



Teacher–student training and triplet loss to reduce the effect of drastic face occlusion

Application to emotion recognition, gender identification and age estimation

Mariana-Iuliana Georgescu^{1,2} · Georgian-Emilian Duță^{1,2} · Radu Tudor Ionescu^{1,3}

Received: 15 April 2021 / Revised: 13 August 2021 / Accepted: 19 November 2021 / Published online: 22 December 2021
© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2021

Abstract

We study a series of recognition tasks in two realistic scenarios requiring the analysis of faces under strong occlusion. On the one hand, we aim to recognize facial expressions of people wearing virtual reality headsets. On the other hand, we aim to estimate the age and identify the gender of people wearing surgical masks. For all these tasks, the common ground is that half of the face is occluded. In this challenging setting, we show that convolutional neural networks trained on fully visible faces exhibit very low performance levels. While fine-tuning the deep learning models on occluded faces is extremely useful, we show that additional performance gains can be obtained by distilling knowledge from models trained on fully visible faces. To this end, we study two knowledge distillation methods, one based on teacher–student training and one based on triplet loss. Our main contribution consists in a novel approach for knowledge distillation based on triplet loss, which generalizes across models and tasks. Furthermore, we consider combining distilled models learned through conventional teacher–student training or through our novel teacher–student training based on triplet loss. We provide empirical evidence showing that, in most cases, both individual and combined knowledge distillation methods bring statistically significant performance improvements. We conduct experiments with three different neural models (VGG-f, VGG-face and ResNet-50) on various tasks (facial expression recognition, gender recognition, age estimation), showing consistent improvements regardless of the model or task.

This work was supported by a grant of the Romanian Ministry of Education and Research, CNCS - UEFISCDI, project number PN-III-P1-1.1-TE-2019-0235, within PNCDI III. This article has also benefited from the support of the Romanian Young Academy, which is funded by Stiftung Mercator and the Alexander von Humboldt Foundation for the period 2020–2022.

✉ Radu Tudor Ionescu
raducu.ionescu@gmail.com
Mariana-Iuliana Georgescu
georgescu_lily@yahoo.com
Georgian-Emilian Duță
georgian.duta@securifai.ro

- ¹ SecurifAI, Bd. Mircea Vodă 21D, Bucharest, Romania
- ² Department of Computer Science, University of Bucharest, 14 Academiei, Bucharest, Romania
- ³ Department of Computer Science and Romanian Young Academy, University of Bucharest, 14 Academiei, Bucharest, Romania

1 Introduction

We aim to study and develop a generic framework suitable for solving various recognition tasks requiring the analysis of faces under strong occlusion. We underline that the studied framework could be useful in several realistic scenarios. In this work, we focus on two specific yet diverse scenarios to demonstrate the practical applicability of our framework. Our first scenario is related to the recognition of facial expressions of people wearing virtual reality (VR) headsets. A system able to solve this task with high accuracy provides the means to control and change the VR environment with respect to the user's emotions, in real time. This could be useful for adjusting the level of exposure for VR applications designed for the treatment of various types of phobia. Our second scenario stems from the regulations imposed by many countries around the world to minimize the spread of the SARS-CoV-2 virus, requiring people in public indoor and even outdoor environments to wear surgical masks. In the context of the COVID-19 pandemic, estimating the age and identifying the

gender of customers wearing surgical masks is very useful to generate customer demographics for retail stores and supermarkets. Such demographics are necessary for businesses to estimate the impact of advertisement campaigns or to create strategic plans with respect to current trends in customer demand.

In the scenarios enumerated above, the common denominator is the fact that the automated analysis needs to be performed on faces with an occlusion rate of about 50%, i.e., either the upper half or the lower half of the face is occluded. We consider this level of occlusion as *drastic*, given the significant performance damage implied when a deep convolutional neural network (CNN) trained on completely visible faces is applied on half-visible faces. Perhaps the most natural solution to close the performance gap is to fine-tune the model on occluded faces. Yet, we conjecture that the performance gap can be reduced even further by distilling knowledge from a teacher model trained on fully visible faces into a student model fine-tuned on partially visible faces. In this work, we study a conventional knowledge distillation technique based on teacher–student training [30,59] as well as a newly developed technique based on triplet loss [17]. In our novel knowledge distillation framework, we formulate the objective such that the model reduces 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. In addition, we consider combining distilled student models, learned with conventional teacher–student training or triplet loss, into ensemble models.

We conduct experiments on multiple tasks and data sets to demonstrate the generality of using knowledge distillation to reduce the negative effect of face occlusion on accuracy. More precisely, we present facial expression recognition results on AffectNet [51] and FER+ [5], gender identification results on UTKFace [83] and age estimation results on UTKFace. To simulate occlusions, we replace the pixel values in the upper or lower half of an image with zeros, blacking out the corresponding region as necessary. To demonstrate the generality across models, we consider two VGG architectures, namely VGG-f [7] and VGG-face [55], and one residual architecture, namely ResNet-50 [28]. Our empirical results indicate that knowledge distillation yields superior results across the investigated tasks and models.

With respect to our preliminary work [17], we make the following contributions:

- We provide a more comprehensive description of the proposed methods.

- We demonstrate the generality across multiple tasks, adding gender recognition and age estimation as new tasks.
- We demonstrate the generality across multiple models, considering both residual and VGG-like architectures.

The rest of the paper is organized as follows. Section 2 contains an overview of the related scientific articles. The investigated methods are detailed in Sect. 3. The experiments and results are described in Sect. 4. Our conclusions and future work directions are provided in Sect. 5.

2 Related work

2.1 Facial expression recognition

Over the past few years, the research efforts on facial expression recognition have concentrated on building and training deep neural networks aimed at obtaining state-of-the-art results [11,18,20,27,31,34,40,42,44–47,50,52,64–67,70,71,74]. Engineered models based on handcrafted features [3,35,62,63] have attracted comparably less attention, since such models typically attain less accurate results than deep learning models. In works such as [5,25], the authors adopted VGG-like architectures. Barsoum et al. [5] proposed a CNN particularly for the FER+ benchmark, formed of 13 layers (VGG-13). Guo et al. [25] concentrated on recognizing facial expressions on mobile devices, designing a lightweight VGG architecture. In order to minimize computational costs, the authors reduced the input size, the number of filters and the number of layers and replaced the fully connected layers with global average pooling. Their architecture is formed of 12 layers divided into 6 blocks.

Although the majority of works studied facial expression recognition from static images, there is a body of works focusing on video [27,39]. Hasani et al. [27] designed a neural architecture that is formed of 3D convolutional layers followed by a long short-term memory (LSTM) network, extracting spatial relations within facial images and temporal relations between different frames in a video.

In a different direction from the aforementioned methods, Meng et al. [50] and Liu et al. [47] presented identity-aware facial expression recognition models. Meng et al. [50] suggested to jointly estimate expression and identity features using a neural architecture consisting of two identical CNN streams, aiming to alleviate inter-subject variations introduced by personal attributes and attain superior facial expression recognition performance. Liu et al. [47] considered deep metric learning, jointly optimizing a deep metric loss and the softmax loss. They obtained an identity-invariant model by using a scheme based on identity-aware hard-negative mining and online positive mining. Li et al. [45]

optimized a CNN architecture using an upgraded back-propagation algorithm that uses a locality preserving loss aiming to pull the neighboring faces from the same class together. Zeng et al. [81] designed a model that addresses the labeling inconsistencies across data sets. In their approach, images are annotated with multiple pseudo-labels, either given by human annotators or predicted by trained models. Then, a facial expression recognition model is trained to fit the latent ground truth from the inconsistent pseudo-labels. Hua et al. [34] presented 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 [3,5,11,18,20,27,31,33–35,40,42,44–47,50,52,62–67,70,71,73,74,80,81], which recognize facial expressions from fully visible faces, we concentrate on recognizing the emotion by looking only at the lower part of the face. The number of works that focus on facial expression recognition under occlusion is considerably smaller [16,29,46]. Li et al. [46] applied a model on synthetically occluded images. They designed an end-to-end trainable Patch-Gated CNN to automatically detect the occluded regions and concentrate on the most discriminative non-occluded regions. Unlike Li et al. [46], we consider a more difficult scenario in which half of the face is completely occluded. In order to learn effectively in this difficult scenario, we propose to transfer knowledge from teacher models that are trained on fully visible (non-occluded) faces.

Closer to our method are the approaches designed for the difficult VR setting [16,29,32], in which a VR headset covers the upper side of the face. Hickson et al. [29] proposed a method that analyzes expressions from the eye region. The eye region is captured by an infrared camera mounted inside the VR headset, making the method less generic. Georgescu et al. [16] proposed an approach that analyzes the mouth region captured with a standard camera. The same type of approach is adopted by Houshmand et al. [32], but the experiments are conducted on a data set that uses face landmarks to apply the occlusion. In this work, we consider the same setting as [16,32], investigating the task of facial expression recognition when the upper half of the face is occluded. Unlike these closely related papers [16,32], we propose to perform knowledge distillation to produce more accurate CNN models. We study two knowledge distillation techniques 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. Moreover, we show the generality of the studied methods across different tasks and neural models.

2.2 Gender recognition

Gender prediction models are widely used across different domains such as advertising, security and human–computer interaction. Similar to [2,6,19,26,36,37,58,77], we focus on gender prediction from facial images. Some of these works [2,19,26,36,58] proposed the use of CNN models to predict the gender. Abirami et al. [2] used a CNN model to jointly predict the gender and the age of a person from facial images. Priadana et al. [57] proposed the use of CNN models to predict the gender based on profile pictures posted on a social media platform. Jhang et al. [37] proposed an ensemble of CNNs using a weighted-softmax voting scheme to address the gender prediction task. The weighted-softmax voting scheme is obtained by applying a fully connected layer on top of the models' predictions. The fully connected layer is further trained to learn the weights for each model. The weights of the other models are frozen during the additional training process. Georgescu et al. [19] proposed a ResNet-50 model based on pyramidal neurons with apical dendrite activations (PyNADA) to address the gender prediction task, among several other tasks. The neurons are equipped with a novel activation function inspired by some recent neuroscience discoveries.

Similarly to the main body of recent works on gender recognition in images, we also employ a CNN model. Different from all the works mentioned so far [2,19,26,36,58], which identify the gender in fully visible faces, we focus on a harder problem, that of predicting the gender in strongly occluded faces. More precisely, we aim to identify the gender of a person wearing a surgical mask, half of the face being occluded. While there are studies on face recognition under occlusion caused by surgical masks [10], or studies on gender recognition under random occlusions [38], to the best of our knowledge, we are the first to propose knowledge distillation for the gender prediction task under consistently strong occlusion of the lower half of the face.

