Teacher-Student Training and Triplet Loss for Facial Expression Recognition under Occlusion

Mariana-Iuliana Georgescu[*,†], Radu Tudor Ionescu[†]
*Novustech Services, 12B Aleea Ilioara, Bucharest, Romania
Email: georgescu_lily@yahoo.com
†University of Bucharest, 14 Academiei, Bucharest, Romania
Email: raducu.ionescu@gmail.com

Abstract—**In this paper, we study the task of facial expression recognition under strong occlusion. We are particularly interested in cases where 50% of the face is occluded, e.g. when the subject wears a Virtual Reality (VR) headset. While previous studies show that pre-training convolutional neural networks (CNNs) on fully-visible (non-occluded) faces improves the accuracy, we propose to employ knowledge distillation to achieve further improvements. First of all, we employ the classic teacher-student training strategy, in which the teacher is a CNN trained on fully-visible faces and the student is a CNN trained on occluded faces. Second of all, we propose a new approach for knowledge distillation based on triplet loss. During training, the goal is to reduce the distance between an anchor embedding, produced by a student CNN that takes occluded faces as input, and a positive embedding (from the same class as the anchor), produced by a teacher CNN trained on fully-visible faces, so that it becomes smaller than the distance between the anchor and a negative embedding (from a different class than the anchor), produced by the student CNN. Third of all, we propose to combine the distilled embeddings obtained through the classic teacher-student strategy and our novel teacher-student strategy based on triplet loss into a single embedding vector.**

We conduct experiments on two benchmarks, FER+ and AffectNet, with two CNN architectures, VGG-f and VGG-face, showing that knowledge distillation can bring significant improvements over the state-of-the-art methods designed for occluded faces in the VR setting. Furthermore, we obtain accuracy rates that are quite close to the state-of-the-art models that take as input fully-visible faces. For example, on the FER+ data set, our VGG-face based on concatenated distilled embeddings attains an accuracy rate of 82.75% on lower-half-visible faces, which is only 2.24% below the accuracy rate of a state-of-the-art VGG-13 that is evaluated on fully-visible faces. Given that our model sees only the lower-half of the face, we consider this to be a remarkable achievement. In conclusion, we consider that our distilled CNN models can provide useful feedback for the task of recognizing the facial expressions of a person wearing a VR headset.

I. INTRODUCTION

Facial expression recognition is an actively studied topic in computer vision, having many practical applications in various domains, such as detection of mental disorders, human behavior understanding and human-computer interaction. In this paper, we focus on the application of facial expression recognition to human-computer interaction in a virtual reality (VR) environment. More specifically, our aim is to design a system able to recognize the facial expressions of a user wearing a VR headset in order to automatically control and adjust the VR environment according to the user's emotions.

To our knowledge, there are only a few approaches in this direction [1], [2], in which the main challenge is to deal with severe occlusions of the face caused by the fact that the user is wearing a VR headset covering the entire upper part of the user's face. Hickson et al. [2] proposed an approach that analyzes expressions from the eyes region captured with a camera mounted inside the VR headset, while Georgescu et al. [1] proposed an approach that analyzes the lower-half of the face captured with a standard camera. We place ourselves in the same setting as Georgescu et al. [1], studying the task of facial expression recognition when the upper-half of the face is occluded. Different from Georgescu et al. [1], we propose to employ knowledge distillation to obtain more accurate convolutional neural networks (CNNs). We study two knowledge distillation approaches in order to distill information from CNNs trained on fully-visible faces to CNNs trained on occluded faces. First of all, we employ the classic teacher-student training strategy [3], [4]. Second of all, we propose a new approach for knowledge distillation based on triplet loss [5]. During training, the goal is to reduce the distance between an anchor embedding, produced by a student CNN that takes occluded faces as input, and a positive embedding (from the same class as the anchor), produced by a teacher CNN trained on fully-visible faces, so that it becomes smaller than the distance between the anchor and a negative embedding (from a different class than the anchor), produced by the student CNN. To our knowledge, we are the first to apply triplet loss in order to distill knowledge into neural networks. Last but not least, we propose to combine the distilled face embeddings obtained through the classic teacher-student strategy and our novel teacher-student strategy based on triplet loss into a single face embedding vector, further boosting the performance gains.

We conduct experiments on two benchmark data sets, FER+ [6] and AffectNet [7], comparing our models based on knowledge distillation with closely-related models [1], [2] designed for the VR setting, as well as with state-of-the-art models [6], [7], [8], [9], [10] that work on fully-visible faces. We note that the latter comparison with works such as [6], [7], [8], [9], [10] is not entirely fair, as these state-of-the-art models get to see the entire (non-occluded) faces at test time. This comparison is rather intended to provide some upper-bounds to the results that could be obtained on occluded faces. We

also note that FER+ and AffectNet do not contain occluded faces, so need to modify the original images by blacking out regions to simulate the occlusions caused by the VR headset.

In all the experiments, the empirical results show that our distilled CNN models obtain superior results compared to the other models [1], [2] evaluated on occluded faces. Furthermore, the accuracy rates of our best models tested on lower-half-visible faces are between 2% and 8% under accuracy rates of the state-of-the-art CNN models, which are tested on fully-visible faces. We consider noteworthy the fact that the gap between our distilled models (evaluated on occluded faces) and the state-of-the-art ones (evaluated on fully-visible faces) is so small.

