

MOON: A Mixed Objective Optimization Network for the Recognition of Facial Attributes

Ethan M. Rudd^(✉), Manuel Günther, and Terrance E. Boulton

Vision and Security Technology (VAST) Lab,
University of Colorado at Colorado Springs, Colorado Springs, USA
{erudd,mgunther,tboulton}@vast.uccs.edu

Abstract. Attribute recognition, particularly facial, extracts many labels for each image. While some multi-task vision problems can be decomposed into separate tasks and stages, e.g., training independent models for each task, for a growing set of problems joint optimization across all tasks has been shown to improve performance. We show that for deep convolutional neural network (DCNN) facial attribute extraction, multi-task optimization is better. Unfortunately, it can be difficult to apply joint optimization to DCNNs when training data is imbalanced, and re-balancing multi-label data directly is structurally infeasible, since adding/removing data to balance one label will change the sampling of the other labels. This paper addresses the multi-label imbalance problem by introducing a novel mixed objective optimization network (MOON) with a loss function that mixes multiple task objectives with domain adaptive re-weighting of propagated loss. Experiments demonstrate that not only does MOON advance the state of the art in facial attribute recognition, but it also outperforms independently trained DCNNs using the same data. When using facial attributes for the LFW face recognition task, we show that our balanced (domain adapted) network outperforms the unbalanced trained network.

Keywords: Facial attributes · Deep neural networks · Multi-task learning · Multi-label learning · Domain adaptation

1 Introduction

Given an input image or video, there are often multiple vision tasks to be accomplished, i.e., multiple objectives to be optimized. Under certain constraints, e.g., when tasks feed into each other, or when there is need to share computed features or representations, then multiple task objectives can benefit from being mixed and jointly optimized. This kind of multi-objective learning has affected many areas of computer vision including scene/object classification and annotation [1–4], tracking [5], facial landmark estimation [6, 7], face verification [8], and face detection with head pose estimation [9–11].

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46454-1_2](https://doi.org/10.1007/978-3-319-46454-1_2)) contains supplementary material, which is available to authorized users.

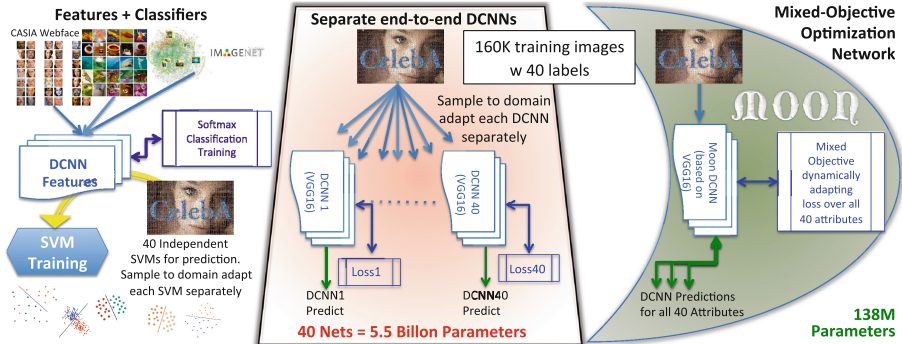


Fig. 1. Three approaches to attribute learning (and other multi-task problems). In the left is a conceptual model of previous state-of-the-art approaches, with features trained for classification problems and then adapted as inputs to independent SVMs for prediction. The middle approach attacks the problem with separately trained deep convolution neural networks (DCNNs). While we demonstrate that this advances the state of the art in attribute accuracy, it is cost prohibitive. This paper shows that for attributes, joint multi-task learning does better. However, for multi-label learning there is no way simply to re-weight or sample inputs to deal with imbalance or domain adaptation because each input defines values for all attributes. On the right is our answer, the mixed objective optimization network (MOON) architecture with a domain adaptive multi-task DCNN loss. To adapt, for each input the MOON objective re-weights each part of the loss associated with each attribute. MOON learns to balance its multi-task output predictions with reduced training and storage costs, while producing better accuracy than independently trained DCNNs.

