

OpticalDR: A Deep Optical Imaging Model for Privacy-Protective Depression Recognition

Yuchen Pan, Junjun Jiang*, Kui Jiang, Zhihao Wu, Keyuan Yu, Xianming Liu
 Harbin Institute of Technology

Abstract

Depression Recognition (DR) poses a considerable challenge, especially in the context of the growing concerns surrounding privacy. Traditional automatic diagnosis of DR technology necessitates the use of facial images, undoubtedly expose the patient identity features and poses privacy risks. In order to mitigate the potential risks associated with the inappropriate disclosure of patient facial images, we design a new imaging system to erase the identity information of captured facial images while retain disease-relevant features. It is irreversible for identity information recovery while preserving essential disease-related characteristics necessary for accurate DR. More specifically, we try to record a de-identified facial image (erasing the identifiable features as much as possible) by a learnable lens, which is optimized in conjunction with the following DR task as well as a range of face analysis related auxiliary tasks in an end-to-end manner. These aforementioned strategies form our final Optical deep Depression Recognition network (OpticalDR). Experiments on CelebA, AVEC 2013, and AVEC 2014 datasets demonstrate that our OpticalDR has achieved state-of-the-art privacy protection performance with an average AUC of 0.51 on popular facial recognition models, and competitive results for DR with MAE/RMSE of 7.53/8.48 on AVEC 2013 and 7.89/8.82 on AVEC 2014, respectively. Code is available at <https://github.com/divertingPan/OpticalDR>.

1. Introduction

Automated depression diagnosis (ADD) approaches have emerged based on various biometric features, including acoustics, expressions, and electroencephalogram [1, 8, 31, 41]. Among these approaches, facial-based ADD has become mainstream [10, 11, 26], offering the advantages of easy access and non-contact assessment. However, facial images from depression patients are particularly sensitive, as they inherently encompass biometric identifiers and pose

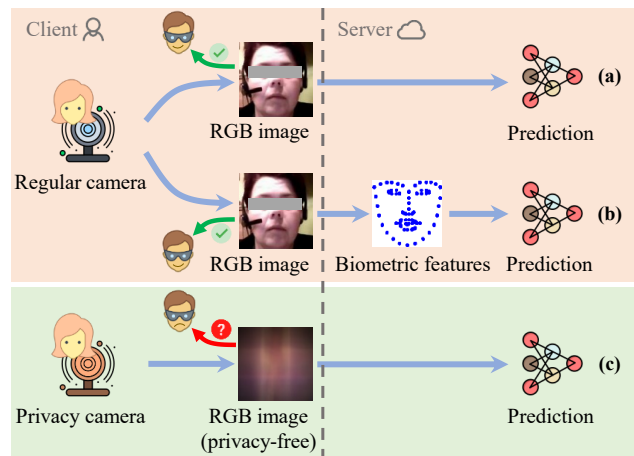


Figure 1. Different DR approaches: (a) conventional facial recognition with no privacy preservation, (b) facial features-based approaches with limited privacy preservation, and (c) our full privacy-preserving approach that doesn't generate facial images. The DR approaches with regular cameras could be at the risk of sensitive information access by attackers after the images are captured. However, with our approach, no sensitive information would be captured or stored, whether on the client or the server.

privacy risk. Consequently, it is of paramount importance to address the safeguarding of facial information when developing ADD systems to uphold medical privacy and security.

To minimize these risks, previous studies [10, 11, 26] introduce facial image anonymization techniques. Rather than the full facial images, facial biometric features (identity-protected) such as facial landmarks are used for DR (Fig. 1). These endeavors, however, are still dependent on the presence of visually accessible images for feature extraction, which poses a significant threat to privacy. Therefore, it is crucial to develop a more secure and robust facial imaging approach at the initial stage of data collection.

In the field of Deep Optics [32], the lens can be parameterized to a distinct layer within the deep learning model. Inspired by this, in this paper we integrate the Deep Optics into deep learning network and advises an Optical deep DR network – which we term a OpticalDR – that provides

*Correspondence author: Junjun Jiang (junjunjiang@hit.edu.cn)

a graceful solution for privacy-protective DR. Compared to the fixed blurring or defocus image processing methods, which might compromise depression-related features when removing privacy information, the proposed lens is learnable and can be optimized in an end-to-end manner, like a specialized feature extractor on the front of deep network. As demonstrated in Fig. 2, we develop a learnable single-lens optical model to extract depression-related features while erasing sensitive privacy information from the observation. In the subsequent deep learning model, we design a progressive learning tactic to train our network. To harness abundant supervision (label) information from various face datasets, we incrementally incorporate different tasks, *i.e.*, identity recognition, emotion recognition and DR, to assist our deep DR network training. It is worth noting that acquiring a complete dataset with labels for identity, emotion, and depression is extremely challenging in practice. Consequently, the adoption of progressive learning offers a viable strategy to leverage the available supervision information from existing datasets. Our contributions can be summarized as follows:

1. We have successfully simulated a privacy lens with parameterized model. The simulated lens is co-optimized alongside the subsequent model to ensure facial privacy protection while retaining valuable features.
2. We introduce OpticalDR, a novel framework that integrates an optical lens with a deep learning model to extract depression-related information from privacy-preserving images. The entire pipeline operates without the need for storing or transmitting facial images, ensuring a heightened level of privacy protection.
3. Our experimental results demonstrate competitiveness on the AVEC 2013 and AVEC 2014 datasets in comparison to existing DR approaches. Furthermore, the optical component in OpticalDR exhibits robustness against various types of attacks.