In summary, our contribution is threefold:

- We propose a novel knowledge distillation method based on triplet loss.
- We propose to combine the classic teacher-student strategy with our strategy based on triplet loss using a late fusion strategy, i.e. by concatenating the distilled face embeddings.
- We conduct experiments on two benchmarks, showing that knowledge distillation brings significant performance improvements in facial expression recognition under strong occlusion.

We organize the rest of this paper as follows. We discuss related work in Section II. We present our knowledge distillation approaches in Section III. We describe the empirical results in Section IV. Finally, we draw our conclusions in Section V.

II. RELATED WORK

A. Facial Expression Recognition

In the past few years, most works on facial expression recognition have focused on building and training deep neural networks in order to obtain state-of-the-art results [6], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29]. Engineered models based on handcrafted features [9], [30], [31], [32] have drawn very little attention, since such models usually yield less accurate results compared to deep learning models. In [6], [8], the authors adopted VGG-like architectures. Barsoum et al. [6] designed a convolutional neural network specifically for the FER+ data set, consisting of 13 layers (VGG-13). Guo et al. [8] focused on detecting the emotion on mobile devices, proposing a light-weight VGG architecture. In order to gain computational performance, they reduced the input size, the number of filters and the number of layers, and replaced the fully-connected layers with global average pooling. Their network consists of 12 layers organized into 6 blocks.

We note that most works studied facial expression recognition from static images, but there are also some works designed for video [14], [33]. Hasani et al. [14] proposed a network architecture that consists of 3D convolutional layers followed by a Long Short-Term Memory (LSTM) network, extracting the spatial relations within facial images and the temporal relations between different frames in the video.

Unlike other approaches, Meng et al. [23] and Liu et al. [21] presented identity-aware facial expression recognition models. Meng et al. [23] proposed to jointly estimate expression and identity features through a neural architecture composed of two identical CNN streams, in order to alleviate inter-subject variations introduced by personal attributes and to achieve better facial expression recognition performance. Liu et al. [21] employed deep metric learning and jointly optimized a deep metric loss and the softmax loss. They obtained an identity-invariant model by using an identity-aware hard-negative mining and online positive mining scheme. Li et al. [20] trained a CNN model using a modified back-propagation algorithm which creates a locality preserving loss aiming to pull the neighboring faces of the same class together. Zeng et al. [29] proposed a model that addresses the labeling inconsistencies across data sets. In their framework, images are tagged with multiple (pseudo) labels either provided by human annotators or predicted by learned models. Then, a facial expression recognition model is trained to fit the latent ground-truth from the inconsistent pseudo-labels. Hua et al. [16] proposed a deep learning algorithm consisting of three sub-networks of different depths. Each sub-network is based on an independently-trained CNN.

Different from all the works mentioned so far and many others [6], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], that recognize facial expressions from fully-visible faces, we focus on recognizing the emotion using only the lower part of the image. The number of works that focus on facial expression recognition under occlusion is considerably smaller [1], [2], [22]. Li et al. [22] trained and tested a model on synthetically occluded images. They proposed an end-to-end trainable Patch-Gated CNN that automatically detects the occluded regions and focuses on the most discriminative non-occluded regions. Different from Li et al. [22], we consider a more difficult setting in which half of the face is completely occluded. In order to learn effectively in this difficult setting, we transfer knowledge from teacher networks that are trained on fully-visible (non-occluded) faces.

Closer to our approach are the works designed for the difficult VR setting [1], [2], in which a VR headset covers the upper side of the face. Hickson et al. [2] proposed a method that analyzes expressions from the eyes region. The eyes region is captured by an infrared camera mounted inside the VR headset, making the method less generic. Georgescu et al. [1] proposed an approach that analyzes the mouth region captured with a standard camera. In this work, we use the same setting as Georgescu et al. [1], studying the task of facial expression recognition when the upper-half of the face is occluded. Different from Georgescu et al. [1], we propose to employ knowledge distillation to obtain more accurate CNNs. We study two knowledge distillation approaches in order to distill information from CNNs trained on fully-visible faces to CNNs trained on occluded faces. To our knowledge, we are the first to apply knowledge distillation in the context of facial expression recognition under strong occlusion.

B. Knowledge Distillation

Knowledge distillation [3], [4] is a recently studied approach [34], [35], [36], [37], [38], [39] that enables the transfer of knowledge between neural networks. Knowledge distillation is a framework that unifies model compression [3], [4], [34] and learning under privileged information [35], [40], the former one being more popular than the latter. In model compression, knowledge from a large neural network [3], [36] or an ensemble of large neural models [4], [38], [39] is distilled into a small neural network, that runs efficiently during inference. In learning under privileged information, knowledge from a neural model training on privileged information (additional data representation not available at test time) is transferred to another neural model that does not have access to the privileged information. In our paper, we are not interested in compressing neural models, but in learning under privileged information. In particular, we study teacher-student training strategies, in which the teacher neural network can learn from fully-visible faces and the student neural network can learn from occluded faces only. In this context, hidden (occluded) face regions represent the privileged information.

To our knowledge, we are the first to propose the distillation of knowledge using triplet loss. We note that there are previous works [34], [36], [38], [39] that distilled triplets or the metric space from a teacher network to a student network. Different from these methods, we do not aim to transfer the metric space learned by a teacher network, but to transfer knowledge from the teacher using metric learning, which is fundamentally different.