This paper addresses facial attribute recognition, which we hypothesize is well suited to a multi-objective approach because facial attributes have a shared, albeit latent correlation that imposes soft constraints on the space of attributes, e.g., $p(\text{Male}|\text{Mustache}) \cong 1$. Despite the fact that facial attribute recognition inherently seeks multiple labels for the same image, multi-objective learning has not been widely applied to facial attributes. One potential reason is that balancing the training for the labels is difficult. Prior approaches to facial attribute recognition independently optimize a choice of features and recognition model (Features+Classifiers in Fig. 1). For example, the original approach taken by Kumar et al. [12] used AdaBoost to select a separate feature space for each attribute and independent SVMs to perform classification. Likewise, the current state of the art [13] trains DCNN features with facial identity recognition and localization datasets and then trains independent SVMs in this feature space for attribute classification. In both cases the separation makes it easy to re-balance training per attribute.

In this work, we show that a joint optimization with respect to all attributes offers performance superior to the state-of-the-art Features+Classifiers approach. We also show that joint optimization over all attributes outperforms training a single independent network of similar topology per attribute, in which

the feature space is optimized along with the classifier on a per-attribute basis, both in terms of accuracy and storage/processing efficiency. This result suggests that the multi-task approach is far more effective at distilling latent correlations than relying on independent classifiers to learn them implicitly. Thus, not only is a multi-objective approach far more intuitive, it is also far more effective.

It is unlikely that the source distribution of binary facial attributes for the training set will match the target distribution of the test set. We would like facial attribute classifiers trained on a dataset of one demographic to still work well for discrimination on a different demographic; thus some sort of domain adaptation is required. Many approaches to domain adaptation exist [14], with input sampling or re-weighting being common. Unfortunately, multi-objective training introduces challenges because balanced training is difficult or impossible via input sampling or weighting. Given a target distribution, domain adaptation is easy for separately trained attribute classifiers, e.g., by re-weighting errors in the cost function in each classifier. However, it is less immediately obvious how to do this in training for a multi-objective classifier. To this end, we introduce the MOON (Mixed-Objective Optimization Network) architecture. MOON is a novel multi-objective neural network architecture, which *mixes* the tasks of multi-label classification and domain adaptation under one unified objective function.

In summary, the contributions of this paper include:

- A mixed objective optimization network (MOON) architecture, which advances face attribute recognition by learning multiple attribute labels simultaneously via a single DCNN that supports domain adaptation for multi-task DCNNs.
- A fair evaluation technique which incorporates source and target distributions into the classification measure, leading to the balanced CelebA (CelebAB) evaluation protocol,
- Experiments demonstrating that the MOON architecture significantly advances state-of-the-art attribute recognition on the CelebA dataset, improving both accuracy and efficiency. These experiments also demonstrate that optimizing over all attributes simultaneously offers a noticeable reduction in classification error compared to optimizing single attributes over the same dataset and network topology.
- Experiments showing that domain adaptation on attribute classifiers trained on CelebA enhances the recognition capacity of MOON attributes on LFW, advancing attribute-based face recognition.
- Evaluation of stability of the MOON architecture to fiducial perturbations and data set imbalance.

2 Related Work

Multi-task learning has been applied to several areas that rely on learning fine-grained discriminations or localizations under the constraint of a global

correlating structure. In these problems, multiple target labels or objective functions must simultaneously be optimized. In object recognition problems, multiple objects may be present in a training image whose co-occurrences should be explicitly learnt [15]. In text classification problems, joint inference across all characters in a word yields performance gains over independent classification [16]. In multi-label image tagging/retrieval [4,17], representations of the contents of an image across modalities (e.g., textual descriptions, voice descriptions) are jointly inferred from the images. The resulting classifiers can then be used to generate descriptions of novel images (tagging) or to query images based on their descriptions (retrieval). Closer to this work, facial model fitting and landmark estimation [18,19] is another multi-task problem, which requires a fine-grained fit due to tremendous diversity in facial features, poses, lighting conditions, expressions, and many other exogenous factors. Solutions also benefit from global information about the space of face shapes and textures under different conditions. Optimization with respect to local gradients and textures is necessary for a precise fit, while considering the relative locations of all points is important to avoid violating facial topologies.