2. Related Works

Visual Depression Recognition. Contemporary research in DR primarily relies on audio and video signals obtained from volunteers during interviews. Datasets such as AVEC 2013 [40] and AVEC 2014 [39] contain original facial images within video recordings. Some studies [3, 25] have concentrated on recognizing depression from video frames that contain the original facial images. Early approaches were centered on hand-crafted designed feature extraction and appropriate regression models [13, 40] for DR using visual data. Growing concerns about privacy have emerged. As a result, the AVEC 2017 [28] dataset takes a privacy-centric approach, providing only facial landmarks and other

manually engineered facial features. Some research employed facial landmarks and other physiological data as facial features, often considering the motion of landmark points [27] or incorporating factors like Action Units and head pose [10, 35]. Furthermore, studies have shown that human behavior representations [33, 34] can significantly contribute to DR from facial data. More recently, deep learning-based methods [42, 43], or a combination of hand-crafted and deep features [30, 37], have been applied, resulting in substantial improvements in recognition performance. Nevertheless, approaches relying on facial features may not necessitate the use of original images during inference, but these features still need to be extracted from facial images, exposing potential risks of facial data leakage.

Deep Optics for Privacy Protection. Recently, Deep Optics [32] has gained more of the spotlight in research communities with the advancement of deep learning. It involves the joint optimization of optical systems and downstream tasks, and has demonstrated success in various computer vision tasks such as depth estimation [18], 3D object detection [5], extended depth of field, super-resolution imaging [32], and image classification [4]. The underlying philosophy is to capture the downstream task friendly images to improve the performance of high-level applications.

To achieve a privacy protection high-level visual task, the optical imaging system is designed to intentionally degrade image quality and obscure private information, while still enabling downstream tasks. This approach has been successfully demonstrated in lensless camera facial recognition [14, 19], deep optimized lens cameras for action recognition [15], and pose estimation [16]. This work focuses on exploring facial privacy protection using Deep Optics, primarily on facial images. The goal is to extract facial depression features from privacy-preserving images for a DR system with enhanced privacy.

3. The Proposed Method

We are primarily focused on the task of privacy-preserving DR. The main objective is to create a personal information erased “face” image using a specially optimized lens, which is achieved through joint optimization with the DR model. Therefore, OpticalDR offers privacy preservation at a hardware level, ensuring that no privacy information is generated or stored in digital devices. This entire process is illustrated in Fig. 2. The OpticalDR comprises two key components: the optical model and the deep feature extraction model. In the optical model, we parameterize a learnable single optical element for generating privacy preserving image (Sec. 3.1). The deep feature model is responsible for the subsequent extraction of emotional and depression features and the fusion of these features for the final prediction (Sec. 3.2). Following that, we explain

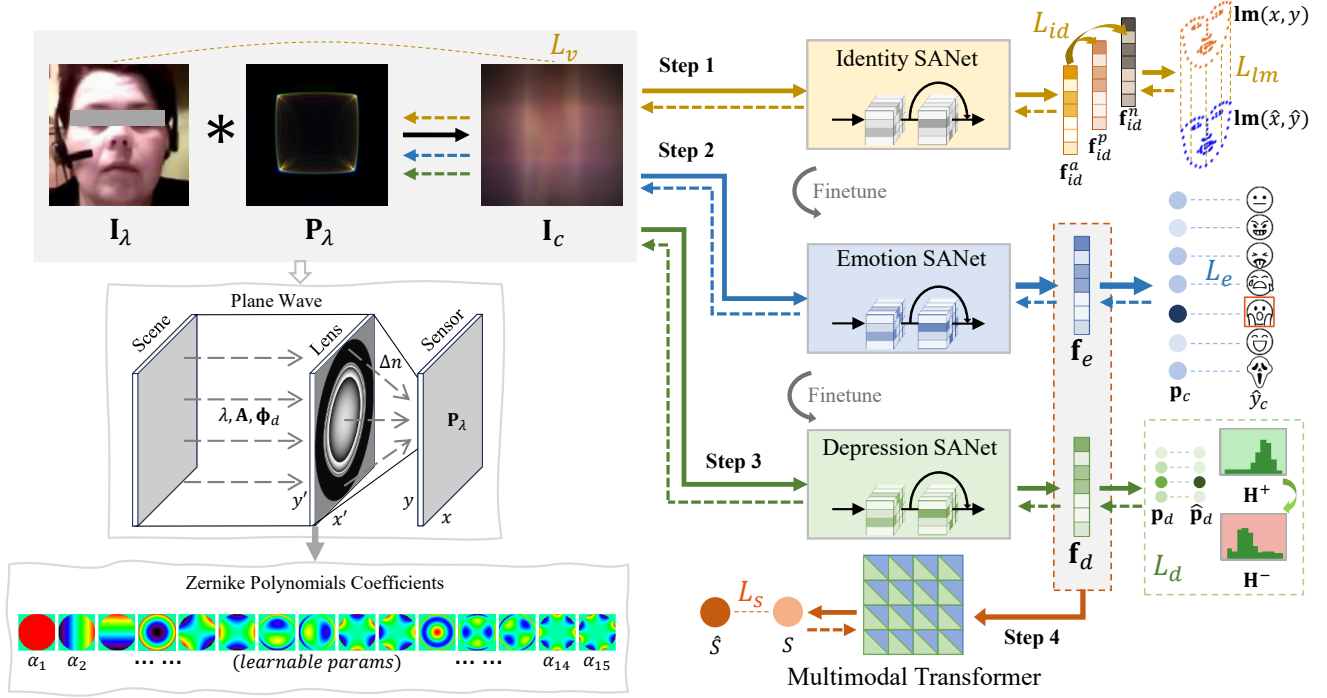


Figure 2. The architecture of the proposed OpticalDR. The optical model comprises a thin lens for image capture in front of the sensor. The deep learning model utilizes the captured image and is jointly optimized with the optical component for depression score recognition.

the auxiliary tasks and the corresponding loss functions designed for optimizing the OpticalDR (Sec. 3.3).

3.1. Privacy Erasing Optical Model

The optical model corresponds to the image generation process captured by the camera sensor. The objective of this model is to produce an image that is visually unrecognizable and doesn't reveal the identity of the subject. Simultaneously, this process must preserve the important features required for DR. This is achieved through end-to-end joint optimization with the feature extraction model, utilizing appropriate loss functions.