III. METHODS

To demonstrate that our knowledge distillation methods generalize across neural architectures, we employ two CNN models, namely VGG-f [41] and VGG-face [42]. We opted for these particular models to allow a direct and fair comparison with Georgescu et al. [1]. We note that VGG-f is pre-trained on object recognition, while VGG-face is pre-trained on face recognition. To obtain the teacher networks, we fine-tune VGG-f and VGG-face on facial expression recognition from fully-visible faces. Similarly, we obtain the student networks by fine-tuning VGG-f and VGG-face on facial expression recognition from lower-half-visible faces. In order to fine-tune the models, we employ the same hyperparameters as Georgescu et al. [1]. Unlike Georgescu et al. [1], we next employ one of two teacher-student training strategies to learn privileged information from a teacher CNN to a student CNN. We note that the two architectures, VGG-f and VGG-face, are never mixed in the training process. In other words, we use the teacher VGG-f to distill knowledge into the student VGG-f, and independently, the teacher VGG-face to distill knowledge into the student VGG-face. Our goal is not to compress the models, but to improve performance using the privileged information accessible to the teacher networks. In our case, the privileged information is represented by the upper-half of the face, which is occluded (not available) at test time.

We next describe in detail the teacher-student strategies for facial expression recognition under strong occlusion.

A. Standard Teacher-Student Training

Ba et al. [3] discussed the idea of model compression in the context of deep learning. Model compression refers to training a compact (shallow) model to approximate the function learned by a more complex (deeper) model. Hinton et al. [4] further developed the idea and proposed to distill the knowledge from an ensemble of models into a single neural network, to achieve faster inference time and improved performance. Hinton et al. [4] suggested that knowledge can be transferred in a straightforward manner by training the distilled (student) network using a soft target distribution that is produced by the ensemble (teacher) model. Improving this approach is possible if the correct labels for the chosen training samples are also known by the student network. In this case, we can employ a weighted average of two different loss functions, thus obtaining a single knowledge distillation (KD) loss function:

$$\mathcal{L}_{KD'}(\theta_S) = (1 - \lambda)\mathcal{L}(y, N_S) + \lambda\mathcal{L}(N_T^\tau, N_S^\tau), \quad (1)$$

where θ_S are the weights of the student network S, y are the target labels, N_T and N_S are the outputs of the teacher network T and the student network S, respectively, N_T^τ and N_S^τ are the softened outputs of T and S, respectively, and $\tau > 1$ is a temperature parameter for the softening operation. We note that N_T^τ and N_S^τ are derived from the pre-softmax activations A_T and A_S of the teacher network and the student network, respectively:

$$N_T^\tau = softmax\left(\frac{A_T}{\tau}\right), N_S^\tau = softmax\left(\frac{A_S}{\tau}\right). \quad (2)$$

In Equation (1), the first loss function $\mathcal{L}(y, N_S)$ is the cross-entropy with respect to the correct labels. The second loss function $\mathcal{L}(N_T^\tau, N_S^\tau)$ is the cross-entropy with respect to the soft targets provided by the teacher network. Hinton et al. [4] suggested that the second objective function must have a higher weight with respect to the first objective function in Equation (1). Therefore, we set $\lambda = 0.7$ for both VGG-f and VGG-face. We also set the same temperature ($\tau = 20$) for both networks.

Lopez-Paz et al. [35] proposed a generalized knowledge distillation paradigm that allows to learn not only from a different (teacher) network, but also by using a different data representation. Indeed, the teacher-student paradigm presented in Equation (1) can also be applied to learn privileged information [40] that is available only to the teacher network. This approach is suitable in our setting that involves learning from severely occluded faces, as we aim to learn from a teacher network that has access to privileged information, i.e. non-occluded faces. We therefore choose as teacher networks the CNN models that are trained on fullly-visible face images. We stress out that VGG-f and VGG-face are used both as teacher networks and as student networks, but without mixing the architectures, i.e. the teacher and the student networks always have the same architecture. The fact that the teacher's

Fig. 1. The standard teacher-student training pipeline for facial expression recognition on severely occluded faces. The teacher CNN takes as input non-occluded (fully-visible) faces, having access to privileged information. The student CNN takes as input only occluded (lower-half-visible) faces, but learns useful information from the teacher CNN model. The loss functions $\mathcal{L}(y, N_S)$ and $\mathcal{L}(N_T^\tau, N_S^\tau)$ are defined in Equation (1). Best viewed in color.

architecture is not deeper (as commonly agreed in [3], [4]) is less important in our context. What matters most is that the teacher network is trained on non-occluded faces and its soft labels provide additional information that is not seen by the student network in the input images, since the faces are occluded. Our first teacher-student framework is illustrated in Figure 1.

B. Teacher-Student Training with Triplet Loss

In standard teacher-student training, the goal is to train the student network to reproduce the output class probabilities of the teacher network, in addition to predicting the correct labels. We believe that the same effect can be achieved by training the student network to reproduce the face embeddings (the activations from the penultimate layer) of the teacher network. However, we do not necessarily want to obtain similar embeddings when the student network takes an input image from a different class than the teacher network. In other words, it is sufficient to ensure that the embeddings produced by the student and the teacher are similar only when the input examples belong to the same class. In order to model these constraints properly, we propose to employ triplet loss [5]. Triplet loss has been previously used [5] to obtain close embeddings for objects belonging to the same class and distant embeddings for objects belonging to different classes, while ignoring other factors of variation such as pose, illumination and rotation. We hereby explain in detail how triplet loss can be adapted and used to distill knowledge from a teacher neural network to a student neural network.