This paper applies multi-task learning to facial attributes. Applications of facial attributes include searches based on semantically meaningful descriptions (e.g., “Caucasian female with blond hair”) [12,20,21], verification systems that explain in a human-comprehensible form *why* verification succeeded or failed [22], relative relations among attributes [23], social relation/sentiment analysis [24], and demographic profiling. Facial attributes also provide information that is more or less independent of that distilled by conventional recognition algorithms, potentially allowing for the creation of more accurate and robust systems, narrowing down search spaces, and increasing efficiency at match time.

The classification of facial attributes was first pioneered by Kumar et al. [22]. Their classifiers depended heavily on face alignment, with respect to a frontal template, with each attribute using AdaBoost-learned combinations of features from hand-picked facial regions (e.g., cheeks, mouth, etc.). The feature spaces were simplistic by today’s standards, consisting of various normalizations and aggregations of color spaces and image gradients. Different features were learnt for each attribute, and a single RBF-SVM per attribute was independently trained for classification. Although novel, the approach was cumbersome due to high dimensional varying length features for each attribute, leading to inefficiencies in feature extraction and classification [25].

In recent years, approaches have been developed to leverage more sophisticated feature spaces. For example, gated CNNs [26] use cross-correlation across an aligned training set to determine which areas of the face are *most relevant* to particular attributes. The outputs of an ensemble of CNNs, one trained for each of the relevant regions, are then joined together into a global feature vector. Final classification is performed via independent binary linear SVMs. Zhang et al. [24] use CNNs to learn facial attributes, with the ultimate goal of using these features as part of an intermediate representation for a Siamese network to infer social relations between pairs of identities within an image. Liu et al. [13]

use three CNNs – a combination of two *localization networks* (LNet), and an *attribute recognition network* (ANet) to first localize faces and then classify facial attributes in the wild. The localization network proposes locations of face images, while the attribute network is trained on face identities and attributes, and is used to extract features, which are fed to independent linear SVMs for final attribute classification. Their approach was the state-of-the-art on the CelebA dataset at the time of the submission of this paper – and serves as a basis of comparison. In contrast to our approach, Liu et al. and many other recent works do not directly use attribute data in learning a feature space representation, but instead use truncated networks trained for other tasks. While research suggests that coarse-grained attribute data (e.g., image-level) can be indirectly embedded into the hidden layers of large-scale identification networks [27], the efficiency of this approach has not been well studied for inferring fine-grained (e.g., facial) attribute representations, and findings from [28] suggest that optimal implicit representations reside across different layers depending on the attribute.

Surprisingly, multi-task learning has not been widely applied to the problem of facial attribute recognition. Only very recently has it been addressed, e.g., Ehrlich et al. [29] developed a Multi-Task Restricted Boltzmann Machine (MT-RBM). In terms of joint inference for facial attributes, it is the first we could find in the literature, but the approach deviates radically from DCNN approaches in many other respects as well: the MT-RBM is generative and non-convolutional and it is unclear what contributed most to their improvement over [13].

While there has been significant prior work in visual domain adaptation [14], including more recent work for CNNs [30], the main problem that we address in this paper – incorporating domain adaptation into the training procedure for multi-objective attribute classifiers – has heretofore not been addressed, either in DCNN multi-task learning or in facial attribute research. For facial attributes in particular, we contend that domain adaptation is essential when building classifiers fit to chosen target demographics. Recently, Wang et al. [31] demonstrated that even throughout New York City, a relatively compact geographic region, differences in demographic profile are so prominent as a function of geolocation that binned geolocation can be used to derive a powerful unsupervised facial attribute feature space representation. In order to leverage attribute data we have for training demographic-specific classifiers, domain adaptation during training is vital to provide a balanced representation and mitigate problems from an over-correlated representation [32].