2.3 Age estimation

Age estimation is a very important task across many domains such as advertising, human–computer interaction and security. There are many works [14,15,23,24,43,53,72,75,82] addressing the problem of age estimation of people from facial images. Before the era of deep learning, the age estimation task was tackled by extracting handcrafted features from facial images, then applying a regressor or a classifier on top of the extracted features. Guo et al. [23] proposed the use of biologically inspired features based on Gabor filters to estimate the age of a person from an image. After extracting the bio-inspired features, Guo et al. [23] applied the PCA algorithm to reduce the number of components. In the end, an SVR model is employed on top of the extracted features.

In another work based on biologically inspired features [24], people were separated based on gender and age groups, proving that the segregation approach improves the performance of the model by a significant margin.

Wang et al. [72] employed a CNN model to estimate the age of a person, given a facial image. Instead of estimating the age based on the features resulting from the last layer, Wang et al. [72] extracted features from different layers of the CNN model and concatenated them to obtain an aging pattern. In the end, both SVM and SVR models are applied on top of the aging patterns. Nam et al. [53] also used a CNN model to estimate the age, but before applying the CNN model, the authors employ a generative adversarial network [22] to increase the resolution of the image in order to improve the accuracy of the age predictor.

Similarly to the recent age prediction methods [53,72], we employ a CNN model to estimate the age of people from facial images. However, the main difference is that we are interested in estimating the age under severe occlusion. To the best of our knowledge, we are the first to study the age estimation task on severely occluded images (50% of the face being occluded).

2.4 Knowledge distillation

Knowledge distillation [4,30] is a recently studied technique [13,48,54,76,78,79] that enables the transfer of knowledge between neural networks. Knowledge distillation is a framework that unifies model compression [4,13,30] and learning under privileged information [48,68], the former one being more popular than the latter. In model compression, knowledge from a large neural model [4,54] or an ensemble of large neural networks [30,78,79] is distilled into a shallower or thinner neural network that runs efficiently during inference. In learning under privileged information, knowledge from a learning model trained with privileged information (some additional data representation not available at test time) is transferred to another learning 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 the best of our knowledge, we are the first to propose the distillation of knowledge using triplet loss. We should underline that there are a few previous works [13,54,78,79] 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.

3 Methods

To demonstrate the effectiveness of our approach across various models, we consider three neural network architectures namely, VGG-f [7], VGG-face [55] and ResNet-50 [28]. Each of the considered neural architectures is applied on three different tasks that involve drastic face occlusions. Regardless of the task, our training procedure based on knowledge distillation [4,17,30,68] is organized in a curriculum composed of the following three steps:

1. Training a teacher neural model on fully visible faces.
2. Fine-tuning the teacher on half-visible faces, resulting in a pre-trained student model.
3. Distill knowledge from the teacher model at step 1 into the student model obtained at step 2.

We emphasize that, due to catastrophic forgetting [49], the model fine-tuned at step 2 forgets useful information learned on fully visible faces. We conjecture that the lost information can be recovered at step 3, by distilling knowledge from the model trained on fully visible faces. In this context, we refer to the model trained on fully visible faces as the *teacher* and the model fine-tuned on half-visible faces as the *student*. In our experiments detailed in Sect. 4, we present ablation results after pruning steps 2 and 3, one by one.

Our first step is to train the teacher models on fully visible faces. In the second step, the networks are fine-tuned on occluded faces, thus obtaining pre-trained student models ready to undergo knowledge distillation. In the final training step, the students are fine-tuned on occluded faces, simulating either the setting in which people wear a VR headset or the setting in which people wear a surgical mask.

We underline that each student architecture is identical to the corresponding teacher architecture, as our goal is to learn privileged information from the teacher [68]. Thus, during the training process, we never mix the network architectures, such that when the teacher is a VGG-f network, the corresponding student is also a VGG-f network. The same applies to VGG-face and ResNet-50. We do not aim to compress the models, i.e., to train a shallow and efficient model by distilling a deeper network. Instead, we aim to distill the knowledge from the teacher as privileged information for the student. The privileged information available for the student is the other half of the face (the occluded half) seen only by the teacher. More specifically, for people wearing VR headsets, the upper half of the face represents the privileged information, while for people wearing surgical masks, the privileged information is the lower half of the face.

In the remainder of this section, we describe the two teacher–student training strategies for learning to predict the facial expression, gender or age of people from images with strong face occlusion. The two teacher–student strategies are alternatively employed at step 3 in our training process. As an attempt to increase robustness and stability, we also combine the resulting student models into an ensemble based on meta-learning.

3.1 Conventional teacher–student training

Ba et al. [4] proposed a method to compress a deeper model into a shallow one, their purpose being that of training an efficient network to mimic the complex representation of a deeper neural network. Using a similar idea, Hinton et al. [30] proposed to distill the knowledge of an ensemble of models into a single neural architecture. In their study, the authors demonstrated that even when an entire class of samples is missing during training, through knowledge distillation, the student is still able to recognize the examples belonging to the missing class due to the knowledge received from the teacher. The knowledge distillation is performed by training the student model using a soft target distribution produced by the teacher model. Besides the knowledge distillation loss, the cross-entropy loss between the ground-truth labels and the student predictions is also added, obtaining the following weighted loss:

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

where θ_S are the weights of the student model S , y is the ground-truth label, N_T is the prediction of the teacher network T , N_S is the prediction of the student network S , and λ is a hyperparameter that controls the importance of each loss function forming the final loss \mathcal{L}_{KD} . The first term of \mathcal{L}_{KD} is the cross-entropy with respect to the soft prediction of the teacher, and the second term is the cross-entropy with respect to the ground-truth label. N_T^τ and N_S^τ are the softened predictions of the teacher T and the student S , respectively, where $\tau > 1$ is a temperature parameter for the softening operation. More precisely, 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, as follows:

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

Lopez et al. [48] introduced the generalized distillation, where a model can learn from a different teacher model, but also from a different data representation. When training the student model by optimizing the loss defined in Equation (1), the student can learn privileged information available only to the teacher. In our case, we employ the generalized distilla-

tion method to learn a student to recognize facial expressions on strongly occluded faces with privileged information coming from a teacher that has access to fully visible faces [68]. The application of the conventional teacher–student framework for facial expression recognition under drastic occlusion is illustrated in Fig. 1.

We underline that the aforementioned knowledge distillation methods [4,30,48] are suitable for classification problems with multiple classes, in which the softmax output has sufficient components to contain edifying information. As a solution for our regression and binary classification tasks, we distill the knowledge in the penultimate layer of the student. To distill knowledge at any given layer, including the penultimate one, we can employ the approach of Romero et al. [59]. In order to train a deeper and thinner student than the teacher, Romero et al. [59] provided *hints* learned by the teacher network to the student network. A *hint* is the output of a teacher’s hidden layer, that is used to guide the corresponding output of a student’s hidden layer. The loss function used to guide the training of the student is the L_1 distance between the output H_T of a teacher’s hidden layer and the output H_S of a student’s hidden layer. More precisely, the following weighted loss is proposed by Romero et al. [59]:

$$\mathcal{L}_{HT}(\theta_S) = \lambda \|H_T - H_S\|_1 + (1 - \lambda) \mathcal{L}(y, N_S), \quad (3)$$

where θ_S are the weights of the student model S , y is the ground-truth label, H_S is the output of a student’s hidden layer, H_T is the output of the teacher’s hint (hidden) layer, N_S is the final output of the student network S , and λ is a hyperparameter that controls the importance of the two components.

For the age estimation and gender recognition tasks, we employ the knowledge distillation method of Romero et al. [59]. We prefer the hint teacher–student paradigm [59] over the standard teacher–student method [30], due to the low number of available predictions in the final layers for gender recognition (at most two components for male versus female classification) and age estimation (one component for estimating the age on a continuous scale).

3.2 Teacher–student training with triplet loss

When employing the teacher–student paradigm, aside from training the student to predict the correct labels, related techniques add a loss term to minimize the difference between the output of the student and that of the teacher [4,30,48], or the difference between some intermediate layers [59] of the student and teacher networks. To our knowledge, we are the first to perform distillation by adding a triplet loss term.

In general, triplet loss [60] is employed in neural network training to produce close (similar) embeddings for objects belonging to the same class and distant embeddings when

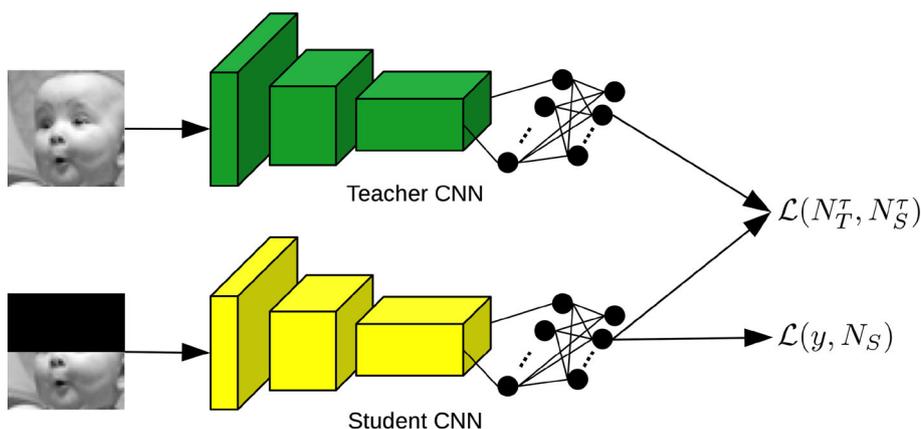


Fig. 1 The standard teacher–student training pipeline for facial expression recognition on severely occluded faces. The teacher CNN takes as input non-occluded (fully visible) faces, having access to privileged information. The student CNN takes as input only occluded (lower-half-visible) faces, but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

the objects belong to different classes. Our approach is to employ triplet loss on the face embeddings (the activations from the layer immediately before the final classification or regression layer). We aim to obtain similar face embeddings when the student network and the teacher network take an input image from the same class, and different embeddings otherwise.

We next present in detail how triplet loss can be applied to train the student network to learn privileged information provided by the teacher network. Throughout the remainder of this section, we use the prime symbol to denote an occluded face. Let x be a fully visible face and x' be an occluded face. Let $E_T(x)$ and $E_S(x')$ be the face embeddings produced by the teacher network T and the student network S , respectively. In order to employ the triplet loss, we need triplets of input images of the form (a', p, n') , where a' is an occluded image from a class k , p is a fully visible image from class k (the positive sample), and n' is an occluded image from a different class than k (the negative sample). During training, our goal is to reduce the distance between the anchor embedding $E_S(a')$ and the positive embedding $E_T(p)$ until it becomes smaller than the distance between the anchor embedding $E_S(a')$ and the negative embedding $E_S(n')$. In order to accomplish this goal, we use the following triplet loss function:

$$\mathcal{L}_{\text{triplet}}(\theta_S) = \sum_{i=1}^m \left[\|E_S(a'_i) - E_T(p_i)\|_2^2 - \|E_S(a'_i) - E_S(n'_i)\|_2^2 + \alpha \right]_+, \tag{4}$$

where θ_S are the learnable parameters of the student network S , m is the number of training examples, $[\cdot]_+ = \max(0, \cdot)$, and α is the margin (minimum distance) between the positive

pair of embeddings $(E_S(a'_i), E_T(p_i))$ and the negative pair of embeddings $(E_S(a'_i), E_S(n'_i))$.