In the following, we use the prime symbol to distinguish between fully-visible and occluded images. For a fully-visible input image x and an occluded input image x', let $E_T(x)$ and $E_S(x')$ denote the embedding vectors (face embeddings or embeddings) produced by the teacher network T and the student network S, respectively. In order to distill knowledge using triplet loss, we need triplets of input images of the form (a', p, n'), where a' is an occluded image from a class k, called anchor example, p is a fully-visible image from the same class k as the anchor, called positive example, and n' is another

occluded image from a class $j \neq k$, called negative example. During training, our goal is to reduce the distance between the anchor embedding $E_S(a')$ and the positive embedding $E_T(p)$, so that it becomes smaller than the distance between the anchor embedding $E_S(a')$ and the negative embedding $E_S(n')$. We thus optimize the following triplet loss function:

$$
\mathcal{L}_{triplet}(\theta_S) = \sum_{i=1}^{m} \Big[\Delta(E_S(a'_i), E_T(p_i)) - \\
- \Delta(E_S(a'_i), E_S(n'_i)) + \alpha \Big]_+ ,
\tag{3}
$$

where θ_S are the weights of the student network S, m is the number of training examples, Δ is the Euclidean distance, $[\cdot]_+$ is equal to $max(0, \cdot)$ and α is a margin that is enforced between the positive pair (a'_i, p_i) and the negative pair (a'_i, n'_i). In the experiments, we set the margin α to 0.1.

In a similar fashion to the standard teacher-student loss from Equation (1), we also want to make sure that the student network predicts the labels correctly. Therefore, our final loss function is:

$$
\mathcal{L}_{KD''}(\theta_S) = \mathcal{L}(y, N_S) + \lambda \mathcal{L}_{triplet}(\theta_S),
\tag{4}
$$

where θ_S are the weights of the student network S, $\mathcal{L}(y, N_S)$ is the cross-entropy with respect to the correct labels y, $\mathcal{L}_{triplet}$ is the triplet loss defined in Equation (3) and λ is a parameter that controls the importance of the triplet loss with respect to the cross-entropy loss. In the experiments, we set λ to 0.5.

For an efficient training process, we generate the triplets as follows. Each occluded image from the training set is selected as the anchor a'. For each anchor, we sample a random subset of fully-visible images belonging to the same class as the anchor, we compute the distances between the anchor and each image in the resulting subset and we select the image that is farthest from the anchor as the positive example p. In a similar fashion, we sample a random subset of occluded images belonging to different classes than the anchor, we compute the distances between the anchor and each image in the resulting subset and we select the image that is closest
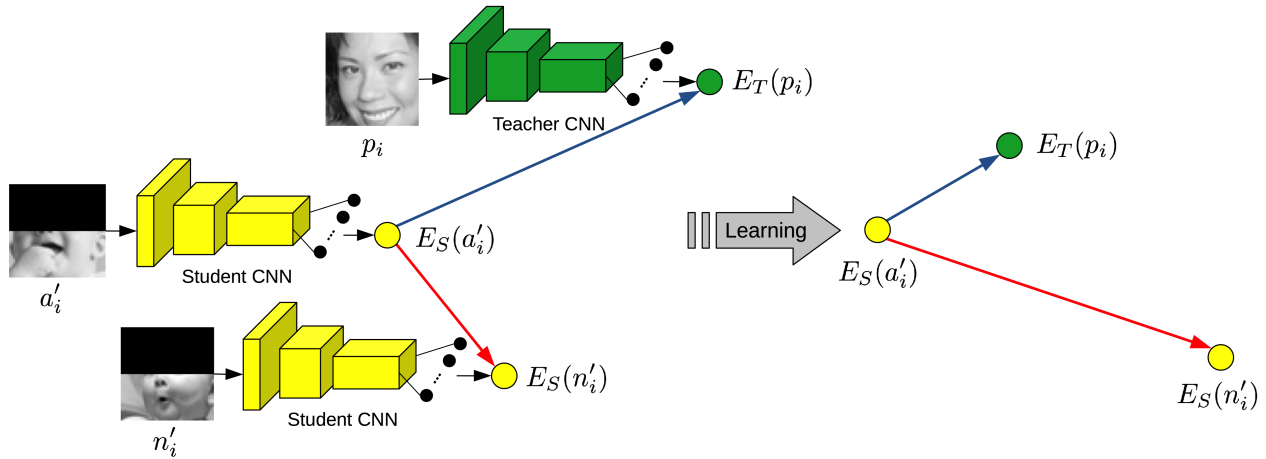
Fig. 2. The teacher-student training based on triplet loss for facial expression recognition on severely occluded faces. During training, we modify the weights of the student network such that the distance $\Delta(E_S(a_i'), E_T(p_i))$ becomes smaller than the distance $\Delta(E_S(a_i'), E_S(n_i'))$. Best viewed in color.

to the anchor as the negative example n'. In both cases, the random subsets represent only one tenth of the full training set, reducing the processing time by an order of magnitude. The random subsets are sampled once per epoch. The effect of training using the defined triplet loss is illustrated in Figure 2. We note that only the weights θ_S of the student model are updated during training, hence the embedding $E_T(p_i)$ cannot be modified.

IV. EXPERIMENTS

A. Data Sets

FER+. The FER+ data set [6] is a newer version of the FER 2013 data set [43] which contains images with wrong labels as well as images not containing faces. Barsoum et al. [6] cleaned up the FER 2013 data set by relabeling images and by removing those without faces. In the relabeling process, Barsoum et al. [6] added a new class of emotion, *contempt*, while also keeping the other 7 classes from FER 2013: *anger*, *disgust*, *fear*, *happiness*, *neutral*, *sadness*, *surprise*. The FER+ data set is composed of 25045 training images, 3191 validation images and 3137 test images. The size of each image is 48×48 pixels.