To accomplish this, we derive a wave-based image formation model using a single thin lens, as discussed in [32]. The key variables in the optical model include the surface profile denoted as \mathbf{H} , represented by Zernike term coefficients α , and the PSF represented as \mathbf{P} of the lens, which can be derived from \mathbf{H} . The image passing through the lens can be calculated from \mathbf{P} using the following equation:

$$\mathbf{I}_c(x, y) = \int (\mathbf{I}_\lambda * \mathbf{P}_\lambda)(x, y) \kappa_c(\lambda) d\lambda + \eta, \quad (1)$$

where \mathbf{I}_λ represents the original scene from the real world with the particular light wavelength λ which corresponds to the red, green, and blue light in the RGB image model.

κ_c represents the sensitivity of sensors to these three wavelengths. \mathbf{I}_c represents the generated image on the sensor plane. Additionally, we account for sensor noise by introducing the term η , which is modeled as Gaussian noise with $\eta \sim \mathcal{N}(0, \sigma^2)$.

The phase change induced by the thin lens is given by

$$\Phi(x', y') = \frac{2\pi\Delta n}{\lambda} \mathbf{H}(x', y'). \quad (2)$$

Here, Δn represents the refractive index difference between air and the lens material. \mathbf{H} can be expressed as:

$$\mathbf{H} = \sum_{j=1}^q \alpha_j \mathbf{Z}_j, \quad (3)$$

where \mathbf{Z}_j represents the j -th term of the Zernike polynomial in Noll notation [24].

Assuming a point light source at optical infinity with amplitude \mathbf{A} and phase Φ_d . The PSF \mathbf{P} for the sensor at a distance z can be formulated as follows:

$$\mathbf{P}_\lambda(x, y) = \left| \mathcal{F} \left\{ \mathbf{A}(x', y') e^{i\Phi(x', y')} e^{i\frac{\pi}{\lambda z}(x'^2 + y'^2)} \right\} \right|^2. \quad (4)$$

In this way, the PSF \mathbf{P} of the learnable optical element can be simulated through the introduced wave-based image formation module. Since the \mathbf{P} is determined by the optical

height map \mathbf{H} , we learn from Eq. 3 that we only need to learn the optimization variable $\{\alpha_j\}_{j=1}^q$.

3.2. Recognition from Privacy-Preserving Images

In this section, we elaborate on the proposed multi-task deep network for the extraction of intrinsic facial features. In the realm of DR, the incorporation of spatial feature relationships of face image is very important and has shown great efficacy [25, 30]. Therefore, we propose to leverage the Spatial Attention Network (SANet) for the extraction of intrinsic facial features. Importantly, this module can be seamlessly integrated into any convolutional network architecture. In the context of our SANet implementation, we have chosen to adopt a ResNet [12] backbone enhanced with the SA module.

Our approach unfolds in four progressive learning stages. Initially, we pretrain the SANet for recognizing facial features from optically blurred images. This step will harness the benefits of large-scale face recognition database to erase the identity information as well as preserve the facial structure information by the landmark prediction loss. Subsequently, the SANet undergoes fine-tuning to extract emotional features from the blurred image to further extract the DR-related information. Finally, we progressively train the SANet by the DR score loss to represent the depression-related features. The emotional and depression information extracted from these two SANets is harmoniously fused using a multimodal transformer [36] layer to obtain the final depression score.

The fusion of the emotional and depression feature is executed through the multimodal transformer. This operation can be succinctly expressed in the following simplified formulation for the two modalities, representing emotional and depression conditions:

$$\mathbf{Z}_{\{E,D\}} = [\mathbf{Z}_{E \rightarrow D}; \mathbf{Z}_{D \rightarrow E}], \quad (5)$$

where $\mathbf{Z}_{E \rightarrow D}$ and $\mathbf{Z}_{D \rightarrow E}$ correspond to the cross-modality attention information, respectively. In this way, the adoption of progressive learning offers an efficient strategy to leverage the available supervision information from different tasks.

3.3. Joint Optimization of Optical Model and DR

In practical, it is very hard to obtain a full dataset with the labels of identity, emotion and depression, simultaneously. Therefore, in this paper we introduce the progressive learning to train our OpticalDR network, which requires progressively training of the optical lens, a SANet for emotion recognition and a SANet for DR. The final score is predicted by a fusion model with extra training with only the DR dataset. In the following, we will present the details to obtain our DR network.

Step 1: Pretrain the privacy preserving lens. To achieve privacy preservation, we try to render the images that are visually indistinct from the input image, and unidentifiable by deep learning face recognition models. For visual degradation, we maximize the Mean Squared Error (MSE) between the generated image and the real image, $L_v = \|\mathbf{I}_c - \mathbf{I}_\lambda\|_2^2$.

To further obfuscate the identity features in the images, we adopt the inverse triplet loss defined as:

$$L_{id} = \max(d(\mathbf{f}_{id}^a, \mathbf{f}_{id}^n) - d(\mathbf{f}_{id}^a, \mathbf{f}_{id}^p) + m, 0), \quad (6)$$

where m is a margin value, $\mathbf{f}_{id}^a, \mathbf{f}_{id}^p, \mathbf{f}_{id}^n$ denote the feature embeddings extracted by SANet from the anchor image, the image with the same identity as the anchor, and the image with a different identity than the anchor after passing through the lens. $d(\cdot)$ represents the feature distance. We employ hypersphere distance [21] with a margin r to measure it, *i.e.*, $d(x, y) = \|\|x\|_2^2 - \|y\|_2^2\| + r$.

Landmarks have proven to be crucial for conveying essential features in DR [26]. To ensure that the blurred image retains these landmarks and the vital depression-related features, we introduce a landmark recognition loss L_{lm} , which can be calculated by the differences of coordinates of the ground truth and predicted landmarks. We obtain the landmarks followed the extracted facial embedding and using non-linear mapping from embedding to landmarks.

The total loss for optimizing the identity information can be formulated as:

$$L_i = \alpha L_v + \beta L_{id} + \gamma L_{lm}, \quad (7)$$

where α, β , and γ are the balancing parameters and are set to $-0.1, 1$ and 1 , respectively.