Similarly to the standard teacher–student loss expressed in Equation (1), we want the student network to be able to reproduce the correct labels. Thus, our final loss function becomes:

$$\mathcal{L}_{KDT}(\theta_S) = (1 - \lambda)\mathcal{L}(y, N_S) + \lambda\mathcal{L}_{\text{triplet}}(\theta_S), \tag{5}$$

where θ_S are the learnable parameters of the student network S , $\mathcal{L}(y, N_S)$ is the loss function with respect to the ground-truth labels (cross-entropy for classification tasks and mean absolute error for regression tasks), $\mathcal{L}_{\text{triplet}}$ is the triplet loss function, and λ is a hyperparameter that controls the importance of the second objective function with respect to the first one. We illustrate our knowledge distillation paradigm based on triplet loss in Fig. 2. We underline that only the weights θ_S of the student are updated during training, while the embeddings $E_T(p_i)$ are kept unchanged during the whole knowledge distillation process.

Following [60], we propose a fast hard example mining scheme, as described next. In order to speed-up the training process, we generate the triplets offline, at the beginning of each epoch. One by one, each sample x' from the training set is selected as the anchor a' . For each anchor a' belonging to the class k , we randomly select a subset S_{pos} of fully visible training faces belonging to the same class. Next, we compute the distance between the embeddings produced by the teacher network for the fully visible faces in S_{pos} and the embedding produced by the student for the anchor a' . The sample which is the farthest from the anchor is selected as the positive example p :

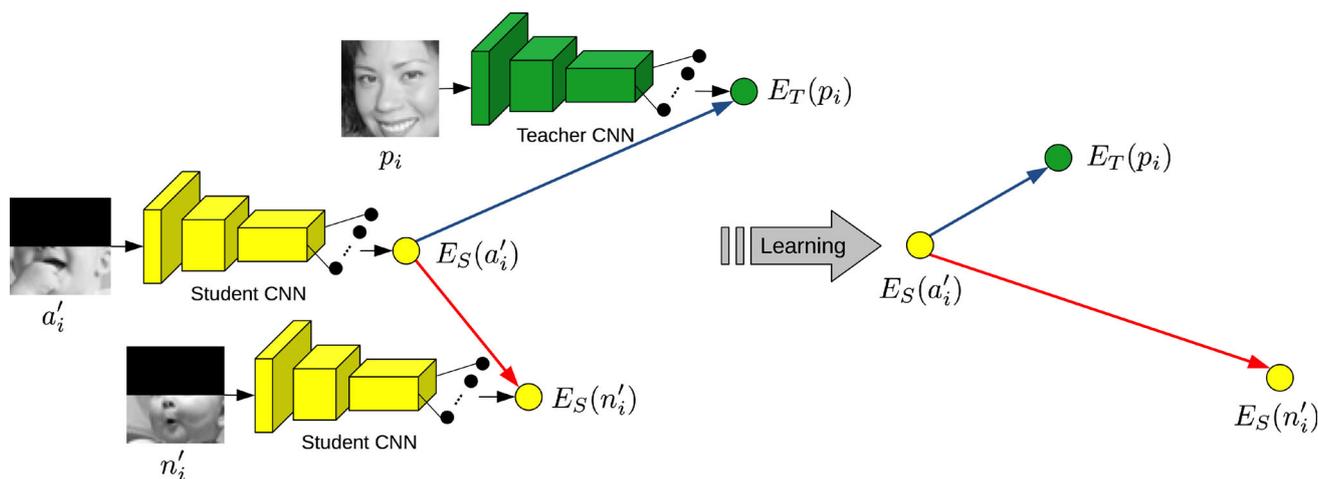


Fig. 2 The teacher–student training based on triplet loss for facial expression recognition on severely occluded faces. During training, we modify the weights of the student network such that the distance $\|E_S(a'_i) - E_T(p_i)\|_2^2$ becomes smaller than the distance $\|E_S(a'_i) - E_S(n'_i)\|_2^2$. Best viewed in color

$$p = \operatorname{argmax}_j \left\{ \|E_S(a') - E_T(p_j)\|_2^2 \right\}, \forall p_j \in S_{\text{pos}}. \quad (6)$$

In order to select the negative example n' , we randomly select a subset S_{neg} of half-visible training faces from a different class than k . Then, we compute the distance between the anchor embedding and the embeddings produced by the student model for the half-visible faces in S_{neg} . The sample which has the embedding closest to the anchor embedding is selected as the negative example n' :

$$n' = \operatorname{argmin}_j \left\{ \|E_S(a') - E_S(n'_j)\|_2^2 \right\}, \forall n'_j \in S_{\text{neg}}. \quad (7)$$

The random subsets S_{pos} and S_{neg} used in our hard example mining scheme contain only a small percentage of the entire training set (10% for facial expression recognition and 20% for age and gender estimation), speeding up the training time by a large margin. For additional efficiency improvements, we generate the subsets of positive and negative samples only once per epoch. Since age estimation is a regression task, we do not have truly positive or negative examples. Hence, for each anchor example, we consider examples under a difference of 5 years as positive and the other examples as negative. The threshold is set to roughly match the mean absolute error of the teachers on UTKFace.

4 Experiments

4.1 Organization

We hereby present experiments to demonstrate the efficiency of the proposed knowledge distillation methods. We consider

three tasks (facial expression recognition, gender recognition, age estimation) and three models (VGG-f, VGG-face, ResNet-50), evaluating all models across all tasks in three scenarios (fully visible faces, lower-half-visible faces, upper-half-visible faces). For each data set used in our experiments, we also present ablation results, considering as baselines the students trained on occluded faces (without employing any knowledge distillation technique) and the teachers trained on fully visible faces (without fine-tuning on occluded faces). Our knowledge distillation frameworks are employed only in the scenarios of interest (facial expression recognition on lower-half-visible faces, gender prediction on upper-half-visible faces, age estimation on upper-half-visible faces).

4.2 Facial expression recognition

4.2.1 Data sets

FER+ The FER+ data set [5] is a curated version of the FER 2013 data set [21]. The latter data set contains images with incorrect labels as well as images not containing faces. Barsoum et al. [5] cleaned up the FER 2013 data set by relabeling images and by removing those without faces. In the relabeling process, Barsoum et al. [5] added a new class of emotion, *contempt*, while also keeping the other 7 classes from FER 2013: *anger*, *disgust*, *fear*, *happiness*, *neutral*, *sadness* and *surprise*. The FER+ data set is composed of 25,045 training images, 3,191 validation images and 3,137 test images. The size of each image is 48×48 pixels.

AffectNet The AffectNet [51] data set is one of the largest data sets for facial expression recognition, containing 287,651 training images and 4,000 validation images with manual annotations. The images from AffectNet have various sizes. Since the test set is not yet publicly available, meth-

ods [18,51,64,65,70,71,81] 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 are unbalanced. As proposed by Mollahosseini et al. [51], we down-sample the training set for classes with more than 15,000 images. This leaves us with a training set of 88,021 images.

4.2.2 Data preprocessing

In order to train and evaluate the neural models in scenarios with drastic face occlusion, we replace the values of specific pixels with zero to simulate occlusions.

For the setting introduced by Georgescu et al. [16], 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. [29], in which facial expressions are recognized from the eye region, we occlude the entire lower half of the FER+ and AffectNet images. All images are resized to 224×224 pixels, irrespective of the data set, in order to be given as input to VGG-f, VGG-face and ResNet-50.

4.2.3 Evaluation metrics

Our metric for evaluating the classification models for facial expression recognition is the *accuracy* between the ground-truth labels and the predicted labels. Due to the fact that FER+ is highly imbalanced, we also report the *weighted accuracy* for this data set.

4.2.4 Implementation details

We emphasize that all the hyperparameters specified below are tuned on the FER+ validation set. Since the AffectNet validation set is used for the final evaluation, on AffectNet, we use the hyperparameter settings found optimal on the FER+ validation set. The VGG-f and VGG-face 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, since the latter model has a lower memory footprint. We use the same mini-batch sizes in all training stages. The ResNet-50 model is trained using the Adam optimizer [41] on mini-batches of 16 images. *Preliminary training of teachers and students* For the preliminary fine-tuning of the teacher and the student models, we use the MatConvNet [69] 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. 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 on occluded faces for 80 epochs. Further details about training these VGG-face and VGG-f baselines on fully visible or occluded faces are provided in [16]. The teacher ResNet-50 is fine-tuned for 75 epochs using a learning rate of 10^{-4} . Similarly, the student ResNet-50 is fine-tuned on lower-half-visible faces for 75 epochs with a learning rate set to 10^{-4} .

Standard teacher–student training For the conventional 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} . The student ResNet-50 is trained for 10 epochs starting with a learning rate of 10^{-5} . Hinton et al. [30] suggested to use a lower weight on the second objective function defined in Equation (1). Therefore, we set λ to a value of 0.9 for ResNet-50 and a value of 0.7 for both VGG-f and VGG-face models. The parameter λ is validated on the FER+ validation set.

Teacher–student training with triplet loss To implement the teacher–student training based on triplet loss, we switch to TensorFlow [1], exporting the VGG-f and VGG-face 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 VGG-f and ResNet-50 students for 10 epochs using a learning rate of 10^{-5} . The parameter λ in Equation (5) is set to a value of 0.5 for the VGG-f and VGG-face models, and a value of 0.1 for the ResNet-50 model. The value of the margin α from Equation (4) is chosen based on the performance measured on the validation sets, considering values in the interval $[0, 0.5]$ at a step of 0.1.

Combining distilled embeddings After training the student models using the two teacher–student strategies independently, we concatenate the corresponding face embeddings into a single embedding vector. The concatenated embeddings are provided as input to a Support Vector Machines (SVM) model [8]. The regularization parameter C of the resulting SVM models is chosen according to the performance on the validation sets, considering values between 10^{-1} and 10^3 , at a step of 10^1 . We use the SVM implementation from Scikit-learn [56].

