ViT_base	87.93	3.96	11.37	403

Table 9. Efficiency comparison among CPU-only mode and GPU mode in our proposed pipeline. We report the average (Avg.) and standard deviation (Std.) of durations of each round of test, and averaged frame rate per second (FPS)

Method	Avg.↓	Std.	FPS↑
CPU-only	37.20	0.66	26.88
GPU	6.07	0.13	164.82

Facial Expression Recognition. We show the FER results in Table 5, where overall accuracy (\uparrow) is the evaluation metric. We compare to other state-of-the-art methods according to their reported performance in the paper. We achieve comparable results to the state-of-the-art methods, which require far more computation and training/inference time than LibreFace.

Ablation Study and Efficiency Analysis. We provide ablation analysis when different encoders are implemented in our pipeline and their corresponding efficiency analysis. We also compare the computation cost of LibreFace to OpenFace 2.0 [4] in this section. For the efficiency comparison, we run LibreFace with different encoders and OpenFace 2.0 [4]. We run all experiments with **Intel i9-13900K CPU, 64GB DDR5 4800MHz RAM, and Nvidia GTX1080** on **Windows 11** Operating System. E-cores and hyper-threading are disabled. We randomly pick 1000 images in a resolution of 1024×768 from our data and feed them to these toolboxes. We conduct five rounds of tests for each toolbox and calculate averaged running time (Avg.) in second and standard deviation (Std.). We also report the FPS and exported binary model size.

As shown in Table 6, utilizing a more complicated and heavier structure as an encoder, *e.g.*, Swin-Transformer and Masked Auto-encoder, yields better performance. However,

considering the efficiency factor shown in Table 8, we finally adopt ResNet-18 as the encoder and apply feature-wise knowledge distillation to boost the performance. The original knowledge distillation yields better PCC than our baseline (Res-18), but MAE and MSE are significantly worse. In contrast, feature-wise knowledge distillation apparently boosts the performance for all evaluation metrics.

From Table 7, we observe the efficiency improvement of LibreFace compared to OpenFace 2.0 [4]. LibreFace provides more accurate facial expression analysis while running **two times** faster. Noted that *AU only* stands for we only turn on the function of AU intensity estimation and AU Detection, with other output features turned off so that the efficiencies are comparable. We also report the efficiency performance of LibreFace while predicting AU, FER, and other outputs at the same time, where LibreFace still runs **1.5 times** faster than [4]. In Table 9, we observed that running our pipeline on GPUs is generally much faster than running on CPUs, which is in accordance with our expectations.

6. Conclusions and Future work

In this work, we introduce LibreFace, an accurate and flexible framework for facial expression analysis. With proposed feature-wise knowledge distillation and extensive pre-training, we offer a user-friendly toolbox for facial expression analysis. We leverage large-scale networks pre-trained on relevant tasks, then distillate the model to a lightweight structure with feature-wise knowledge to improve efficiency. We offer an easy-to-use toolbox better than OpenFace 2.0 [4] in terms of wider applications, more accurate predictions, and comparable efficient inference.

The current build of our pipeline is tested on Windows platform. In the future, we plan to test our .NET assemblies on Linux and MacOS. All model checkpoints, including ResNet-18 after distillation, pre-trained Swin-Tiny, and pre-trained MAE, together with their corresponding training/inference code and the binary executable program are available on GitHub (link will be provided upon acceptance). As mentioned in the paper, LibreFace will be able to use in a CPU-only environment. We will further provide documents for LibreFace when the user is equipped with a GPU so that more efficient and faster inference will be accomplished with detailed installation and utilization instructions.

7. Acknowledgement

This work is sponsored by the U.S. Army Research Laboratory (ARL) under contract number W911NF-14-D-0005. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