Step 2: Integrate emotional information. In this step, we incorporate emotional information for two primary purposes. First, it serves as emotional prior knowledge to aid in DR, as depression patients often exhibit abnormal emotional arousal. Second, it facilitates the model in learning emotional features from the image captured by the lens. The SANet utilizes the weights obtained from the SANet trained for identity obfuscation in *Step 1* and then is trained to discriminate the facial emotion classes \hat{y}_c with its prediction probability of each class p_c , using the following cross-entropy loss:

$$L_e = - \sum_c \hat{y}_c \log(p_c). \quad (8)$$

The last feature layer of the SANet is retained as the emotional feature \mathbf{f}_e .

Step 3: Integrate depression information. To introduce a coarse-grained feature of depression, and bridge the features of emotion and depression, we employ a SANet

Table 1. Mapping of depression levels to scores based on the BDI-II criteria in the depression stage training for L_d

Score	0-13	14-19	20-28	29-63
Level	0 (minimal)	1 (mild)	2 (moderate)	3 (severe)

that utilizes the weights obtained from the SANet trained in *Step 2*. The SANet is fine-tuned to gather information about depression levels.

In contrast to prior label distribution designs [43], we convert the real-valued depression score into a discrete probability distribution, represented as $\mathbf{p} = [p_1, p_2, \dots, p_c]$, where c represents the classes of depression levels based on depression scores, as outlined in Tab. 1. The fine-tuning process minimizes the distribution similarity between the SANet’s prediction \mathbf{p}_d and $\hat{\mathbf{p}}_d$, and employs a histogram loss for the label-aware similarity between positive feature pairs (where both samples are from the same depression level) and negative pairs, then define the histogram \mathbf{H}^+ and \mathbf{H}^- as defined in [38]. Then the depression metric loss can be defined as:

$$L_d = LAH(\hat{\mathbf{p}}_d, \mathbf{p}_d, \mathbf{H}^+, \mathbf{H}^-). \quad (9)$$

The depression feature is extracted from the last feature layer of the SANet and is represented as \mathbf{f}_d .

Step 4: Fuse emotional and depression features. Finally, a fusion model is expected to fuse the emotional feature and the coarse-grained depression feature for fine-grained depression score prediction. The complete OpticalDR comprises the pretrained weights of the optics model, the emotion SANet, and the depression SANet. At this stage, the network is frozen, and a final fusion layer is introduced. This layer is trained using the depression score \hat{S} to enable the transformer to learn the mapping from emotional and depression features to the predicted depression score S . The optimization of the fusion layer is achieved using the MSE loss:

$$L_s = \sum \frac{1}{n} (\hat{S} - S)^2. \quad (10)$$

Above-mentioned four steps are sequentially applied to optimize the corresponding part of OpticalDR, as depicted in Fig. 2, by the specified loss functions. The emotion SANet is fine-tuned from the identity SANet, and the depression SANet is fine-tuned from the emotion SANet to incorporate prior information.

4. Experiments

4.1. Datasets

In our experiments, we utilized the following datasets: i) **CelebA** [22]: This dataset comprises human face images

with corresponding identity labels, totaling 10,177 identities and 202,599 samples. We assigned the last 1,000 samples for the validation set and the remainder for training. ii) **CK+** [23]: It comprises 593 series of facial images from 123 identities, spanning 7 emotion classes for all samples. We allocated 300 samples for validation and utilized the remainder for training. iii) **AVEC 2013** [40]: This dataset includes 150 video clips with 82 unique participants, exclusively employed in the validation and testing phases of our experiment. iv) **AVEC 2014** [39]: Comprising 300 video clips from 83 individuals, both AVEC 2013 and AVEC 2014 record data in a human-computer interaction scenario. The labels, assessed using the Beck Depression Inventory-II [2], range from 0 to 63. Frames extracted from both datasets are cropped, aligned, and employed for training, validation, and testing based on the official split.

4.2. Analysis of Optical Lens

To assess the impact of DR when employing an optimizable lens to distort images in comparison to conventional image blurring methods, we established a pipeline for images subjected to various privacy-preserving strategies. We conducted comparisons of the results obtained from Gaussian blurring with different σ of the Gaussian kernel, and different defocus strengths using an approximate bokeh PSF with different bokeh ball diameters. To evaluate the privacy-preserving efficacy of each solution, we employed four widely used face recognition models: IR152 [12], IRSE50 [17], ArcFace [7], and Facenet [29], to measure the recognizability in terms of the Area Under the Curve (AUC). Furthermore, Fig. 3 presents the Receiver Operating Characteristic (ROC) curves for our OpticalDR and alternative strategies. Detailed comparisons of DR and privacy-preserving performance are provided in Fig. 4. It is observed that increasing blurring strength enhances privacy preservation (closer to random guess ROC), but this comes at the cost of a substantial decrease in DR performance. Notably, our learnable privacy-preserving approach demonstrates a significant advantage by achieving the lowest prediction error in Mean Absolute Error (MAE) and Root Mean Square Error (RMSE). This suggests that the optical lens retains more useful features even when the image is significantly distorted.

From the colored lines in Tab. 2, it is evident that while Gaussian blur and defocus approaches show similar performance to the optics solution in privacy preservation, the recognition error increases (yellow lines), and vice versa (blue lines). In the case of defocus with a diameter of 275 and the Gaussian blurring with σ is 50, the privacy preserving performance is quite similar to the optics solution, but the visualization in Fig. 5 reveals additional identify information such as skin color and a potential gender of the subject. In the broader context of privacy, this information

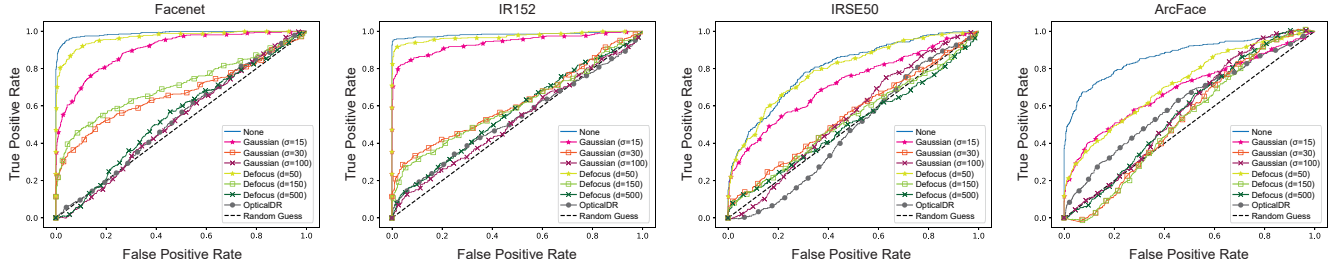


Figure 3. ROC curves depicting the performance of facial recognition models under privacy-preserving approaches, including Gaussian blur and defocus methods, alongside our OpticalDR.