4.2.5 Baselines

As baselines for facial expression recognition, we consider two state-of-the-art methods [16,29] 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) [16] or the upper half of the face (eye region) [29]. 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 [5,12,18,25,35,42,51,64–66,70,71]. We underline that these state-of-the-art methods are trained and tested on fully visible faces. Hence, the comparison to our approach or other approaches applied on occluded faces [16,29] is unfair, but we included it as a relevant indicator of the upper bound for the models applied on occluded faces. We also underline that the included state-of-the-art methods [64–66,70,71] are not based on standard modeling choices, using sophisticated loss functions, label smoothing, ensembles of multiple neural architectures or more of the above. Our approach is closer to the methods trained under more simple settings [5,25,51], since our teachers and students are trained using fairly well-known loss functions and the evaluation is performed using only one model.

4.2.6 Results

In Table 1, we present the empirical results obtained on AffectNet [51] and FER+ [5] by the VGG-f, VGG-face and ResNet-50 models based on our teacher–student training strategies in comparison with the results obtained by the state-of-the-art methods [5,12,25,35,42,51,64–66,70,71] tested on fully visible faces and by the methods [16,29] designed for the VR setting (tested on occluded faces).

Comparison with the state of the art First of all, we note that it is natural for the state-of-the-art methods [5,12,25,35,42,51,64–66,70,71] to achieve better accuracy rates (on fully visible faces) than our approach or the other approaches applied on occluded faces [16,29]. One exceptional case is represented by the VGG-face model of Georgescu et al. [16] and our student VGG-face model, as both of them surpass the bag-of-visual-words model [35] on both data sets. Another exception of the above observation is represented by the fine-tuned or distilled ResNet-50 students, both surpassing the bag-of-visual-words on FER+.

Comparison between lower-half- and upper-half-visible faces With respect to the baselines [16,29] designed for the VR setting, all our teacher–student training strategies provide superior results. We observe that the accuracy rates of Hickson et al. [29] are considerably lower than the accuracy rates of Georgescu et al. [16] (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 eye region (denoted by ☹) than from the mouth region (denoted by ☺). To test this hypothesis, we evaluate the teachers (VGG-f, VGG-face and ResNet-50) 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 24% and 49%) and the results on lower-half-visible faces (accuracy rates are between 37% and 71%), confirming our hypothesis. We also underline that the results attained by the teacher models on occluded faces are considerably lower than the results of the baselines [16,29] designed for the VR setting, although the teacher models attain results close to the state-of-the-art methods [5,12,25,35,42,51,64–66,70,71] 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.

Comparison with closely related methods The results presented in the last nine rows of Table 1 indicate that the teacher–student learning strategies provide very good results on lower-half-visible faces, surpassing the other methods [16,29] 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. We also 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 some state-of-the-art methods [5,25,35,42,51]. For example, our VGG-face with triplet loss and standard teacher–student training yields an accuracy rate of 82.75% on FER+, which is 2.24% under the state-of-the-art VGG-13 [5]. 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).

Statistical significance testing We also performed statistical significance testing to compare our models (VGG-f, VGG-face and ResNet-50) based on teacher–student training with the models of Georgescu et al. [16], which are equivalent with our students before undergoing distillation. Notably, the combined teacher–student strategies provide significant improvements for the VGG-f, VGG-face and ResNet-50 models on both data sets, with a significance level of 0.05.

Grad-CAM visualizations In order to better understand how our models make decisions, we used the Grad-CAM [61] approach to provide visual explanations for some image examples illustrated in Fig. 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 each face depicted in the second row of Fig. 3 is occluded. We observe that the neural architectures turn their attention on the lower part of the face, particularly on the mouth region. This indi-

Table 1 Accuracy rates of various models on AffectNet [51] and FER+ [5], for fully visible faces (denoted by ○), lower-half-visible faces (denoted by ◐) and upper-half-visible faces (denoted by ◑)

Model	Train faces	Test faces	AffectNet (%)	FER+	
				Accuracy (%)	Weighted accuracy (%)
VGG-13 [5]	○	○	–	84.99	–
DAFL [12]	○	○	65.20	–	–
CNNs+BOVW+LC [18]	○	○	59.58	87.76	–
VGG-12 [25]	○	○	58.50	–	–
Bag of visual words [35]	○	○	48.30	80.65	–
MT-VGG [42]	○	○	54.00	–	–
AlexNet [51]	○	○	58.00	–	–
Res-50IBN [64]	○	○	63.11	89.51	–
MBCC-CNN [65]	○	○	–	88.10	–
ESR-9 [66]	○	○	59.30	87.15	–
SCN [71]	○	○	60.23	89.35	–
PSR [70]	○	○	60.68	89.75	–
Teacher VGG-f	○	○	57.37	85.05	59.71
Teacher VGG-face	○	○	59.03	84.79	66.15
Teacher ResNet-50	○	○	56.07	85.91	65.67
Teacher VGG-f	○	◐	41.58	70.00	43.24
Teacher VGG-face	○	◐	37.70	68.89	39.69
Teacher ResNet-50	○	◐	40.50	70.89	44.56
Teacher VGG-f	○	◑	26.85	40.07	32.82
Teacher VGG-face	○	◑	31.23	48.29	37.36
Teacher ResNet-50	○	◑	24.12	44.21	30.01
VGG-f [16]	◐	◐	47.58	78.23	50.52
VGG-face [16]	◐	◐	49.23	82.28	58.69
ResNet-50	◐	◐	45.90	81.79	60.57
VGG-f [29]	◑	◑	42.45	66.18	44.66
VGG-face [29]	◑	◑	43.18	70.19	48.83
ResNet-50	◑	◑	43.37	72.26	54.62
VGG-f (standard T-S)	○+◐	◐	48.75 [†]	80.17 [†]	53.00 [†]
VGG-face (standard T-S)	○+◐	◐	49.75	82.37	59.46
ResNet-50 (standard T-S)	○+◐	◐	46.95 [†]	82.37 [†]	59.19 [†]
VGG-f (triplet loss T-S)	○+◐	◐	48.13	80.05 [†]	52.87 [†]
VGG-face (triplet loss T-S)	○+◐	◐	49.71	82.57	59.12
ResNet-50 (triplet loss T-S)	○+◐	◐	46.17	81.28	60.93
VGG-f (triplet loss + standard T-S)	○+◐	◐	48.70 [†]	81.09 [†]	58.90 [†]
VGG-face (triplet loss + standard T-S)	○+◐	◐	50.09 [†]	82.75 [†]	61.23 [†]
ResNet-50 (triplet loss + standard T-S)	○+◐	◐	47.00 [†]	82.37 [†]	59.30 [†]

The VGG-f, VGG-face and ResNet-50 models based on our teacher–student (T–S) training strategies are compared with state-of-the-art methods [5, 12, 18, 25, 35, 42, 51, 64–66, 70, 71] tested on fully visible faces and with methods [16, 29] designed for the VR setting (tested on occluded faces). The test results of our student networks that are significantly better than the stronger baseline [16], according to a paired McNemar’s test [9], are marked with † for a significance level of 0.05

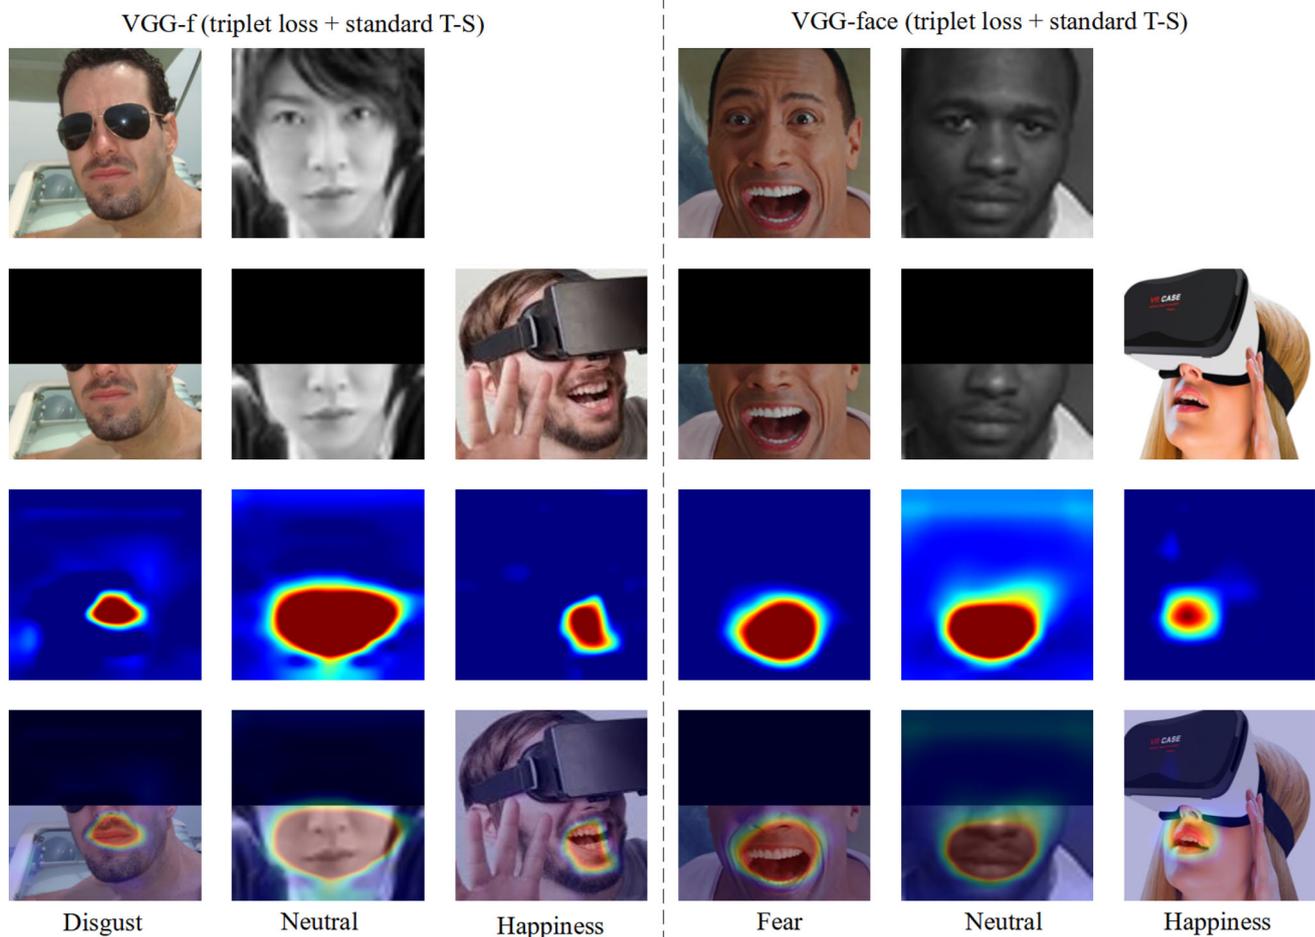


Fig. 3 Fully visible images (○) on top row, lower-half-visible faces (◐) on second row, Grad-CAM [61] explanation masks on third row and lower-half-visible faces with superimposed Grad-CAM masks on bottom row. The predicted labels provided by the distilled VGG-face (left-hand side) or VGG-f (right-hand side) models are also provided at the bottom. The first two examples from each side are selected from AffectNet [51] and FER+ [5], respectively. The third example from each side is a person wearing an actual VR headset. Best viewed in color

cates that our neural networks can properly handle situations in which people wear VR headsets. We note that the predicted labels for the first five samples presented in Fig. 3 are correct.