- [1] Tadas Baltrušaitis, Ntombikayise Banda, and Peter Robinson. Dimensional affect recognition using continuous conditional random fields. In *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, pages 1–8. IEEE, 2013. [3](#)
- [2] Tadas Baltrušaitis, Marwa Mahmoud, and Peter Robinson. Cross-dataset learning and person-specific normalisation for automatic action unit detection. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 6, pages 1–6. IEEE, 2015. [3](#)
- [3] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. Openface: an open source facial behavior analysis toolkit. In *2016 IEEE winter conference on applications of computer vision (WACV)*, pages 1–10. IEEE, 2016. [1](#), [3](#), [7](#)
- [4] Tadas Baltrušaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 59–66. IEEE, 2018. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [5] Mina Bishay, Kenneth Preston, Matthew Strafuss, Graham Page, Jay Turcot, and Mohammad Mavadati. Affdex 2.0: A real-time facial expression analysis toolkit. *arXiv preprint arXiv:2202.12059*, 2022. [3](#)
- [6] Dan Bohus, Sean Andrist, Ashley Feniello, Nick Saw, Mihai Jalobeanu, Patrick Sweeney, Anne Loomis Thompson, and Eric Horvitz. Platform for situated intelligence. *arXiv preprint arXiv:2103.15975*, 2021. [6](#)
- [7] Zhixi Cai, Shreya Ghosh, Kalin Stefanov, Abhinav Dhall, Jianfei Cai, Hamid Rezatofighi, Reza Haffari, and Munawar Hayat. Marlin: Masked autoencoder for facial video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1493–1504, 2023. [2](#)
- [8] Celso M de Melo, Jonathan Gratch, and Peter J Carnevale. Humans versus computers: Impact of emotion expressions on people’s decision making. *IEEE Transactions on Affective Computing*, 6(2):127–136, 2014. [1](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [2](#)
- [10] David DeVault, Ron Artstein, Grace Benn, Teresa Dey, Ed Fast, Alesia Gainer, Kallirroi Georgila, Jon Gratch, Arno Hartholt, Margaux Lhomme, et al. Simsensei kiosk: A virtual human interviewer for healthcare decision support. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1061–1068, 2014. [1](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#)
- [12] Paul Ekman. Facial action coding system, 1977. [1](#)
- [13] C Fabian Benitez-Quiroz, Ramprakash Srinivasan, and Aleix M Martinez. Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5562–5570, 2016. [2](#), [3](#), [4](#), [5](#), [6](#)
- [14] Jonathan Gratch, Ron Artstein, Gale Lucas, Giota Stratou, Stefan Scherer, Angela Nazarian, Rachel Wood, Jill Boberg, David DeVault, Stacy Marsella, et al. The distress analysis interview corpus of human and computer interviews. Technical report, University of Southern California Los Angeles, 2014. [1](#)
- [15] Amogh Gudi, H. Emrah Tasli, Tim M. den Uyl, and Andreas Maroulis. Deep learning based face action unit occurrence and intensity estimation. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 06, pages 1–5, 2015. [7](#)
- [16] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [2](#), [3](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#), [3](#)
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [3](#), [5](#)
- [19] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [20] Yury Kartynnik, Artsiom Ablavatski, Ivan Grishchenko, and Matthias Grundmann. Real-time facial surface geometry from monocular video on mobile gpus. *arXiv preprint arXiv:1907.06724*, 2019. [3](#)
- [21] Vahid Kazemi and Josephine Sullivan. One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1867–1874, 2014. [3](#)
- [22] Davis E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009. [3](#)
- [23] Dimitrios Kollias, Viktoriia Sharmanska, and Stefanos Zafeiriou. Distribution matching for heterogeneous multi-task learning: a large-scale face study. *arXiv preprint arXiv:2105.03790*, 2021. [1](#)
- [24] Dimitrios Kollias and Stefanos Zafeiriou. Affect analysis in-the-wild: Valence-arousal, expressions, action units and a unified framework. *arXiv preprint arXiv:2103.15792*, 2021. [2](#), [3](#), [4](#), [5](#)
- [25] Shan Li, Weihong Deng, and JunPing Du. Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2584–2593. IEEE, 2017. [6](#), [7](#)