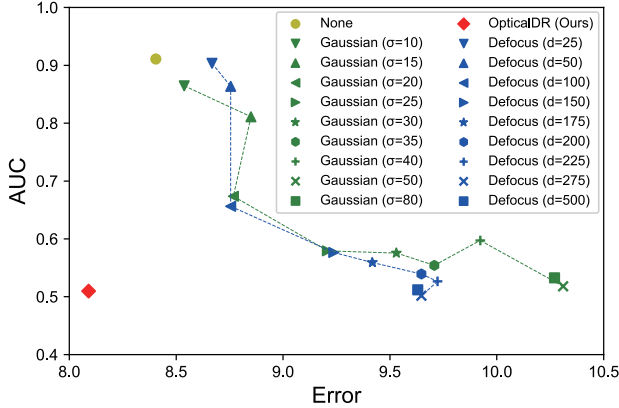


Figure 4. Trade-off between privacy preservation and DR performance among different privacy-preserving strategies.

is also considered sensitive and should be safeguarded. In contrast, the optics solution exhibits a more robust image against various sensitive information leakage.

4.3. Analysis of Privacy Robustness

We conducted a series of tests to evaluate the deconvolution method on facial images captured through the lens. These tests encompassed the following scenarios: i) **Non-blind deconvolution with known PSF**: This test simulates a scenario in which an attacker has access to the device and can obtain the PSF of the camera. We applied the non-blind Wiener deconvolution method under these conditions; ii) **Deblur GAN**: We investigated the effectiveness of blind restoration approaches using GAN-based image restoration techniques, as outlined in [20]. This method simulates a scenario in which an attacker has access to the camera and can collect clear and blurred image pairs, but doesn't get the PSF. The attacker can employ these pairs to train a deblurring model for image recovery.

The input data for these tests consisted of blurred images and their corresponding restored images, as illustrated in Fig. 6. It is evident that both the deconvolution and deblurring approaches exhibited limitations in restoring the images, characterized by a multitude of aberrations and significant loss of visual information.

4.4. Analysis of Facial Features

We explored the cumulative impact of incorporating information from identity (I), emotions (E), depression features (D), and the fusion model for depression score prediction (S). Throughout the auxiliary task, we assessed the performance of the respective tasks, and the results are presented in Tab. 3. Regarding *Step 1* (I), our results reveal that the generated images, after training the lens, effectively erase identity information, as evidenced by an AUC of identity recognition near 0.5.

Subsequently, the emotion recognition task in *Step 2* and the DR task in *Step 3* involve the incorporation of emotion and depression-related information, which is utilized in the final depression score prediction. We conducted tests using four training methods at *Step 3*: i) Training both the lens and the SANet from scratch (D); ii) Training the lens from scratch while initializing the SANet with the weights obtained in *Step 3* ($E \rightarrow D$); iii) Training the SANet from scratch while fine-tuning the lens acquired in *Step 1* ($I \rightarrow D$); iv) Loading both the lens and SANet weights acquired in *Step 2* ($I \rightarrow E \rightarrow D$). These trained models only contain the lens and the depression SANet, as shown in *Step 3* in Fig. 2. Then the models underwent evaluation to assess their performance in predicting depression scores.

Comparing cases with emotion recognition task ($E \rightarrow D$ and $I \rightarrow E \rightarrow D$) to those without emotion recognition pre-training information (D and $I \rightarrow D$) signifies whether emotional information stored in the network is utilized in extracting depression-related features. It is observed that utilizing emotional information significantly improves the DR performance, with an approximately 10% improvement, underscoring the importance of incorporating emotional information for extracting depression-related features.

Contrasting cases with lens pretraining ($I \rightarrow D$ and $I \rightarrow E \rightarrow D$) and without lens pretraining (D and $E \rightarrow D$) evaluates the impact of the blend of identity features in the optical lens on DR. The incorporation of these identity features extracted by the lens can boost the recognition performance, but it is less crucial compared to the emotional features extracted from the SANet (7% vs. 10% improvement). This suggests that information stored within the SANet is more

Table 2. DR performance in terms of MAE and RMSE of different privacy-preserving methods evaluated on AVEC 2013 and AVEC 2014 datasets. Additionally, privacy-preserving performance assessed on the CelebA dataset, measuring the AUC under various facial recognition models. Average error and AUC values are computed for a comprehensive comparison.

Methods	AVEC 2013		AVEC 2014		AUC (CelebA)				$\overline{\text{Error}}$	$\overline{\text{AUC}}$
	MAE	RMSE	MAE	RMSE	Facenet	IR152	IRSE50	ArcFace		
None	7.25	9.57	7.27	9.53	0.99	0.98	0.80	0.87	8.41	0.91
Gaussian ($\sigma=15$)	7.81	9.94	7.82	9.83	0.90	0.94	0.72	0.69	8.85	0.81
Gaussian ($\sigma=25$)	8.37	10.06	8.47	9.93	0.67	0.61	0.55	0.51	9.21	0.59
Gaussian ($\sigma=50$)	9.44	11.33	9.13	11.34	0.58	0.57	0.54	0.52	10.31	0.55
Defocus ($d=50$)	7.67	9.84	7.78	9.73	0.97	0.97	0.79	0.72	8.76	0.86
Defocus ($d=150$)	8.13	10.39	8.14	10.29	0.71	0.60	0.53	0.53	9.24	0.59
Defocus ($d=275$)	8.66	10.73	8.64	10.56	0.53	0.54	0.50	0.51	9.65	0.52
OpticalDR	7.92	8.39	7.48	8.57	0.51	0.51	0.51	0.51	8.09	0.51




Figure 5. The visual effects of (a) the original facial image, (b) Gaussian blur with different sigma values, (c) defocus with varying diameters of bokeh balls, and (d) the optical solution for facial privacy preservation.