4.3 Age and gender estimation

4.3.1 Data sets

The UTKFace [83] data set contains images with faces of people of various age, gender and ethnicity. The data set consists of 23,689 images. We randomly divide the data set obtaining 14,213 (60%) images for training, 4,738 (20%) images for validation and 4,738 (20%) images for testing. The size of each image is 200×200 pixels. We perform two types of experiments on UTKFace: gender recognition and age estimation.

4.3.2 Data preprocessing

We adopt a similar technique as in the VR setting to simulate occlusions caused by surgical masks, thus replacing the values of pixels in the lower half of each face in the UTKFace data set with zero. The images from the data set are resized to 224×224 pixels to correspond to the input size of the convolutional networks.

4.3.3 Evaluation metrics

Our metric for evaluating the gender prediction models is the standard *classification accuracy*. To evaluate the regression models for age estimation, we consider the *mean absolute error* (MAE) between the predicted age and the target age of each test sample.

4.3.4 Implementation details

The networks are trained using the same optimizer and batch size as in the facial expression recognition experiments. We tune all other hyperparameters on the UTKFace validation set.

Preliminary training of teachers and students The teacher and student VGG-f models are fine-tuned for 200 epochs starting with a learning rate of 10^{-4} . We note that there is a teacher and a student for each of the two tasks (gender recognition and age estimation). Both teacher and student VGG-face networks are fine-tuned for 250 epochs with the learning rate set to 10^{-4} . The teacher and student ResNet-50 models are trained from scratch to predict the gender or the age of people in images, for 100 epochs. The learning rate for all teacher and student ResNet-50 models is set to 10^{-4} .

Standard teacher–student training In order to apply the standard teacher–student strategy to estimate the age and the gender from an image containing a face, the VGG-f students are each trained for 20 epochs starting with a learning rate of 10^{-6} . We set the parameter λ in Equation (3) to a value of 0.3 for gender prediction and a value of 0.5 for the age estimation task. The VGG-face students are trained for 40 epochs, setting the learning rate to a value of 10^{-5} for the age estimation task and a value of 10^{-6} for the gender prediction task. The parameter λ in Equation (3) is set to 0.3 for both tasks. The ResNet-50 students are each trained for a number of 40 epochs to estimate the age or the gender of a person, respectively. The starting learning rate is 10^{-5} , and we opted to decrease it when the validation error does not improve for 10 consecutive epochs. The parameter λ in Equation (3) is set to 0.1 for both tasks, based on the performance observed on the UTKFace validation set.

Teacher–student training with triplet loss The VGG-f model is trained for 30 epochs for the gender prediction task, starting with a learning rate of 10^{-5} . For the age estimation task, the model is trained for 40 epochs and the learning rate is set to 10^{-6} . The VGG-face model is trained for 40 epochs with the learning rate set to 10^{-6} for both tasks. The parameter λ in Equation (5) is set to 0.7 for both tasks and both VGG networks. Each of the two student ResNet-50 models is trained for 40 epochs with a learning rate of 10^{-5} . The value of the parameter λ is 0.7 for the age estimation task and 0.4 for the gender prediction task. The value of the margin α in Equation (4) is tuned on the validation set, considering values in the range $[0, 0.5]$ at a step of 0.1.

Combining distilled embeddings We concatenate the embeddings obtained from our two teacher–student training strategies and provide them as input either to an SVM model [8] for the classification task (gender recognition) or to an ϵ -support vector regression (SVR) for the regression task (age estimation). The regularization parameter C of these models is chosen according to the performance on the validation

set, considering values between 10^{-1} and 10^3 , at a step of 10^1 . We keep the default value for the parameter ϵ of the SVR, that is $\epsilon = 0$. As for the SVM, we employ the SVR implementation from Scikit-learn [56].

4.3.5 Baselines

As reference, we include the state-of-the-art results of the ResNet-50 based on pyramidal neurons with apical dendrite activations (PyNADA) reported in [19], although the results are not directly comparable due to the different splits applied on UTKFace. As baselines, we also include the student trained on lower-half-visible faces and the teacher trained on fully visible faces.

4.3.6 Gender recognition results

Comparison with the state of the art We present the gender prediction results in Table 2. The teacher ResNet-50 trained and evaluated on fully visible faces (denoted by \circ) reaches an accuracy of 90.88%, which is quite close to the top result reported in [19]. The VGG-f and VGG-face teachers surpass the performance reported in [19] by at least 2%.

Comparison between lower-half- and upper-half-visible faces When we evaluate the teachers on half-visible faces (denoted by \odot and \ominus), the accuracy rates drop by considerable margins. On the one hand, evaluating the teachers on the lower half (\odot) of the faces (the upper half being occluded) induces a performance drop between 14% and 19%. On the other hand, when we black out the lower half of each face, evaluating the teachers on upper-half-visible faces (\ominus), the accuracy decreases by less than 7.68% (the maximum drop being observed for the ResNet-50 teacher), suggesting that it is easier to perform gender recognition on the upper side of the face. Despite the significant performance drop, this represents an encouraging result for the surgical mask scenario.

When we train the students on half-visible faces and evaluate them in the same setting, the accuracy rates improve. The ResNet-50 student trained and evaluated on the lower-half-visible faces reaches an accuracy of 86.47%. The other ResNet-50 student, the one trained and evaluated on upper-half-visible faces, yields an accuracy of 88.75%. Both students obtain better performance compared with the teacher trained on fully visible faces, when the evaluation is conducted on half-visible faces.

Comparison with the baseline Further, we observe that, with respect to the baseline students, both teacher–student strategies attain superior performance. When we fine-tune the students using the standard teacher–student strategy, we obtain improvements ranging between 0.19% and 0.70%, reaching a top accuracy rate of 89.45% with the ResNet-50 model. The privileged information received from the teacher helps the student to outperform its ablated version trained

Table 2 Accuracy rates for gender prediction on UTKFace [83], for fully visible faces (denoted by ○), lower-half-visible faces (denoted by ◐) and upper-half-visible faces (denoted by ◑)

Method	Train faces	Test faces	Accuracy (%)
ResNet-50+PyNADA [19]	○	○	90.80
Teacher VGG-f	○	○	92.78
Teacher VGG-face	○	○	92.20
Teacher ResNet-50	○	○	90.88
Teacher VGG-f	○	◐	78.13
Teacher VGG-face	○	◐	73.05
Teacher ResNet-50	○	◐	72.69
Teacher VGG-f	○	◑	85.69
Teacher VGG-face	○	◑	88.18
Teacher ResNet-50	○	◑	83.20
VGG-f	◐	◐	88.70
VGG-face	◐	◐	90.62
ResNet-50	◐	◐	86.47
VGG-f	◐	◐	88.92
VGG-face	◐	◐	88.26
ResNet-50	◐	◐	88.75
VGG-f (standard T-S)	○+◐	◐	89.13
VGG-face (standard T-S)	○+◐	◐	88.45
ResNet-50 (standard T-S)	○+◐	◐	89.45 [†]
VGG-f (triplet loss T-S)	○+◐	◐	89.55 [†]
VGG-face (triplet loss T-S)	○+◐	◐	88.31
ResNet-50 (triplet loss T-S)	○+◐	◐	89.19
VGG-f (triplet loss + standard T-S)	○+◐	◐	89.82 [†]
VGG-face (triplet loss + standard T-S)	○+◐	◐	90.35 [†]
ResNet-50 (triplet loss + standard T-S)	○+◐	◐	89.63 [†]

A state-of-the-art model [19] is included as reference. The results of distilled models that are significantly better than the student trained on upper-half-visible faces, according to a paired McNemar's test [9] at a significance level of 0.05, are marked with †

with the standard loss (binary cross-entropy between the predicted labels and the ground-truth labels). The VGG-f student fine-tuned using the teacher–student training strategy based on triplet loss obtains an accuracy of 89.55%, surpassing its ablated version by 0.63%. Concatenating the VGG-face embeddings of the two knowledge distillation strategies provides an accuracy of 90.35%. We emphasize that the performance of the ensemble formed by the two ResNet-50 students evaluated on upper-half-visible faces (◐) is only 1.25% below the ResNet-50 teacher trained and evaluated on fully visible faces (○). We thus conclude that putting on a mask does not represent a significant problem for gender recognition when both fine-tuning and knowledge distillation are employed.

Statistical significance testing Furthermore, we also performed statistical significance testing to compare our distilled models with the ablated version of the student (the version before undergoing distillation). For all network types (VGG-f, VGG-face and ResNet-50), the results obtained by concatenating the two strategies and those obtained by

the model fine-tuned using the standard teacher–student paradigm are indeed statistically significant at a significance level of 0.05.

Grad-CAM visualizations To further investigate how our models make decisions, we employ Grad-CAM [61] to visualize what parts of the image are important in the gender recognition process. A set of representative Grad-CAM [61] visualizations are shown in Fig. 4. We observe that our model tends to concentrate on the upper half of the face, especially on the eye region. We notice that the focus area usually extends until it covers some hair. Thus, it is likely that our model considers the shape of the eyes and the length of the hair as discriminative features for gender prediction. Additionally, we underline that the predicted classes for the samples presented in Fig. 4 are all correct.