3 Approach

For multi-task problems, the high level goal is to maximize accuracy over all tasks, where each task has its own objective. In our case, the task is attribute prediction, and we seek to simultaneously maximize prediction accuracy over all attributes.

Formally, let \mathbb{I} be the space of allowable images, and let M be the number of attributes. For a given sample $x \in \mathbb{I}$, let $y_i \in \{-1, +1\}$ be the binary ground

truth label for x 's i th attribute, where $i \in \{1, \dots, M\}$ is the attribute index. Let \mathcal{H} be the space of allowable decision functions and $f_i(x; \theta_i) \in \mathcal{H}$ be the decision function, with parameters θ_i , learnt for the i th attribute classifier. Given a set of loss functions $L_i(f_i(x; \theta_i), y_i)$, each of which defines the cost of an error on input x with respect to attribute i , let $\mathbb{E}(f_i(x; \theta_i), y_i)$ be the expected value of that loss over the range of inputs \mathbb{I} . Then the idealized problem is to minimize the loss for each attribute, i.e.:

$$\forall i: f_i^* = \operatorname{argmin}_{f_i \in \mathcal{H}} \mathbb{E}(f_i(x; \theta_i), y_i). \quad (1)$$

For input x and attribute i , the classification result $c_i(x)$ and its corresponding error $e_i(x, y_i)$ are obtained by thresholding the associated prediction:

$$c_i(x) = \begin{cases} +1 & \text{if } f_i(x) > 0 \\ -1 & \text{otherwise,} \end{cases} \quad \text{and} \quad e_i(x, y_i) = \begin{cases} 0 & \text{if } y_i c_i(x) > 0 \\ +1 & \text{otherwise.} \end{cases} \quad (2)$$

Intuitively, this appears to lead to M independent optimization problems, for which one should be able to optimize each f_i separately. Accordingly, the most common approach to attribute classification in prior work is to use independent binary classifiers in some characteristic feature space to classify each attribute [13, 22]. Both approaches in [22] and [13] learn M independent binary classifiers trained with a hinge-loss objective. The hinge-loss objective function is:

$$\operatorname{argmin}_{\theta_i} L_i(x, \theta_i, y_i) = \max(0, 1 - y_i f_i(x; \theta_i)). \quad (3)$$

When the classifier is a dot product, i.e., $f_i(x) = \theta_i^T (1, x^T)^T$, solving this objective function results in a binary *support vector machine* (SVM) – the hyperplane that separates the two binary classes of data (+1 and -1) with maximum soft-margin. Given M attributes, this approach leads to M binary classifiers, each of which outputs a decision score. A positive decision score corresponds to the predicted presence of an attribute, while a negative decision score corresponds to its absence.

In order to learn latent correlations, it is also important to use attribute data directly to derive the feature space. Although Liu et al. [13] claim that latent features of attributes are learnt by their feature space representation while optimizing over a dataset for an identification task, the extent to which this is true for attributes that have little to do with identity (e.g., **Smiling**) is questionable. Rather, intuition suggests the opposite – that networks trained for identification of individuals would learn to ignore such attributes. To uncover such correlations, the network used to learn the feature space should be directly trained on attribute data and the distribution of attributes in training should match the operational or testing distribution.

This leads to the problem of how to appropriately balance the dataset used to learn attribute features. A perfectly balanced dataset can be obtained by collecting separate images for each attribute, but this leads to an enormous

dataset, with different identities for different attributes, effectively yielding a relatively small number of training images per attribute in proportion to the size of the dataset [22]. This approach also does not account for label correlations. Using a multi-label dataset, e.g., CelebA [13] allows us to leverage multiple labels in a mixed objective, but the distribution is highly imbalanced for many attributes (cf. Sect. 4). Unfortunately, the attribute distribution of a given target population does not always follow the dataset bias.