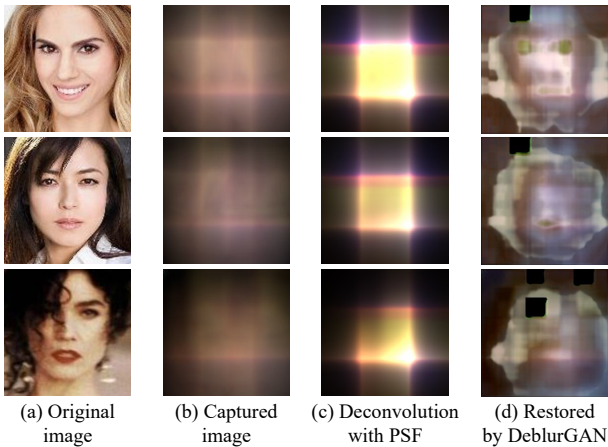


Figure 6. A quantitative comparison of samples from CelebA is provided. The displayed images include the simulated captured images along with their corresponding deblurred results. It is noteworthy that the restored images generally exhibit poor visual recognition of human faces and appearances.

critical in DR than that stored in the lens, although both contribute to performance enhancement.

Furthermore, comparing the model with lens and SANet trained from scratch (D), a noticeable performance improvement is observed when pretraining the emotion SANet and concurrently training the lens from scratch ($E \rightarrow D$). However, pretraining the lens while omitting information

Table 3. Performance of the proposed auxiliary tasks during the pretraining stages. I , E , D , S represent the *Step 1* to *Step 4* from the previous joint optimization stages, respectively.

Task	CelebA	AVEC 2013		AVEC 2014		$\overline{\text{Error}}$
	AUC	MAE	RMSE	MAE	RMSE	
I	0.5036	-	-	-	-	-
D	0.5352	10.35	11.28	10.31	11.25	10.80
$E \rightarrow D$	0.5328	9.18	10.04	9.06	9.88	9.54
$I \rightarrow D$	0.5026	10.15	11.03	10.27	11.15	10.65
$I \rightarrow E \rightarrow D$	0.5038	8.54	9.63	8.87	9.97	9.25
S	0.5367	8.46	11.12	8.62	11.29	9.87
$I \rightarrow S$	0.5013	8.24	10.19	8.60	10.62	9.41
$I \rightarrow E \rightarrow D \rightarrow S$	0.5016	7.92	8.39	7.48	8.57	8.09

in the emotion SANet ($I \rightarrow D$) results in only a slight performance improvement. This underscores that the emotion model already contains advanced features suitable for recognizing images obscured by the lens, particularly beneficial for DR.

To demonstrate the necessity of the designed auxiliary tasks in OpticalDR, we conducted tests with i) training the entire model without any auxiliary task but only with L_s (S), ii) pretraining only using the privacy-preserving lens in *Step 1* ($I \rightarrow S$), and iii) the full pretraining strategy ($I \rightarrow E \rightarrow D \rightarrow S$). Comparison of DR errors indicates the effectiveness of auxiliary tasks for improved DR performance. Moreover, AUC values in Tab. 3 show that the model maintains privacy preservation as training steps progress, evidenced by close

Table 4. Comparison with previous methods for DR on the testing sets of AVEC 2013 and AVEC 2014.

Method	Privacy-Preserving	AVEC 2013		AVEC 2014	
		MAE	RMSE	MAE	RMSE
Baseline [39,40]	×	10.88	13.61	8.86	10.86
ViT [9]	×	8.14	10.67	8.37	10.59
STA-DRN (SA) [25]	×	7.25	9.57	7.27	9.53
LQGDNet [30]	×	6.38	8.20	6.08	7.84
LQGDNet + lens	✗	10.56	13.43	10.27	13.41
ViT + lens	✓	8.16	10.49	8.25	10.53
STA-DRN (SA) + lens	✓	8.62	9.53	8.67	9.68
OpticalDR	✓	7.53	8.48	7.89	8.82

facial recognition AUC between $I \rightarrow S$ and $I \rightarrow E \rightarrow D \rightarrow S$. This demonstrates that training with emotional and depression information does not compromise privacy-preserving ability.

4.5. Comparison with the State-of-the-Arts

We compared OpticalDR with state-of-the-art depression recognition (DR) methods on the AVEC 2013 and AVEC 2014 test sets. The comparisons include (1) *ViT* [9]: A state-of-the-art universal vision model. (2) *LQGDNet* [30]: A state-of-the-art single-image facial DR approach. (3) *STA-DRN (SA)* [25]: A state-of-the-art 3D video-based DR approach utilizing the proposed spatial module. Additionally, we jointly trained these models with a lens for privacy-preserving purposes.

Analyzing Tab. 4 yields key insights. Firstly, integrating optical components directly into state-of-the-art approaches, without employing additional learning strategies for feature enhancement, does not arise a significant improvement in performance. Furthermore, *LQGDNet*, the utilization of quaternion representation of features within facial images, still requires a visible facial image as input. Thus, we denoted this approach as ✗ in Tab. 4. The non-differentiable nature of the quaternion feature extractor poses a challenge for joint optimization. Additionally, attempting to optically blur the global facial image can negatively impact the performance of DR. Conversely, when *ViT* is coupled with a lens, there is a slight improvement in performance, especially in terms of RMSE. This improvement indicates that optical interventions can offer additional information prior to *ViT*'s processing, which can have a positive impact on DR.