4.3.7 Age estimation results

Comparison with the state of the art In Table 3, we present the results for the age estimation task. The teacher ResNet-

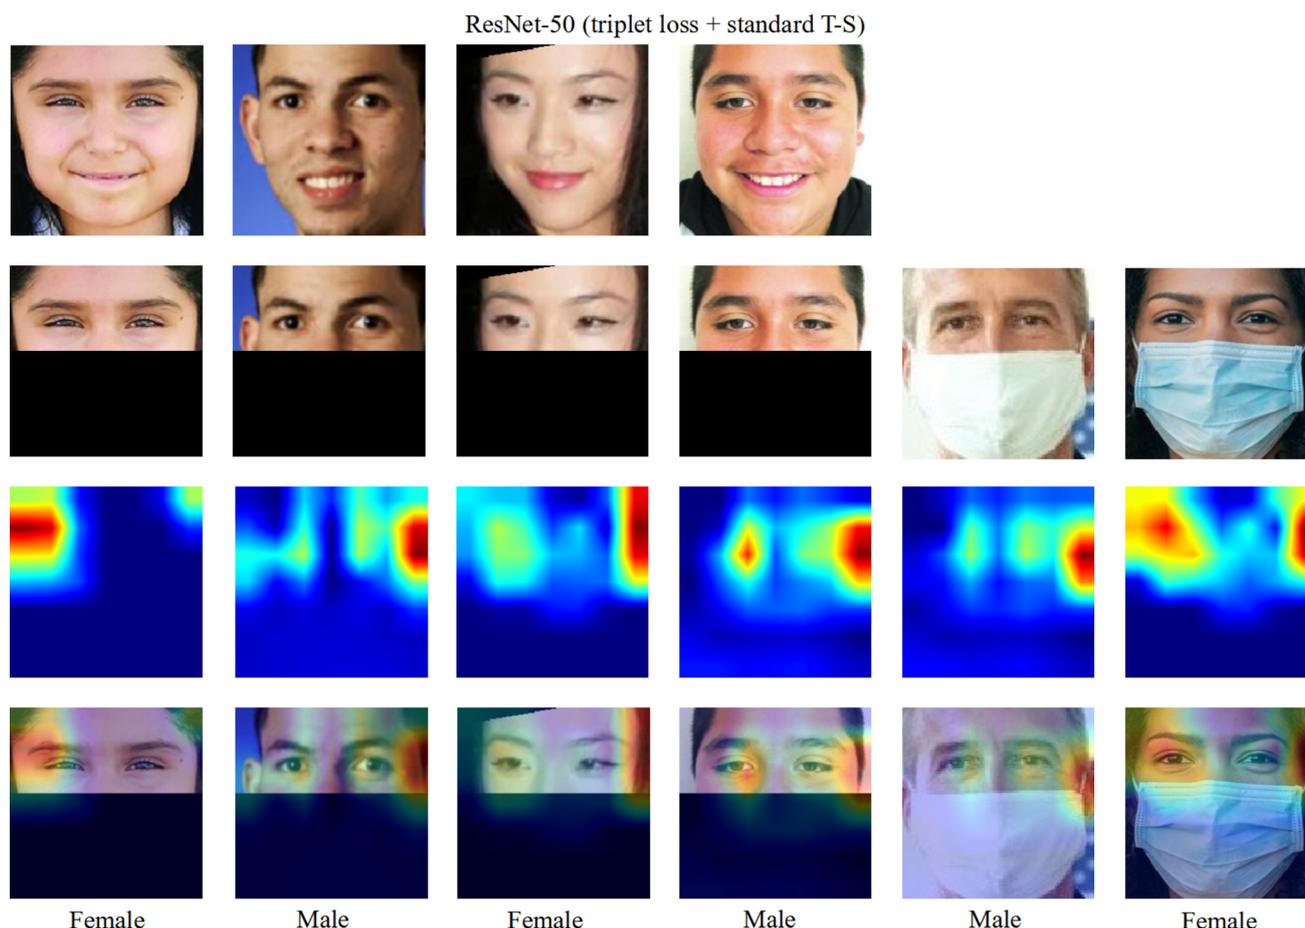


Fig. 4 Fully visible images (○) on top row, upper-half-visible faces (◐) on second row, Grad-CAM [61] explanation masks on third row and upper-half-visible faces with superimposed Grad-CAM masks on bottom row. The predicted gender provided by the distilled ResNet-50 model is shown at the bottom. The first four examples are selected from the UTKFace [83] data set. The last two examples are people wearing surgical masks. Best viewed in color

50 trained and evaluated on fully visible faces (denoted by ○) obtains an error of 5.27 years, surpassing the reference model proposed in [19].

Comparison between lower-half- and upper-half-visible faces When we evaluate the teacher models on the half-visible faces, the error increases by considerable margins. Indeed, evaluating the teacher models on lower-half-visible faces (denoted by ◑) increases the error up to 14.23 years. Similarly, when we evaluate them on upper-half-visible faces (denoted by ◐), the MAE grows up to 11.92 years.

Training the neural networks on half-visible faces and evaluating them in the same manner reduces the error by large margins. More precisely, the MAE goes down to 6.15 years on lower-half-visible faces and 5.53 years on upper-half-visible faces. Based on the results discussed above, we conclude that the upper half of the face is more informative for age estimation than the lower half.

Comparison with the baseline When fine-tuning the student models with the standard teacher–student training strategy,

the error decreases from 6.44 to 6.35 years for the ResNet-50 model, and from 5.53 to 5.40 years for the VGG-face model. Fine-tuning the ResNet-50 student using our teacher–student training paradigm based on triplet loss reduces the error to 6.34 years. By concatenating the embeddings of the two VGG-f students, we obtain a MAE of 6.22, this being our best improvement over the ablated version which scores a MAE of 6.36 years. In the end, the maximum difference between the teacher model trained and evaluated on fully visible faces and the student evaluated on lower-half-visible faces is 1.07 years. Hence, we conclude that the negative performance impact on age estimation generated by putting on a mask can be reduced by employing both fine-tuning and knowledge distillation.

Grad-CAM visualizations To explain how our distilled ResNet-50 models make decisions, we employ Grad-CAM [61] to observe what parts of the image are seen as important when estimating the age. We provide some relevant Grad-CAM [61] visualizations in Fig. 5. First, we observe that

Table 3 Mean absolute error (MAE) values for age estimation on UTKFace [83], for fully visible faces (denoted by ○), lower-half-visible faces (denoted by ◐) and upper-half-visible faces (denoted by ◑)

Method	Train faces	Test faces	MAE
ResNet-50+PyNADA [19]	○	○	5.79
Teacher VGG-f	○	○	5.63
Teacher VGG-face	○	○	5.11
Teacher ResNet-50	○	○	5.27
Teacher VGG-f	○	◐	11.16
Teacher VGG-face	○	◐	13.08
Teacher ResNet-50	○	◐	14.23
Teacher VGG-f	○	◑	9.60
Teacher VGG-face	○	◑	10.30
Teacher ResNet-50	○	◑	11.92
VGG-f	◐	◐	6.80
VGG-face	◐	◐	6.15
ResNet-50	◐	◐	6.66
VGG-f	◑	◑	6.36
VGG-face	◑	◑	5.53
ResNet-50	◑	◑	6.44
VGG-f (standard T-S)	○ + ◐	◐	6.34
VGG-face (standard T-S)	○ + ◐	◐	5.40
ResNet-50 (standard T-S)	○ + ◐	◐	6.35
VGG-f (triplet loss T-S)	○ + ◐	◐	6.34
VGG-face (triplet loss T-S)	○ + ◐	◐	5.42
ResNet-50 (triplet loss T-S)	○ + ◐	◐	6.34
VGG-f (triplet loss + standard T-S)	○ + ◐	◐	6.22
VGG-face (triplet loss + standard T-S)	○ + ◐	◐	5.40
ResNet-50 (triplet loss + standard T-S)	○ + ◐	◐	6.33

A state-of-the-art model [19] is included as reference

our student model tends to focus on the upper half of the face, especially on the forehead region. We conjecture that our model views the wrinkles formed on the forehead as a discriminative feature for age estimation. On the bottom of the figure, we show the predicted age (first number) and the ground-truth age (second number) for each image. For the selected samples, the maximum difference between the predicted and the actual age is 3 years. We consider these predictions as fairly accurate.

4.4 Ablation study regarding neural embeddings

To understand which part of the ensemble of distilled models brings a larger improvement, i.e., the concatenation of embeddings or the SVM model, we conduct an ablation study by training an SVM on top of each type of embedding extracted from the distilled students (without concatenating the embeddings). We present the corresponding results in Table 4. The experiments are performed for the VGG-f and VGG-face models on the FER+, AffectNet and UTK-Face data sets. For the facial expression recognition data sets (FER+, AffectNet), we select the distilled models trained

on lower-half-visible faces, while for the gender prediction data set (UTKFace), we select the distilled models trained on upper-half-visible faces. For the VGG-f model, we observe that the largest improvement is brought by the concatenation of embeddings, not by the SVM model. However, on the UTKFace data set, the SVM model trained on embeddings from the VGG-face based on triplet loss performs on par with the SVM ensemble of distilled models. In this particular case, it seems that the SVM itself makes the difference. Nevertheless, none of the SVMs trained on individual embeddings is able to surpass the performance of the SVM ensemble. We thus conclude that concatenating the embeddings is useful.

5 Conclusion

In this paper, we presented two teacher–student methods to improve the performance of neural network models evaluated in scenarios with strong face occlusion. We demonstrate that our methods generalize across different classification and regression tasks (facial expression recognition, age estimation, gender prediction) and neural architecture types

ResNet-50 (triplet loss + standard T-S)