AffectNet. The AffectNet [7] data set is one of the largest data sets for facial expression recognition, containing 287651 training images and 4000 validation images with manual annotations. The images from AffectNet have various sizes. Since the test set is not yet publicly available, methods [7], [29] are commonly evaluated on the validation set. The data set contains the same 8 classes of emotion as FER+. With 500 images per class in the validation set, the class distribution is balanced. In the same time, the training data is unbalanced. As proposed by Mollahosseini et al. [7], we down-sample the training set for classes with more than 15000 images. This leaves us with a training set of 88021 images.

B. Data Preprocessing

In order to train and evaluate neural models in the VR setting, we replace the values of occluded pixels with zero. For the setting proposed by Georgescu et al. [1], in which facial expressions are recognized from the lower-half of the face, we occlude the entire upper-half of the FER+ and AffectNet images. For the setting proposed by Hickson et al. [2], in which facial expressions are recognized from the eyes region, we occlude the entire lower-half of the FER+ and AffectNet images. Our own approaches are trained and tested in the former setting, in which the upper-half is occluded. All images are resized to 224×224 pixels, irrespective of the data set, in order to be given as input to VGG-f or VGG-face.

C. Implementation Details

All neural models are trained with stochastic gradient descent with momentum. We set the momentum rate to 0.9. The VGG-face model is trained on mini-batches of 64 images, while the VGG-f model is trained on mini-batches of 512 images. We use the same mini-batch sizes in all training stages. **Preliminary training of teachers and students.** For the preliminary fine-tuning of the teacher and the student models, we use the MatConvNet [44] library. The teacher VGG-face is fine-tuned on facial expression recognition from fully-visible faces for a total of 50 epochs. The teacher VGG-f is fine-tuned for 800 epochs. We note that VGG-f requires more training time to converge because it is pre-trained on a more distant task, object recognition, unlike VGG-face, which is pre-trained on face recognition. The student VGG-face is fine-tuned on facial expression recognition from occluded faces for 40 epochs. Similarly, the student VGG-f is fine-tuned for 80 epochs. Further details about training VGG-face and VGG-f on fully-visible or occluded faces are provided in [1]. **Standard teacher-student training.** For the standard teacher-student strategy, the student VGG-face is trained for 50 epochs starting with a learning rate of 10^{-4}, decreasing it when the validation error does not improve for 10 consecutive epochs.

By the end of the training process, the learning rate for the student VGG-face drops to 10^{-5}. In a similar manner, the student VGG-f is trained for 200 epochs starting with a learning rate of 10^{-3}, decreasing it when the validation error does not improve for 10 consecutive epochs. By the end of the training process, the learning rate for the student VGG-f drops to 10^{-4}.

Teacher-student training with triplet loss. In order to apply the teacher-student training based on triplet loss, we switch to TensorFlow [45], exporting the VGG-face and VGG-f models from MatConvNet. We train the student VGG-face for 10 epochs using a learning rate of 10^{-6}. In a similar fashion, we train the student VGG-f for 10 epochs using a learning rate of 10^{-5}.

Combining distilled embeddings. After training the student models using the two teacher-student strategies independently, we combine the corresponding embeddings into a single embedding vector through concatenation. The concatenated embeddings are provided as input to a linear Support Vector Machines (SVM) model [46]. We set the regularization parameter of the SVM to $C = 1$. We use the SVM implementation from Scikit-learn [47].

D. Baselines

As baselines, we consider two state-of-the-art methods [1], [2] designed for facial expression recognition in the VR setting. The key contribution of these methods resides in the region they use to extract features, the lower-half of the face (mouth region) [1] or the upper-half of the face (eyes region) [2]. In order to conduct a fair comparison, we use the same neural architectures for both baselines and our approach.

As reference, we include some results on FER+ and AffectNet from the recent literature [6], [7], [8], [9], [10]. We note that these state-of-the-art methods [6], [7], [8], [9], [10] are trained and tested on fully-visible faces. Hence, the comparison to our approach or other approaches applied on occluded faces [1], [2] is unfair, but we included it as a relevant indicator of the upper bound for the models applied on occluded faces.

E. Results

In Table I, we present the empirical results obtained on AffectNet [7] and FER+ [6] by the VGG-f and VGG-face models based on our teacher-student training strategies in comparison with the results obtained by the state-of-the-art methods [6], [7], [8], [9], [10] tested on fully-visible faces and by the methods [1], [2] designed for the VR setting (tested on occluded faces). First of all, we note that it is natural for the state-of-the-art methods [6], [7], [8], [9], [10] to achieve better accuracy rates than our approach or the other approaches applied on occluded faces [1], [2]. An exceptional case is represented by the VGG-face model of Georgescu et al. [1] and our student VGG-face models, which surpass the bag-of-visual-words model [9] on both data sets.