Our *OpticalDR* achieves superior DR performance compared to *ViT + lens* and *STA-DRN (SA) + lens*, underscoring the effectiveness of our approach over directly joint training of optical and deep learning models. However, it falls short of outperforming *LQGDNet*, which leverages quaternion representation for extracting local features, still depending on visible clear facial images. This indicates that *LQGDNet* benefits from explicit prior information on local features like eyes and mouth. In contrast, *OpticalDR*, relying on the *SANet*, adaptively learns spatial features without

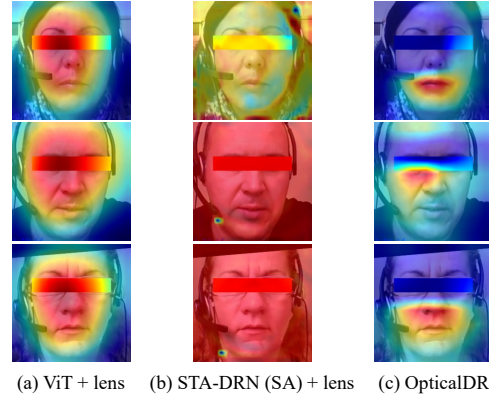


Figure 7. CAM results for optical privacy-preserving DR approaches.

the need for such explicit information, as shown in Fig. 7.

Visualizing CAM [6] results for *OpticalDR* and comparing with *ViT + lens* and *STA-DRN (SA) + lens* (Fig. 7), *OpticalDR* activates on specific regions with clear borders, potentially containing features for DR. Contrastingly, the simple combination of optical and vision deep models results in nonsensical or unchanging activation on the facial area, underscoring the necessity and efficiency of *OpticalDR*'s joint optimization strategy.

Overall, *OpticalDR* has proven to be the most effective solution in terms of privacy preservation and DR. It successfully protects facial privacy while achieving low MAE and RMSE, demonstrating its outstanding performance in DR tasks.

5. Conclusion

We present *OpticalDR*, an innovative privacy-preserving method for DR. Unlike previous image-based approaches, *OpticalDR* prioritizes capturing privacy-preserving images using an optimized lens to preserve depression-related features while concealing identity information. Crucially, *OpticalDR* operates without relying on visually identifiable facial images during system inference, ensuring robust privacy protection. Experimental results showcase *OpticalDR*'s state-of-the-art privacy-preserving performance and competitive DR recognition on established benchmarks. In the future, our efforts will involve the physical deployment of the entire system and the integration of additional modalities, including temporal information, EEG data, and consultation transcripts, all within the framework designed for privacy preservation.

Acknowledgments

The research was supported by the National Natural Science Foundation of China (U23B2009, 92270116).

References

- [1] Sharifa Alghowinem, Roland Goecke, Michael Wagner, Julien Epps, Matthew Hyett, Gordon Parker, and Michael Breakspear. Multimodal depression detection: Fusion analysis of paralinguistic, head pose and eye gaze behaviors. *IEEE Transactions on Affective Computing*, 9(4):478–490, 2018. 1
- [2] Aaron T. Beck, Robert A. Steer, Roberta Ball, and William F. Ranieri. Comparison of beck depression inventories-IA and-II in psychiatric outpatients. *Journal of Personality Assessment*, 67(3):588–597, 1996. 5
- [3] Wheidima Carneiro de Melo, Eric Granger, and Abdenour Hadid. A deep multiscale spatiotemporal network for assessing depression from facial dynamics. *IEEE Transactions on Affective Computing*, pages 1–1, 2020. 2
- [4] Julie Chang, Vincent Sitzmann, Xiong Dun, Wolfgang Heidrich, and Gordon Wetzstein. Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification. *Scientific Reports*, 8(1):12324, Aug 2018. 2
- [5] Julie Chang and Gordon Wetzstein. Deep optics for monocular depth estimation and 3D object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10192–10201, 2019. 2
- [6] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 839–847, 2018. 8
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4685–4694, 2019. 5
- [8] Yizhuo Dong and Xinyu Yang. A hierarchical depression detection model based on vocal and emotional cues. *Neurocomputing*, 441:279–290, 2021. 1
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv e-prints*, page arXiv:2010.11929, Oct. 2020. 8
- [10] Zhengyin Du, Weixin Li, Di Huang, and Yunhong Wang. Encoding visual behaviors with attentive temporal convolution for depression prediction. In *IEEE International Conference on Automatic Face Gesture Recognition (FG)*, pages 1–7, 2019. 1, 2
- [11] Albert Haque, Michelle Guo, Adam S Miner, and Li Fei-Fei. Measuring depression symptom severity from spoken language and 3D facial expressions. *arXiv e-prints*, page arXiv:1811.08592, Nov. 2018. 1
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 4, 5
- [13] Lang He, Dongmei Jiang, and Hichem Sahli. Automatic depression analysis using dynamic facial appearance descriptor and dirichlet process fisher encoding. *IEEE Transactions on Multimedia*, 21(6):1476–1486, 2019. 2
- [14] Chris Henry, M. Salman Asif, and Zhu Li. Privacy preserving face recognition with lensless camera. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023. 2
- [15] Carlos Hinojosa, Miguel Marquez, Henry Arguello, Ehsan Adeli, Li Fei-Fei, and Juan Carlos Niebles. PrivHAR: Recognizing human actions from privacy-preserving lens. In *European Conference on Computer Vision (ECCV)*, page 314–332, 2022. 2
- [16] Carlos Hinojosa, Juan Carlos Niebles, and Henry Arguello. Learning privacy-preserving optics for human pose estimation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2553–2562, 2021. 2
- [17] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-excitation networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(8):2011–2023, 2020. 5
- [18] Hayato Ikoma, Cindy M. Nguyen, Christopher A. Metzler, Yifan Peng, and Gordon Wetzstein. Depth from defocus with learned optics for imaging and occlusion-aware depth estimation. In *IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, 2021. 2
- [19] Yasunori Ishii, Satoshi Sato, and Takayoshi Yamashita. Privacy-aware face recognition with lensless multi-pinhole camera. In *European Conference on Computer Vision (ECCV)*, page 476–493, 2020. 2
- [20] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. DeblurGAN-v2: Deblurring (orders-of-magnitude) faster and better. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8877–8886, 2019. 6
- [21] Zhi Li, Haoliang Li, Kwok-Yan Lam, and Alex Chichung Kot. Unseen face presentation attack detection with hypersphere loss. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2852–2856, 2020. 4
- [22] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3730–3738, 2015. 5
- [23] Patrick Lucey, Jeffrey F. Cohn, Takeo Kanade, Jason Saragih, Zara Ambadar, and Iain Matthews. The extended cohn-kanade dataset (CK+): A complete dataset for action unit and emotion-specified expression. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops (CVPRW)*, pages 94–101, 2010. 5
- [24] Robert J. Noll. Zernike polynomials and atmospheric turbulence*. *Journal of the Optical Society of America*, 66(3):207–211, 1976. 3
- [25] Yuchen Pan, Yuanyuan Shang, Tie Liu, Zhuhong Shao, Guodong Guo, Hui Ding, and Qiang Hu. Spatial-temporal attention network for depression recognition from facial videos. *Expert Systems with Applications*, 237:121410, 2024. 2, 4, 8
- [26] Yuchen Pan, Yuanyuan Shang, Zhuhong Shao, Tie Liu, Guodong Guo, and Hui Ding. Integrating deep facial priors