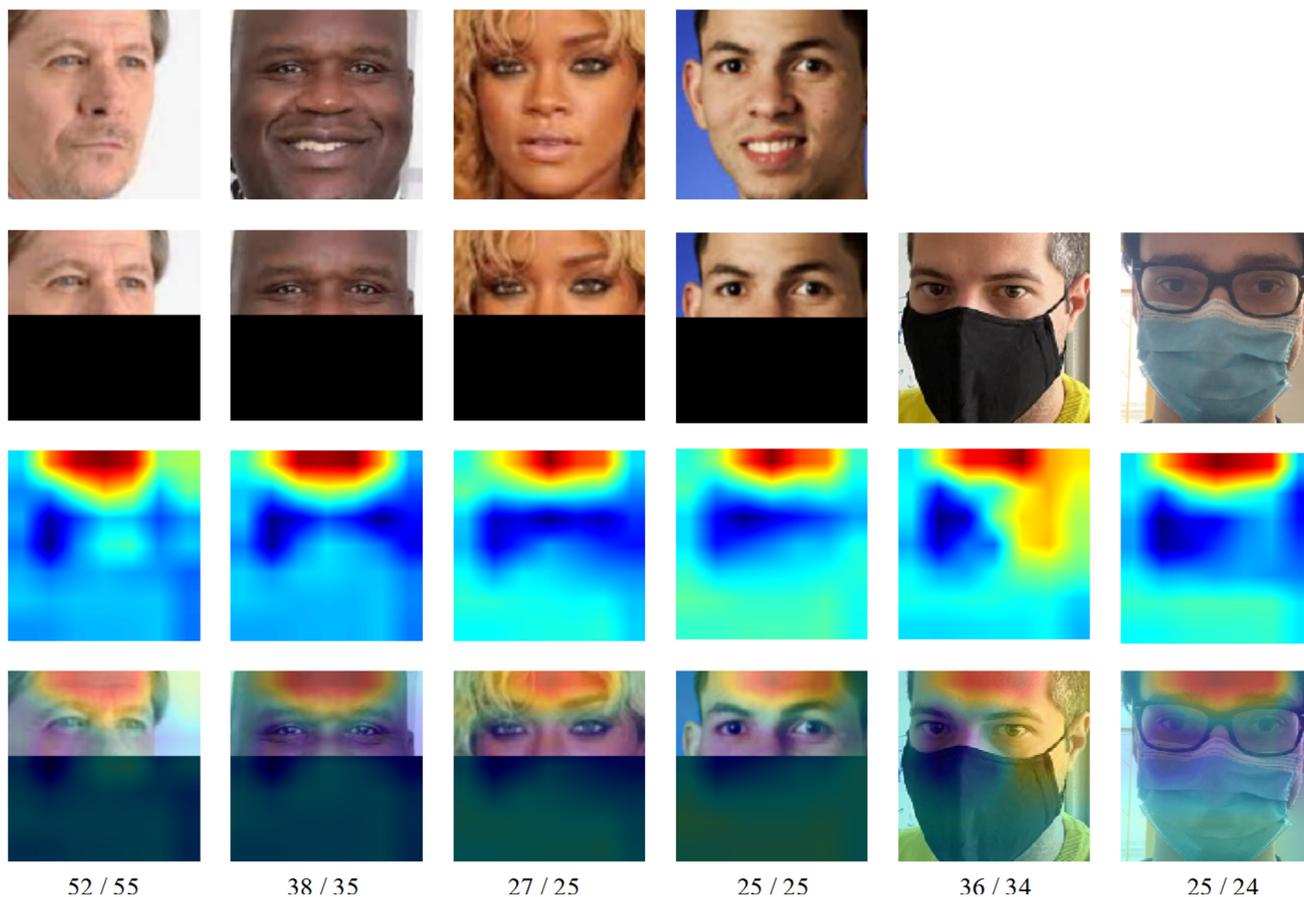


Fig. 5 Fully visible images (○) on top row, upper-half-visible faces (◐) on second row, Grad-CAM [61] explanation masks on third row and upper-half-visible faces with superimposed Grad-CAM masks on bottom row. The estimated age provided by the distilled ResNet-50 model is the first number shown at the bottom. The second number is the ground-truth age. The first four examples are selected from the UTKFace [83] data set, while the last two examples are people wearing masks. Best viewed in color

Table 4 The accuracy rates of SVMs trained on embeddings extracted from students based on standard teacher–student (TS) or triplet loss (TL) strategies

Network	Method	FER+ (%)	AffectNet (%)	UTKFace (%)
VGG-f	TS	80.17	48.75	89.13
	TL	80.05	48.13	89.55
	TS+SVM	80.39	48.52	89.04
	TL+SVM	79.06	47.01	89.70
	TS+TL+SVM	81.09	48.70	89.82
VGG-face	TS	82.37	49.75	88.45
	TL	82.57	49.71	88.31
	TS+SVM	82.34	48.89	90.35
	TL+SVM	82.37	49.90	90.27
	TS+TL+SVM	82.75	50.09	90.35

These models are compared with SVMs trained on concatenated embeddings as well as the students providing the embeddings. Results are reported for two tasks: facial expression recognition (on FER+ and AffectNet) and gender prediction (on UTKFace)

(VGG-face, VGG-f, ResNet-50). To the best of our knowledge, we are the first to study teacher–student strategies to learn privileged information aiming to cope with strong occlusions in images. We also proposed a novel teacher–student method based on triplet loss. The empirical results suggest that our knowledge distillation methods obtain superior performance over the baselines. On facial expression recognition, our ensemble of distilled models is only 2.24% below the performance of VGG-13 [5], when the former method is evaluated on occluded faces and the latter method is evaluated on fully visible faces. Similarly, on gender recognition, our ensemble of distilled ResNet-50 models is only 1.25% below the corresponding teacher. On age estimation, the difference between our ensemble of distilled ResNet-50 models applied on occluded faces and the teacher applied on fully visible faces is 1.06 years. In conclusion, we believe that the performance levels of our distilled models are sufficiently high to be used in real-life scenarios, such as changing the environment in VR applications based on user’s emotions and estimating the age and gender of people wearing surgical masks in supermarkets or retail stores.

In future work, we aim to study if models trained under occlusion can be used to boost the performance of teachers on fully visible faces. Our future goal is to find an effective way to combine students specialized on specific parts of the face with teachers looking at the entire face. We believe that this line of research can lead to significant performance gains.

Acknowledgements We thank reviewers for their valuable feedback which led to important improvements of the article.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: a system for large-scale machine learning. In: Proceedings of OSDI, pp. 265–283 (2016)
- Abirami, B., Subashini, T., Mahavaishnavi, V.: Gender and age prediction from real time facial images using CNN. *Mater. Today Proc.* **33**, 4708–4712 (2020)
- Al Chanti, D., Caplier, A.: Improving bag-of-visual-words towards effective facial expressive image classification. In: Proceedings of VISIGRAPP, pp. 145–152 (2018)
- Ba, J., Caruana, R.: Do deep nets really need to be deep? In: Proceedings of NIPS, pp. 2654–2662 (2014)
- Barsoum, E., Zhang, C., Ferrer, C.C., Zhang, Z.: Training deep networks for facial expression recognition with crowd-sourced label distribution. In: Proceedings of ICMI, pp. 279–283 (2016)
- Bhaskar, A., Aneesh R.P.: Advanced algorithm for gender prediction with image quality assessment. In: Proceedings of ICACCI (2015)
- Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: delving deep into convolutional nets. In: Proceedings of BMVC, pp. 1–12 (2014)
- Cortes, C., Vapnik, V.: Support-vector networks. *Mach. Learn.* **20**(3), 273–297 (1995)
- Dietterich, T.G.: Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Comput.* **10**(7), 1895–1923 (1998)
- Ding, F., Peng, P., Huang, Y., Geng, M., Tian, Y.: Masked face recognition with latent part detection. In: Proceedings of ACMMM, pp. 2281–2289 (2020)
- Ding, H., Zhou, S.K., Chellappa, R.: FaceNet2ExpNet: regularizing a deep face recognition net for expression recognition. In: Proceedings of FG, pp. 118–126 (2017)
- Farzaneh, A.H., Qi, X.: Facial expression recognition in the wild via deep attentive center loss. In: Proceedings of WACV, pp. 2402–2411 (2021)
- Feng, Y., Wang, H., Hu, R., Yi, D.T.: Triplet distillation for deep face recognition. In: Proceedings of ICIP, pp. 808–812 (2020)
- Geng, X., Yin, C., Zhou, Z.: Facial age estimation by learning from label distributions. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(10), 2401–2412 (2013)
- Geng, X., Zhou, Z., Smith-Miles, K.: Automatic age estimation based on facial aging patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(12), 2234–2240 (2007)
- Georgescu, M.I., Ionescu, R.T.: Recognizing facial expressions of occluded faces using convolutional neural networks. In: Proceedings of ICONIP, pp. 645–653 (2019)
- Georgescu, M.I., Ionescu, R.T.: Teacher–student training and triplet loss for facial expression recognition under occlusion. In: Proceedings of ICPR, pp. 2288–2295 (2020)
- Georgescu, M.I., Ionescu, R.T., Popescu, M.: Local learning with deep and handcrafted features for facial expression recognition. *IEEE Access* **7**, 64827–64836 (2019)
- Georgescu, M.I., Ionescu, R.T., Ristea, N.C., Sebe, N.: Non-linear neurons with human-like apical dendrite activations. *arXiv preprint arXiv:2003.03229* (2020)
- Giannopoulos, P., Perikos, I., Hatzilygeroudis, I.: Deep learning approaches for facial emotion recognition: a case study on FER-2013. In: Advances in Hybridization of Intelligent Methods, pp. 1–16. Springer (2018)
- Goodfellow, I.J., Erhan, D., Carrier, P.L., Courville, A., Mirza, M., Hamner, B., Cukierski, W., Tang, Y., Thaler, D., Lee, D.H., Zhou, Y., Ramaiah, C., Feng, F., Li, R., Wang, X., Athanasakis, D., Shave-Taylor, J., Milakov, M., Park, J., Ionescu, R.T., Popescu, M., Grozea, C., Bergstra, J., Xie, J., Romaszko, L., Xu, B., Chuang, Z., Bengio, Y.: Challenges in representation learning: a report on three machine learning contests. In: Proceedings of ICONIP, vol. 8228, pp. 117–124 (2013)
- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Proceedings of NIPS, pp. 2672–2680 (2014)
- Guo, G., Guowang Mu, Fu, Y., Huang, T.S.: Human age estimation using bio-inspired features. In: Proceedings of CVPR, pp. 112–119 (2009)
- Guo, G., Mu, G., Fu, Y., Dyer, C., Huang, T.: A study on automatic age estimation using a large database. In: Proceedings of CVPR, pp. 1986–1991 (2009)
- Guo, Y., Xia, Y., Wang, J., Yu, H., Chen, R.C.: Real-time facial affective computing on mobile devices. *Sensors* **20**(3), 870 (2020)
- Hacıbeyoğlu, M., Ibrahim, M.H.: Human gender prediction on facial mobil images using convolutional neural networks. *Int. J. Intell. Syst. Appl. Eng.* **6**(3), 203–208 (2018)
- Hasani, B., Mahoor, M.H.: Facial expression recognition using enhanced deep 3D convolutional neural networks. In: Proceedings of CVPRW, pp. 2278–2288 (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of CVPR, pp. 770–778 (2016)