- [26] Yong Li, Jiabei Zeng, Shiguang Shan, and Xilin Chen. Occlusion aware facial expression recognition using cnn with attention mechanism. *IEEE Transactions on Image Processing*, 28(5):2439–2450, 2018. 4, 7
- [27] 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, 3
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [29] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuoling Chang, Ming Yong, Juhyun Lee, Wan-Teh Chang, Wei Hua, Manfred Georg, and Matthias Grundmann. Mediapipe: A framework for perceiving and processing reality. In *Third Workshop on Computer Vision for AR/VR at IEEE Computer Vision and Pattern Recognition (CVPR) 2019*, 2019. 3, 4, 6
- [30] Cheng Luo, Siyang Song, Weicheng Xie, Linlin Shen, and Hatice Gunes. Learning multi-dimensional edge feature-based au relation graph for facial action unit recognition. *arXiv preprint arXiv:2205.01782*, 2022. 7
- [31] S Mohammad Mavadati, Mohammad H Mahoor, Kevin Bartlett, Philip Trinh, and Jeffrey F Cohn. Disfa: A spontaneous facial action intensity database. *IEEE Transactions on Affective Computing*, 4(2):151–160, 2013. 4, 5, 6, 7
- [32] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017. 5, 6, 7
- [33] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NeurIPS Autodiff Workshop*, 2017. 6
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 3
- [35] Jason M Saragih, Simon Lucey, and Jeffrey F Cohn. Deformable model fitting by regularized landmark mean-shift. *International journal of computer vision*, 91(2):200–215, 2011. 3
- [36] Zhiwen Shao, Zhilei Liu, Jianfei Cai, and Lizhuang Ma. Deep adaptive attention for joint facial action unit detection and face alignment. In *European Conference on Computer Vision*, pages 725–740. Springer, 2018. 7
- [37] Zhiwen Shao, Zhilei Liu, Jianfei Cai, and Lizhuang Ma. Jaanet: Joint facial action unit detection and face alignment via adaptive attention. *International Journal of Computer Vision*, 129(2):321–340, 2021. 1
- [38] Zhiwen Shao, Zhilei Liu, Jianfei Cai, Yunsheng Wu, and Lizhuang Ma. Facial action unit detection using attention and relation learning. *IEEE transactions on affective computing*, 13(3):1274–1289, 2019. 1
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 7
- [40] Kalin Stefanov, Baiyu Huang, Zongjian Li, and Mohammad Soleymani. Opensense: A platform for multimodal data acquisition and behavior perception. In *Proceedings of the 2020 International Conference on Multimodal Interaction, ICMI '20*, page 660–664, New York, NY, USA, 2020. Association for Computing Machinery. 6
- [41] Lior Wolf, Tal Hassner, and Itay Maoz. Face recognition in unconstrained videos with matched background similarity. In *CVPR 2011*, pages 529–534. IEEE, 2011. 2
- [42] Erroll Wood, Tadas Baltrusaitis, Xucong Zhang, Yusuke Sugano, Peter Robinson, and Andreas Bulling. Rendering of eyes for eye-shape registration and gaze estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3756–3764, 2015. 3
- [43] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 2
- [44] Yinghao Xu, Yujun Shen, Jiapeng Zhu, Ceyuan Yang, and Bolei Zhou. Generative hierarchical features from synthesizing images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4432–4442, 2021. 2
- [45] Jing Yang, Brais Martinez, Adrian Bulat, Georgios Tzimiropoulos, et al. Knowledge distillation via softmax regression representation learning. *International Conference on Learning Representations (ICLR)*, 2021. 2, 3, 5
- [46] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5505–5514, 2018. 4, 7
- [47] Amir Zadeh, Yao Chong Lim, Tadas Baltrusaitis, and Louis-Philippe Morency. Convolutional experts constrained local model for 3d facial landmark detection. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 2519–2528, 2017. 3
- [48] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE signal processing letters*, 23(10):1499–1503, 2016. 3
- [49] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *International Conference on Learning Representations*, 2021. 2, 3
- [50] Tenggao Zhang, Chuanhe Liu, Xiaolong Liu, Yuchen Liu, Liyu Meng, Lei Sun, Wenqiang Jiang, and Fengyuan Zhang. Emotion recognition based on multi-task learning framework in the abaw4 challenge. *arXiv preprint arXiv:2207.09373*, 2022. 4
- [51] Xing Zhang, Lijun Yin, Jeffrey F Cohn, Shaun Canavan, Michael Reale, Andy Horowitz, Peng Liu, and Jeffrey M Girard. Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014. 6, 7
- [52] Kaili Zhao, Wen-Sheng Chu, and Honggang Zhang. Deep region and multi-label learning for facial action unit detection.

In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3391–3399, 2016. [7](#)

- [53] Liming Zhao, Xi Li, Yueting Zhuang, and Jingdong Wang. Deeply-learned part-aligned representations for person re-identification. In *Proceedings of the IEEE international conference on computer vision*, pages 3219–3228, 2017. [4](#), [7](#)