- into landmarks for privacy preserving multimodal depression recognition. *IEEE Transactions on Affective Computing*, pages 1–8, 2023. 1, 4
- [27] Swati Rathi, Baljeet Kaur, and R. K. Agrawal. Enhanced depression detection from facial cues using univariate feature selection techniques. In *International Conference on Pattern Recognition and Machine Intelligence (PRMI)*, pages 22–29, 2019. 2
- [28] Fabien Ringeval, Björn Schuller, Michel Valstar, Jonathan Gratch, Roddy Cowie, Stefan Scherer, Sharon Mozgai, Nicholas Cummins, Maximilian Schmitt, and Maja Pantic. AVEC 2017: Real-life depression, and affect recognition workshop and challenge. In *Annual Workshop on Audio/Visual Emotion Challenge (AVEC)*, page 3–9, 2017. 2
- [29] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015. 5
- [30] Yuanyuan Shang, Yuchen Pan, Xiao Jiang, Zhong Shao, Guodong Guo, Tie Liu, and Hui Ding. LQGDNet: A local quaternion and global deep network for facial depression recognition. *IEEE Transactions on Affective Computing*, pages 1–1, 2021. 2, 4, 8
- [31] Jian Shen, Yanan Zhang, Huajian Liang, Zeguang Zhao, Kexin Zhu, Kun Qian, Qunxi Dong, Xiaowei Zhang, and Bin Hu. Depression recognition from EEG signals using an adaptive channel fusion method via improved focal loss. *IEEE Journal of Biomedical and Health Informatics*, 27(7):3234–3245, 2023. 1
- [32] Vincent Sitzmann, Steven Diamond, Yifan Peng, Xiong Dun, Stephen Boyd, Wolfgang Heidrich, Felix Heide, and Gordon Wetzstein. End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging. *ACM Transactions on Graphics*, 37(4), jul 2018. 1, 2, 3
- [33] Siyang Song, Shashank Jaiswal, Linlin Shen, and Michel Valstar. Spectral representation of behaviour primitives for depression analysis. *IEEE Transactions on Affective Computing*, pages 1–1, 2020. 2
- [34] Siyang Song, Linlin Shen, and Michel Valstar. Human behaviour-based automatic depression analysis using hand-crafted statistics and deep learned spectral features. In *IEEE International Conference on Automatic Face Gesture Recognition (FG)*, pages 158–165, 2018. 2
- [35] Bo Sun, Yinghui Zhang, Jun He, Lejun Yu, Qihua Xu, Dongliang Li, and Zhaoying Wang. A random forest regression method with selected-text feature for depression assessment. In *Annual Workshop on Audio/Visual Emotion Challenge (AVEC)*, page 61–68, 2017. 2
- [36] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal transformer for unaligned multimodal language sequences. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 6558–6569, 2019. 4
- [37] Md Azher Uddin, Joolekha Bibi Joolee, and Young-Koo Lee. Depression level prediction using deep spatiotemporal features and multilayer bi-lstm. *IEEE Transactions on Affective Computing*, pages 1–1, 2020. 2
- [38] Evgeniya Ustinova and Victor Lempitsky. Learning deep embeddings with histogram loss. In *International Conference on Neural Information Processing Systems (NIPS)*, page 4177–4185, 2016. 5
- [39] Michel Valstar, Björn Schuller, Kirsty Smith, Timur Almaev, Florian Eyben, Jarek Krajewski, Roddy Cowie, and Maja Pantic. AVEC 2014: 3D dimensional affect and depression recognition challenge. In *International Workshop on Audio/Visual Emotion Challenge (AVEC)*, page 3–10, 2014. 2, 5, 8
- [40] Michel Valstar, Björn Schuller, Kirsty Smith, Florian Eyben, Bihan Jiang, Sanjay Bilakhia, Sebastian Schnieder, Roddy Cowie, and Maja Pantic. AVEC 2013: The continuous audio/visual emotion and depression recognition challenge. In *ACM International Workshop on Audio/Visual Emotion Challenge (AVEC)*, page 3–10, 2013. 2, 5, 8
- [41] Ziping Zhao, Zhongtian Bao, Zixing Zhang, Nicholas Cummins, Haishuai Wang, and Björn Schuller. Hierarchical attention transfer networks for depression assessment from speech. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7159–7163, 2020. 1
- [42] Xiuzhuang Zhou, Kai Jin, Yuanyuan Shang, and Guodong Guo. Visually interpretable representation learning for depression recognition from facial images. *IEEE Transactions on Affective Computing*, 11(3):542–552, 2020. 2
- [43] Xiuzhuang Zhou, Zeqiang Wei, Min Xu, Shan Qu, and Guodong Guo. Facial depression recognition by deep joint label distribution and metric learning. *IEEE Transactions on Affective Computing*, pages 1–1, 2020. 2, 5