In a separate per-class training, balancing the number of positive and negative examples that are input to the classifier is easy, e.g., by weighting or sampling. However, input balancing is nearly impossible for multi-task training. Furthermore, for many tasks, the training frequencies and the operational/test frequencies will not match. Our solution to both problems is to define a mixed objective function including domain adapted weights that incorporate the difference between the source and target distributions. First, we compute the source distribution S_i from the training set for each attribute i by counting the relative number of occurrences of positive S_i^+ and negative samples S_i^- . Given a binary target distribution, T_i^+ and T_i^- , for each attribute i we assign a probability for each class:

$$p(i|+1) = \begin{cases} 1 & \text{if } T_i^+ > S_i^+ \\ \frac{S_i^- T_i^+}{S_i^+ T_i^-} & \text{otherwise} \end{cases} \quad \text{and} \quad p(i|-1) = \begin{cases} 1 & \text{if } T_i^- > S_i^- \\ \frac{S_i^+ T_i^-}{S_i^- T_i^+} & \text{otherwise.} \end{cases} \quad (4)$$

We would like to incorporate this domain adaptation directly into a loss function, but we need a loss function that additionally mixes all attribute predictions and simultaneously infers latent correlations between attribute labels and image data. One approach would be to combine all of the objective functions for each attribute into one joint objective function, e.g.:

$$\operatorname{argmin}_{\theta} \sum_{i=1}^M L_i(x, \theta, y_i), \quad (5)$$

where θ are the parameters of the joint classifier, which for legibility reasons we omit from the following equations. We can then solve that optimization problem via backpropagation using raw attribute images and labels as a training set. While we could use many potential loss functions, in our formulation we optimize a weighted mixed task squared error. Let M be the number of attributes, \mathbf{X} be a data tensor containing N input images, and \mathbf{Y} be a corresponding $N \times M$ matrix of labels. Then our domain-adapted multitask loss function is given by:

$$L(\mathbf{X}, \mathbf{Y}) = \sum_{j=1}^N \sum_{i=1}^M p(i|Y_{ji}) \|f_i(X_j) - Y_{ji}\|^2. \quad (6)$$

Replacing the standard loss layer of a DCNN with a layer implementing Eq. (6) results in the *mixed objective optimization network* (MOON) architecture, which incorporates attribute correlations and can adapt the bias of the training dataset

to a target distribution. In our custom implementation we obtain the weights $p(i|Y_{ji})$ via sampling. For each attribute i with target value $Y_{ji} \in \{-1, +1\}$ we only backpropagate the error with the probability $p(i|Y_{ji})$, otherwise we set the gradient for attribute i to 0. The more source and target distributions differ, the more elements in the gradient are reset.

4 Experiments

4.1 Dataset

For comparison with other attribute benchmarks, we conducted our experiments on the CelebA dataset [13]. The dataset consists of batches of 20 images from approximately 10K celebrities, resulting in a total of more than 200K images. Following the standard CelebA evaluation protocol, 8K identities (160K images) are used for training, 1K for validation and 1K for testing. Each image is annotated with 5 key points (both eyes, the mouth corners and the nose tip), as well as binary labels of 40 attributes. These attributes are shown in Fig. 2, which also shows the relative number of images in which the attribute is hand-labeled as present (blue) or absent (tan). As one can observe, for many of the attributes, there is a strong bias for either of the two classes. This is especially the case for certain attributes, e.g., relatively few images are labeled as **Bald** or **Wearing Hat**, while the majority of the facial images are labeled as **Young**.