With respect to the baselines [1], [2] designed for the VR setting, all out teacher-student training strategies provide

Model	Test faces	AffectNet	FER+
Bag-of-visual-words [9]	○	48.30%	80.65%
VGG-13 [6]	○	-	84.99%
AlexNet [7]	○	58.00%	-
MT-VGG [10]	○	54.00%	-
VGG-12 [8]	○	58.50%	-
Teacher VGG-f	○	57.37%	85.05%
Teacher VGG-face	○	59.03%	84.79%
Teacher VGG-f	◒	41.58%	70.00%
Teacher VGG-face	◒	37.70%	68.89%
Teacher VGG-f	◓	26.85%	40.07%
Teacher VGG-face	◓	31.23%	48.29%
VGG-f [1]	◒	47.58%	78.23%
VGG-face [1]	◒	49.23%	81.28%
VGG-f [2]	◓	42.45%	66.18%
VGG-face [2]	◓	43.18%	70.19%
VGG-f (standard T-S)	◒	48.75%†	80.17%†
VGG-face (standard T-S)	◒	49.75%	82.37%
VGG-f (triplet loss T-S)	◒	48.13%	80.05%†
VGG-face (triplet loss T-S)	◒	49.71%	82.57%
VGG-f (triplet loss + standard T-S)	◒	48.70%†	81.09%†
VGG-face (triplet loss + standard T-S)	◒	50.09%†	82.75%†

superior results. We observe that the accuracy rates of Hickson et al. [2] are considerably lower than the accuracy rates of Georgescu et al. [1] (differences are between 5% and 12%), although the neural models have identical architectures. We hypothesize that this difference is caused by the fact that it is significantly harder to recognize facial expressions from the eyes region (denoted by ◓) than from the mouth region (denoted by ◒). To test this hypothesis, we evaluate the teacher VGG-f and the teacher VGG-face on upper-half-visible and lower-half-visible faces. We observe even larger differences between the results on upper-half-visible faces (accuracy rates are between 26% and 49%) and the results on lower-half-visible faces (accuracy rates are between 37% and 70%), confirming our hypothesis. We also note that the results attained by the teacher models on occluded faces are much lower than the results of the baselines [1], [2] designed for the VR setting, although the teacher models attain results on par with the state-of-the-art methods [6], [7], [8], [10] when testing is performed on fully-visible faces. This indicates that CNN models trained on fully-visible faces are not particularly suitable to handle severe facial occlusions, justifying the need for training on occluded faces.

The results presented in the last six rows of Table I indicate that the teacher-student learning strategies provide very good results on lower-half-visible faces, surpassing the
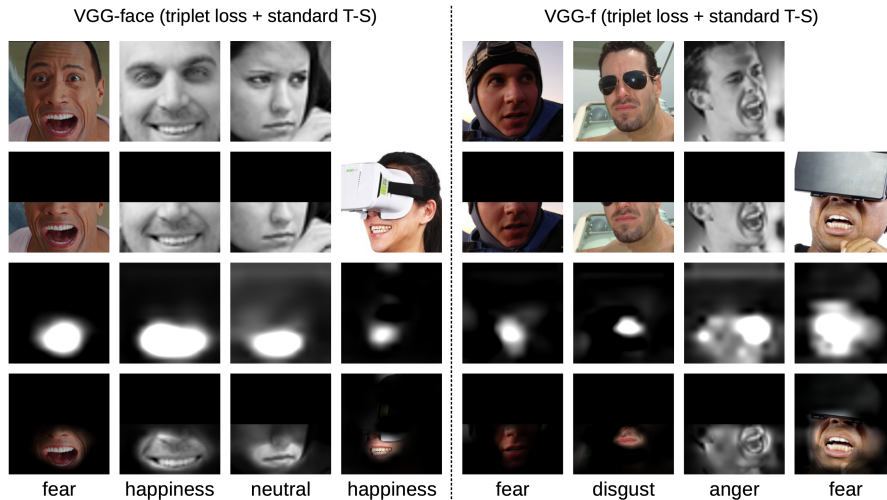
Fig. 3. Fully-visible images (○) on top row, lower-half-visible faces (◑) on second row, Grad-CAM [49] explanation masks on third row and lower-half-visible faces with superimposed Grad-CAM masks on bottom row. The predicted labels provided by the VGG-face (left-hand side) or the VGG-f (right-hand side) models based on the combined teacher-student training strategies are also provided at the bottom. The first three examples from each side are selected from AffectNet [7] and FER+ [6]. The fourth example from each side is a person wearing an actual VR headset. Best viewed in color.

other methods [1], [2] evaluated on occluded faces. We believe that the accuracy gains are due to the teacher neural networks that are trained on fully-visible images, which bring additional (privileged) information from the (unseen) upper half of the training faces. Our teacher-student training strategy based on triplet loss provides results that are comparable to the standard teacher-student training strategy. Nevertheless, we achieve additional performance gains when the two teacher-student strategies are combined through embedding concatenation. Our final models based on the concatenated distilled embeddings attain results that are close to the state-of-the-art methods [6], [7], [8], [10]. For example, our VGG-face with triplet loss and standard teach-student training yields an accuracy rate of 82.75% on FER+, which is 2.24% under the state-of-the-art VGG-13 [6]. We thus conclude that our models can recognize facial expressions with sufficient reliability, despite being tested on faces that are severely occluded (the entire upper-half is occluded).

We also performed statistical significance testing to compare our VGG-f and VGG-face models based on teacher-student training with the VGG-f and VGG-face models of Georgescu et al. [1]. We note that the combined teacher-student strategies provide significant improvements for both models on both data sets, with a significance level of 0.05.