29. Hickson, S., Dufour, N., Sud, A., Kwatra, V., Essa, I.: Eyemotion: classifying facial expressions in VR using eye-tracking cameras. In: Proceedings of WACV, pp. 1626–1635 (2019)
30. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: Proceedings of NIPS Deep Learning and Representation Learning Workshop (2014)
31. Hosseini, S., Cho, N.I.: GF-CapsNet: Using gabor jet and capsule networks for facial age, gender, and expression recognition. In: Proceedings of FG, pp. 1–8 (2019)
32. Houshmand, B., Khan, N.M.: Facial expression recognition under partial occlusion from virtual reality headsets based on transfer learning. In: Proceedings of BigMM, pp. 70–75 (2020)
33. Hu, J., Yu, B., Yang, Y., Feng, B.: Towards facial de-expression and expression recognition in the wild. In: Proceedings of ACII, pp. 157–163 (2019)
34. Hua, W., Dai, F., Huang, L., Xiong, J., Gui, G.: HERO: human emotions recognition for realizing intelligent internet of things. *IEEE Access* **7**, 24321–24332 (2019)
35. Ionescu, R.T., Popescu, M., Grozea, C.: Local learning to improve bag of visual words model for facial expression recognition. In: Proceedings of ICML Workshop on Challenges in Representation Learning (2013)
36. Ito, K., Kawai, H., Okano, T., Aoki, T.: Age and gender prediction from face images using convolutional neural network. In: Proceedings of APSIPA ASC, pp. 7–11 (2018)
37. Jhang, K.: Voting and ensemble schemes based on CNN models for photo-based gender prediction. *J. Inform. Process. Syst.* **16**(4), 809–819 (2020)
38. Juefei-Xu, F., Verma, E., Goel, P., Cherodian, A., Savvides, M.: DeepGender: occlusion and low resolution robust facial gender classification via progressively trained convolutional neural networks with attention. In: Proceedings of CVPRW, pp. 136–145 (2016)
39. Kaya, H., Gürpınar, F., Salah, A.A.: Video-based emotion recognition in the wild using deep transfer learning and score fusion. *Image Vis. Comput.* **65**, 66–75 (2017)
40. Kim, B.K., Roh, J., Dong, S.Y., Lee, S.Y.: Hierarchical committee of deep convolutional neural networks for robust facial expression recognition. *J. Multimodal User Interfaces* **10**(2), 173–189 (2016)
41. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of ICLR (2015)
42. Kollias, D., Zafeiriou, S.: Expression, affect, action unit recognition: Aff-Wild2, multi-task learning and ArcFace. In: Proceedings of BMVC (2019)
43. Lanitis, A., Draganova, C., Christodoulou, C.: Comparing different classifiers for automatic age estimation. *IEEE Trans. Syst. Man Cybern. Part B (Cybernetics)* **34**(1), 621–628 (2004)
44. Li, D., Wen, G.: MRMR-based ensemble pruning for facial expression recognition. *Multimed. Tools Appl.* **77**, 15251–15272 (2017)
45. Li, S., Deng, W., Du, J.: Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In: Proceedings of CVPR, pp. 2584–2593 (2017)
46. Li, Y., Zeng, J., Shan, S., Chen, X.: Patch-gated CNN for occlusion-aware facial expression recognition. In: Proceedings of ICPR, pp. 2209–2214 (2018)
47. Liu, X., Kumar, B., You, J., Jia, P.: Adaptive deep metric learning for identity-aware facial expression recognition. In: Proceedings of CVPRW, pp. 522–531 (2017)
48. Lopez-Paz, D., Bottou, L., Schölkopf, B., Vapnik, V.: Unifying distillation and privileged information. In: Proceedings of ICLR (2016)
49. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: the sequential learning problem. *Psychol. Learn. Motiv.* **24**, 109–165 (1989)
50. Meng, Z., Liu, P., Cai, J., Han, S., Tong, Y.: Identity-aware convolutional neural network for facial expression recognition. In: Proceedings of FG, pp. 558–565 (2017)
51. Mollahosseini, A., Hasani, B., Mahoor, M.H.: AffectNet: a database for facial expression, valence, and arousal computing in the wild. *IEEE Trans. Affect. Comput.* **10**(1), 18–31 (2019)
52. Mollahosseini, A., Hassani, B., Salvador, M.J., Abdollahi, H., Chan, D., Mahoor, M.H.: Facial expression recognition from World Wild Web. In: Proceedings of CVPRW, pp. 1509–1516 (2016)
53. Nam, S.H., Kim, Y.H., Truong, N.Q., Choi, J., Park, K.R.: Age estimation by super-resolution reconstruction based on adversarial networks. *IEEE Access* **8**, 17103–17120 (2020)
54. Park, W., Kim, D., Lu, Y., Cho, M.: Relational knowledge distillation. In: Proceedings of CVPR, pp. 3962–3971 (2019)
55. Parkhi, O.M., Vedaldi, A., Zisserman, A., et al.: Deep face recognition. In: Proceedings of BMVC, pp. 6–17 (2015)
56. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)
57. Priadana, A., Maarif, M.R., Habibi, M.: Gender prediction for instagram user profiling using deep learning. In: Proceedings of DASA, pp. 432–436 (2020)
58. Rafique, I., Hamid, A., Naseer, S., Asad, M., Awais, M., Yasir, T.: Age and gender prediction using deep convolutional neural networks. In: Proceedings of ICIC, pp. 1–6 (2019)
59. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: FitNets: hints for thin deep nets. In: Proceedings of ICLR (2015)
60. Schroff, F., Kalenichenko, D., Philbin, J.: FaceNet: a unified embedding for face recognition and clustering. In: Proceedings of CVPR, pp. 815–823 (2015)
61. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: visual explanations from deep networks via gradient-based localization. In: Proceedings of ICCV, pp. 618–626 (2017)
62. Shah, J.H., Sharif, M., Yasmin, M., Fernandes, S.L.: Facial expressions classification and false label reduction using LDA and threefold SVM. *Pattern Recogn. Lett.* **139**, 166–173 (2017)
63. Shao, J., Gori, I., Wan, S., Aggarwal, J.: 3D dynamic facial expression recognition using low-resolution videos. *Pattern Recogn. Lett.* **65**, 157–162 (2015)
64. She, J., Hu, Y., Shi, H., Wang, J., Shen, Q., Mei, T.: Dive into ambiguity: latent distribution mining and pairwise uncertainty estimation for facial expression recognition. In: Proceedings of CVPR, pp. 6248–6257 (2021)
65. Shi, C., Tan, C., Wang, L.: A facial expression recognition method based on a multibranch cross-connection convolutional neural network. *IEEE Access* **9**, 39255–39274 (2021)
66. Siqueira, H., Magg, S., Wermter, S.: Efficient facial feature learning with wide ensemble-based convolutional neural networks. *Proc. AAAI* **34**(04), 5800–5809 (2020)
67. Tang, Y.: Deep learning using linear support vector machines. In: Proceedings of ICML Workshop on Challenges in Representation Learning (2013)
68. Vapnik, V., Vashist, A.: A new learning paradigm: learning using privileged information. *Neural Netw.* **22**(5–6), 544–557 (2009)
69. Vedaldi, A., Lenc, K.: MatConvNet—convolutional neural networks for MATLAB. In: Proceeding of ACMML, pp. 689–692 (2015)
70. Vo, T.H., Lee, G.S., Yang, H.J., Kim, S.H.: Pyramid with super resolution for in-the-wild facial expression recognition. *IEEE Access* **8**, 131988–132001 (2020)

71. Wang, K., Peng, X., Yang, J., Lu, S., Qiao, Y.: Suppressing uncertainties for large-scale facial expression recognition. In: Proceedings of CVPR, pp. 6897–6906 (2020)
72. Wang, X., Guo, R., Kambhamettu, C.: Deeply-learned feature for age estimation. In: Proceedings of WACV, pp. 534–541 (2015)
73. Wen, G., Hou, Z., Li, H., Li, D., Jiang, L., Xun, E.: Ensemble of deep neural networks with probability-based fusion for facial expression recognition. *Cogn. Comput.* **9**(5), 597–610 (2017)
74. Wikanningrum, A., Rachmadi, R.F., Ogata, K.: Improving lightweight convolutional neural network for facial expression recognition via transfer learning. In: Proceedings of CENIM, pp. 1–6 (2019)
75. Xia, M., Zhang, X., Liu, W., Weng, L., Xu, Y.: Multi-stage feature constraints learning for age estimation. *IEEE Trans. Inf. Forensics Secur.* **15**, 2417–2428 (2020)
76. Yim, J., Joo, D., Bae, J., Kim, J.: A gift from knowledge distillation: fast optimization, network minimization and transfer learning. In: Proceedings of CVPR, pp. 7130–7138 (2017)
77. You, Q., Bhatia, S., Sun, T., Luo, J.: The eyes of the beholder: gender prediction using images posted in online social networks. In: Proceedings of ICDM Workshops, pp. 1026–1030 (2014)
78. You, S., Xu, C., Xu, C., Tao, D.: Learning from multiple teacher networks. In: Proceedings of KDD, pp. 1285–1294 (2017)
79. Yu, L., Yazici, V.O., Liu, X., Weijer, J.v.d., Cheng, Y., Ramisa, A.: Learning metrics from teachers: compact networks for image embedding. In: Proceedings of CVPR, pp. 2907–2916 (2019)
80. Yu, Z., Zhang, C.: Image based static facial expression recognition with multiple deep network learning. In: Proceedings of ICMI, pp. 435–442 (2015)
81. Zeng, J., Shan, S., Chen, X.: Facial expression recognition with inconsistently annotated datasets. In: Proceedings of ECCV, pp. 222–237 (2018)
82. Zeng, X., Huang, J., Ding, C.: Soft-ranking label encoding for robust facial age estimation. *IEEE Access* **8**, 134209–134218 (2020)
83. Zhang, Z., Song, Y., Qi, H.: Age progression/regression by conditional adversarial autoencoder. In: Proceedings of CVPR, pp. 5810–5818 (2017)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Mariana-luliana Georgescu received the B.Sc. degree from the Faculty of Mathematics and Computer Science, University of Bucharest, in 2017, and the M.Sc. degree in artificial intelligence from the University of Bucharest, in 2019, where she is currently pursuing the Ph.D. degree with the Faculty of Mathematics and Computer Science. She has published several articles at international peer-reviewed conferences and journals, including CVPR and TPAMI. Her research interests

include artificial intelligence, computer vision, machine learning, deep learning, and medical image processing.



Georgian-Emilian Duță received the B.Sc. degree from the Faculty of Mathematics and Computer Science, University of Pitești, in 2019. He is currently pursuing the M.Sc. in artificial intelligence from the Faculty of Mathematics and Computer Science, University of Bucharest. His research interests include artificial intelligence, computer vision, machine learning, and deep learning. He has published two articles at international peer-reviewed conferences, NeurIPS and WACV.



Radu Tudor Ionescu is Professor at the University of Bucharest, Romania. He completed his PhD at the University of Bucharest in 2013. He received the 2014 Award for Outstanding Doctoral Research in the field of Computer Science from the Romanian Ad Astra Association. His research interests include machine learning, text mining, computer vision, image processing and medical imaging. He published over 100 articles at international peer-reviewed conferences and journals,

and a research monograph with Springer. Radu received the “Caianiello Best Young Paper Award” at ICIAP 2013 for the paper entitled “Kernels for Visual Words Histograms”. Radu also received the “Young Researchers in Science and Engineering” Prize from prof. Rada Mihalcea and the “Danubius Young Scientist Award 2018 for Romania” from the Austrian Federal Ministry of Education, Science and Research and by the Institute for the Danube Region and Central Europe. Together with other co-authors, he obtained good rankings at several international competitions: 4th place in the Facial Expression Recognition Challenge of WREPL 2013, 3rd place in the NLI Shared Task of BEA-8 2013, 2nd place in the ADI Shared Task of VarDial 2016, 1st place in the ADI Shared Task of VarDial 2017, 1st place in the NLI Shared Task of BEA-12 2017, 1st place in the ADI Shared Task of VarDial 2018.