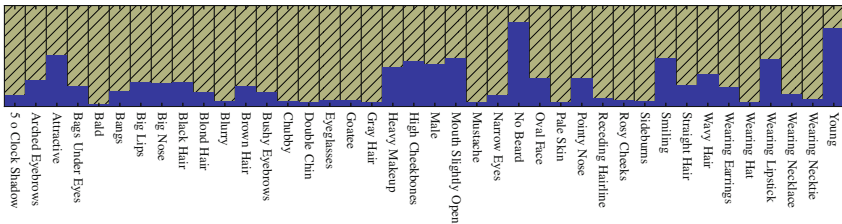


Fig. 2. CelebA Dataset Bias. This figure shows the distribution of the attribute labels throughout the CelebA dataset: presence (blue) or absence (tan). (Color figure online)

The CelebA dataset provides a set of pre-cropped face images, which were aligned using the hand-labeled key points. For our experiments we use these images, but later (cf. Sect. 5.1) we show that the trained classifier can also work with faces which are not perfectly aligned, and we introduce ideas to make our MOON network more robust to mis-alignment.

4.2 Evaluating MOON on CelebA

In order to compare with existing approaches, which do not account for dataset bias, we evaluate MOON on the CelebA dataset, setting the target distribution to the source distribution, i.e., $\forall i T_i \equiv S_i$.

Using the CelebA training set, we trained a DCNN to predict attributes under a MOON architecture. As the basic network configuration, we adopted the 16 layer VGG network from [33], where we replaced the final loss layer with the loss in Eq. (6). We also changed the dimension of the RGB image input layer from 224×224 pixels to 178×218 pixels, the resolution of the aligned CelebA images. In opposition to [33], we do not incorporate any dataset augmentation or mirroring, but train the network purely on the aligned images. Due to memory limitations, the batch size was set to 64 images per training iteration and, hence, the training requires approximately 2500 iterations to run a full epoch on the training set. We selected a learning rate of 0.00001, finding empirically that higher learning rates caused the network to learn only the bias of the training set. During training we update the convolution kernel weights using the backpropagation algorithm with an *RMSProp* update rule and an inverse learning rate decay policy.

We ran two types of network training, one training a separate network for each attribute, and one optimizing the combined MOON network. Separately training classifiers is the most common approach taken in the literature. By training one network per individual attribute, each network can concentrate only on the parts of the image it deems relevant to that attribute. During the separate training, we presented each network with all images from the training set, and a single input to the loss layer encoded with labels that denoted the presence (+1) or the absence (-1) of the attribute. Loss was computed according to Eq. (6). As each network required several hours to train on an NVIDIA Titan-X GPU, we chose to train each network for ≈ 2 epochs (5000 iterations). To check if 2 epochs are sufficient to attain convergence to a maximum validation accuracy, we continued training for four attributes. We selected these attributes – **Attractive**, **Chubby**, **Narrow Eyes**, and **Young** – to have varying statistics from the dataset: While **Attractive** is relatively balanced, images with **Chubby** and **Narrow Eyes** are mostly absent from the dataset, whereas **Young** is over-represented. While errors on the training set further decreased, errors on the validation set *increased* after approximately 4–6 epochs, with little improvement over the 2 epochs networks. This leads us to believe that improvements in validation accuracy beyond 2 epochs are negligible.

When training our MOON network, we use a single network with $M = 40$ outputs to learn all attributes simultaneously. Since CelebA has identical source and target distributions, we define the loss layer in (6) to weight all elements equally during backpropagation – which is equivalent to Euclidean loss between the network output and the 40 binary attribute values. We trained the network for 40 epochs since the validation error after 10 epochs was still decreasing. Based on the minimum validation set error, we chose our final MOON network after 24 epochs. While individual classifiers seem to take fewer training iterations than MOON to minimize their validation error, the total training time of the MOON network is still lower than the sum of the separate network training times. We suspect that the additional iterations required for the MOON network