In order to better understand how our models based on concatenated distilled embeddings make decisions, we used the Grad-CAM [49] approach to provide visual explanations for some image examples illustrated in Figure 3. First, we notice that we, as humans, are still able to recognize the facial expressions in the presented examples, even if the upper half of the faces depicted in the second row of Figure 3 are occluded. We observe that both neural architectures focus their attention on the lower part of the face, particularly on the mouth region. This indicates that our neural networks can properly handle

situations in which people wear VR headsets. Notably, all the predicted labels for the samples presented in Figure 3 are correct.

V. CONCLUSION

In this paper, we proposed to train neural networks for facial expression recognition under strong occlusion, by applying two teacher-student strategies. To our knowledge, this is a novel application of teacher-student training. Another novel contribution is the second teacher-student strategy, which is based on triplet loss. The empirical results indicate that our teacher-student strategies can bring significant performance gains over the baselines [1], [2], particularly when the strategies are combined by concatenating the corresponding distilled embeddings. Notably, on the FER+ data set, our VGG-face based on concatenated distilled embeddings attains an accuracy rate of 82.75% on lower-half-visible faces, which is only 2.24% below the accuracy rate of the state-of-the-art VGG-13 [6] on fully-visible faces. In conclusion, we consider that our results are high enough to deploy our neural models in production.

REFERENCES

[1] M.-I. Georgescu and R. T. Ionescu, "Recognizing Facial Expressions of Occluded Faces Using Convolutional Neural Networks," in *Proceedings of ICONIP*, 2019, pp. 645–653.

[2] S. Hickson, N. Dufour, A. Sud, V. Kwatra, and I. Essa, "Eyemotion: Classifying facial expressions in VR using eye-tracking cameras," in *Proceedings of WACV*, 2019, pp. 1626–1635.

[3] J. Ba and R. Caruana, "Do deep nets really need to be deep?" in *Proceedings of NIPS*, 2014, pp. 2654–2662.

[4] G. Hinton, O. Vinyals, and J. Dean, "Distilling the Knowledge in a Neural Network," in *Proceedings of NIPS Deep Learning and Representation Learning Workshop*, 2014.

[5] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proceedings of CVPR*, 2015, pp. 815–823.

[6] E. Barsoum, C. Zhang, C. C. Ferrer, and Z. Zhang, "Training deep networks for facial expression recognition with crowd-sourced label distribution," in *Proceedings of ICMI*, 2016, pp. 279–283.

[7] A. Mollahosseini, B. Hasani, and M. H. Mahoor, "AffectNet: A Database for Facial Expression, Valence, and Arousal Computing in the Wild," *IEEE Transactions on Affective Computing*, vol. 10, no. 1, pp. 18–31, 2019.

[8] Y. Guo, Y. Xia, J. Wang, H. Yu, and R.-C. Chen, "Real-time facial affective computing on mobile devices," *Sensors*, vol. 20, no. 3, p. 870, 2020.

[9] R. T. Ionescu, M. Popescu, and C. Grozea, "Local Learning to Improve Bag of Visual Words Model for Facial Expression Recognition," in *Proceedings of ICML Workshop on Challenges in Representation Learning*, 2013.

[10] D. Kollias and S. Zafeiriou, "Expression, Affect, Action Unit Recognition: Aff-Wild2, Multi-Task Learning and ArcFace," in *Proceedings of BMVC*, 2019.

[11] H. Ding, S. K. Zhou, and R. Chellappa, "FaceNet2ExpNet: Regularizing a Deep Face Recognition Net for Expression Recognition," in *Proceedings of FG*, 2017, pp. 118–126.

[12] M.-I. Georgescu, R. T. Ionescu, and M. Popescu, "Local learning with deep and handcrafted features for facial expression recognition," *IEEE Access*, vol. 7, pp. 64 827–64 836, 2019.

[13] P. Giannopoulos, I. Perikos, and I. Hatzilygeroudis, "Deep Learning Approaches for Facial Emotion Recognition: A Case Study on FER-2013," in *Advances in Hybridization of Intelligent Methods*. Springer, 2018, pp. 1–16.

[14] B. Hasani and M. H. Mahoor, "Facial expression recognition using enhanced deep 3D convolutional neural networks," in *Proceedings of CVPRW*, 2017, pp. 2278–2288.

[15] S. Hosseini and N. I. Cho, "GF-CapsNet: Using Gabor Jet and Capsule Networks for Facial Age, Gender, and Expression Recognition," in *Proceedings of FG*, 2019.

[16] W. Hua, F. Dai, L. Huang, J. Xiong, and G. Gui, "HERO: Human Emotions Recognition for Realizing Intelligent Internet of Things," *IEEE Access*, vol. 7, pp. 24 321–24 332, 2019.

[17] J. Hu, B. Yu, Y. Yang, and B. Feng, "Towards Facial De-Expression and Expression Recognition in the Wild," in *Proceedings of ACII*, 2019, pp. 157–163.

[18] B.-K. Kim, J. Roh, S.-Y. Dong, and S.-Y. Lee, "Hierarchical committee of deep convolutional neural networks for robust facial expression recognition," *Journal on Multimodal User Interfaces*, vol. 10, no. 2, pp. 173–189, 2016.

[19] D. Li and G. Wen, "MRMR-based ensemble pruning for facial expression recognition," *Multimedia Tools and Applications*, pp. 1–22, 2017.

[20] S. Li, W. Deng, and J. Du, "Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild," in *Proceedings of CVPR*, 2017, pp. 2584–2593.

[21] X. Liu, B. Kumar, J. You, and P. Jia, "Adaptive deep metric learning for identity-aware facial expression recognition," in *Proceedings of CVPRW*, 2017, pp. 522–531.