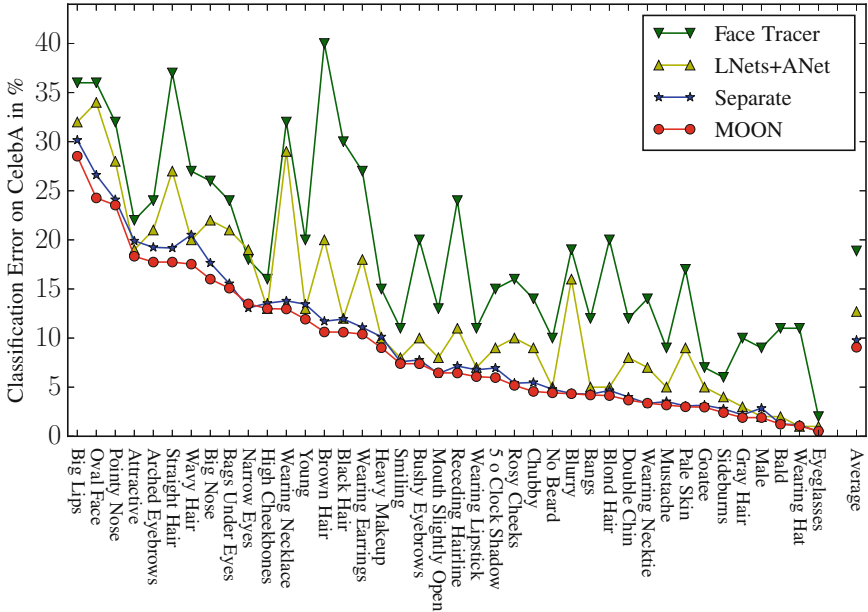


Fig. 3. Error Rates on CelebA. This figure shows the classification errors on the test set of the CelebA dataset for several algorithms, including our Separate networks and MOON. The results of Face Tracer and LNet+ANet are taken from Liu et al. [13]. For a tabular form of these results see the supplement to this paper.

to converge are needed to learn a more sophisticated latent structure than those learnt by the separate networks.

To compare with the results of Liu et al. [13], we measure the success of our training in terms of classification error, i.e., the number of cases, where our classifier f predicted the incorrect label, relative to the total number of test images:

$$E_i(\mathbf{X}, \mathbf{Y}) = \frac{1}{N_{test}} \sum_{j=1}^{N_{test}} e_i(X_j, Y_{ji}). \quad (7)$$

The **Average** classification error is computed by taking the average of the classification errors over all (M) attributes:

$$\bar{E}(\mathbf{X}, \mathbf{Y}) = \frac{1}{M} \sum_{i=1}^M E_i(X, Y). \quad (8)$$

Note that this error does not differentiate between positive and negative values. Hence, for very biased attributes, a random classifier which always predicts the dominant class would reach a low classification error, e.g., for **Bald** the random classification error would be as low as 2.24%!

The classification errors for all the attributes are visually displayed in Fig. 3. There, we also included two results from Liu et al. [13], converting from classification success (reported in [13]) to classification error. The Face Tracer results

reflect the best non-DCNN based algorithm that has been evaluated so far on the CelebA dataset. LNet+ANet represent the state-of-the-art results on this dataset obtained by combining three different deep convolutional neural networks with support vector machines.

The average classification errors over all attributes for each classifier are: Face Tracer: 18.88 %, LNet+ANet: 12.70 %, Separate: 9.78 %, and MOON: 9.06 %. Thus, our MOON network achieves a relative reduction of 28.7 % of the error over the state of the art, and a 7.4 % reduction over the separately trained networks. For almost all attributes, the results of our two approaches outperform the LNet+ANet state-of-the-art results, and the MOON network gives a lower error than the Separate networks trained specifically on a single attribute.