[22] Y. Li, J. Zeng, S. Shan, and X. Chen, "Patch-Gated CNN for occlusion-aware facial expression recognition," in *Proceedings of ICPR*, 2018, pp. 2209–2214.

[23] Z. Meng, P. Liu, J. Cai, S. Han, and Y. Tong, "Identity-aware convolutional neural network for facial expression recognition," in *Proceedings of FG*, 2017, pp. 558–565.

[24] A. Mollahosseini, B. Hassani, M. J. Salvador, H. Abdollahi, D. Chan, and M. H. Mahoor, "Facial expression recognition from World Wild Web," in *Proceedings of CVPRW*, 2016, pp. 1509–1516.

[25] Y. Tang, "Deep Learning using Linear Support Vector Machines," in *Proceedings of ICML Workshop on Challenges in Representation Learning*, 2013.

[26] G. Wen, Z. Hou, H. Li, D. Li, L. Jiang, and E. Xun, "Ensemble of deep neural networks with probability-based fusion for facial expression recognition," *Cognitive Computation*, vol. 9, no. 5, pp. 597–610, 2017.

[27] A. Wikanningrum, R. F. Rachmadi, and K. Ogata, "Improving Lightweight Convolutional Neural Network for Facial Expression Recognition via Transfer Learning," in *Proceedings of CENIM*, 2019, pp. 1–6.

[28] Z. Yu and C. Zhang, "Image based static facial expression recognition with multiple deep network learning," in *Proceedings of ICMI*, 2015, pp. 435–442.

[29] J. Zeng, S. Shan, and X. Chen, "Facial expression recognition with inconsistently annotated datasets," in *Proceedings of ECCV*, 2018, pp. 222–237.

[30] D. Al Chanti and A. Caplier, "Improving Bag-of-Visual-Words Towards Effective Facial Expressive Image Classification," in *Proceedings of VISIGRAPP*, 2018, pp. 145–152.

[31] J. H. Shah, M. Sharif, M. Yasmin, and S. L. Fernandes, "Facial expressions classification and false label reduction using LDA and threefold SVM," *Pattern Recognition Letters*, 2017.

[32] J. Shao, I. Gori, S. Wan, and J. Aggarwal, "3D dynamic facial expression recognition using low-resolution videos," *Pattern Recognition Letters*, vol. 65, pp. 157–162, 2015.

[33] H. Kaya, F. Gürpınar, and A. A. Salah, "Video-based emotion recognition in the wild using deep transfer learning and score fusion," *Image and Vision Computing*, vol. 65, pp. 66–75, 2017.

[34] Y. Feng, H. Wang, R. Hu, and D. T. Yi, "Triplet distillation for deep face recognition," *arXiv preprint arXiv:1905.04457*, 2019.

[35] D. Lopez-Paz, L. Bottou, B. Schölkopf, and V. Vapnik, "Unifying distillation and privileged information," in *Proceedings of ICLR*, 2016.

[36] W. Park, D. Kim, Y. Lu, and M. Cho, "Relational Knowledge Distillation," in *Proceedings of CVPR*, 2019, pp. 3962–3971.

[37] J. Yim, D. Joo, J. Bae, and J. Kim, "A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning," in *Proceedings of CVPR*, 2017, pp. 7130–7138.

[38] S. You, C. Xu, C. Xu, and D. Tao, "Learning from multiple teacher networks," in *Proceedings of KDD*, 2017, pp. 1285–1294.

[39] L. Yu, V. O. Yazici, X. Liu, J. v. d. Weijer, Y. Cheng, and A. Ramisa, "Learning Metrics from Teachers: Compact Networks for Image Embedding," in *Proceedings of CVPR*, 2019, pp. 2907–2916.

[40] V. Vapnik and A. Vashist, "A new learning paradigm: Learning using privileged information," *Neural Networks*, vol. 22, no. 5–6, pp. 544–557, 2009.

[41] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the Devil in the Details: Delving Deep into Convolutional Nets," in *Proceedings of BMVC*, 2014, pp. 1–12.

[42] O. M. Parkhi, A. Vedaldi, A. Zisserman *et al.*, "Deep Face Recognition," in *Proceedings of BMVC*, 2015, pp. 6–17.

[43] I. J. Goodfellow, D. Erhan, P. L. Carrier, A. Courville, M. Mirza, B. Hamner, W. Cukierski, Y. Tang, D. Thaler, D.-H. Lee, Y. Zhou, C. Ramaiah, F. Feng, R. Li, X. Wang, D. Athanasakis, J. Shawe-Taylor, M. Milakov, J. Park, R. T. Ionescu, M. Popescu, C. Grozea, J. Bergstra, J. Xie, L. Romaszko, B. Xu, Z. Chuang, and Y. Bengio, "Challenges in Representation Learning: A report on three machine learning contests," in *Proceedings of ICONIP*, vol. 8228, 2013, pp. 117–124.

[44] A. Vedaldi and K. Lenc, "MatConvNet – Convolutional Neural Networks for MATLAB," in *Proceeding of ACMMM*, 2015.

[45] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: A system for large-scale machine learning," in *Proceedings of OSDI*, 2016, pp. 265–283.

[46] C. Cortes and V. Vapnik, "Support-Vector Networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[47] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[48] T. G. Dietterich, "Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms," *Neural Computation*, vol. 10, no. 7, pp. 1895–1923, 1998.

[49] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," in *Proceedings of ICCV*, 2017, pp. 618–626.