Interestingly, for several attributes that are traditionally not considered to be useful in face recognition, such as hair color (e.g. **Brown Hair**), hair style (e.g. **Straight Hair**), accessories (e.g. **Wearing Necklace**), and non face-related attributes (e.g. **Blurry**), our approach outperforms the LNet+ANet combination by an especially large margin. We suspect that this effect is due to the fact that in [13], the ANet network’s feature space was derived from training on a *face recognition* benchmark, and later adapted to the attribute classification task, which offers little direction for inferring the hidden representations of non facial identity related attributes.

4.3 CelebAB: A Balancing Act

As demonstrated in Sect. 4.2, MOON obtains state-of-the-art classification accuracies on the CelebA dataset. However, it is unclear how meaningful these results are for target distributions with different attribute frequencies, e.g., with more realistic distributions of **Young** or **Chubby** people.

Since our objective is to learn the network outputs to be +1 or -1 corresponding to presence or absence of attributes, respectively, we plotted the score distributions of the validation set for four of the attributes. From Fig. 2 we observe a strong bias for several attributes in the CelebA dataset, which we can find in the score distribution plots of Fig. 4(a), too. Note that the positive and negative score distributions have been normalized independently, otherwise the positive scores for **Narrow Eyes** and **Chubby** would not be visible. For attributes with a balanced number of positive and negative examples, such as **Attractive**, the distributions of negative (tan) and positive (blue) scores are also balanced. On the other hand, for unbalanced attributes, such as **Young**, **Narrow Eyes** or **Chubby**, the dominant class is well distributed around its desired value, but the other class has not been learnt well. Interestingly, a comparably small bias in the training set (for **Young** there are 77 % positives and 23 % negatives) can destroy the capability of the network to learn the inferior class.

Intuitively, when having such unbalanced score distributions, one would expect that the threshold of 0 that we use for classification should be adapted. However, given that the validation and test set follow the same bias as the training set, a threshold of 0 works well for the CelebA dataset. Even more astonishingly, a wide range of thresholds around 0 will lead to approximately the same

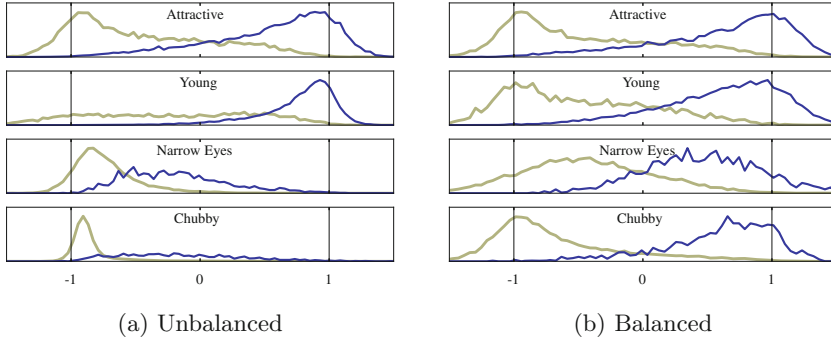


Fig. 4. Score distributions. This figure shows the distributions of the network outputs for four different attributes, when presenting images with present (blue) and absent (tan) attributes. In (a) network outputs after training with unbalanced data are shown, while in (b) the outputs of the network after training with the balancing loss layer are presented. Positive and negative score distributions are normalized independently. (Color figure online)

classification error and, hence, the network has learnt to balance between false positives and false negatives – including the dataset bias.

To obtain balanced score distributions, we chose to have a balanced target distribution, i.e., $T_i^+ = T_i^- = \frac{1}{2}$ for each attribute i . The resulting validation set score distribution for the same four attributes generated by the re-balanced MOON network after 34 training epochs can be seen in Fig. 4(b). Apparently, the score distributions are much more balanced, and the threshold 0 seems to make more sense now. Thus, one would expect that the classification error would be lower, too. However, due to the high dataset bias, which is also present in the validation and test sets, the total average classification error of the balanced network on the (unbalanced) CelebA test set is 13.67%.

Although this classification error is larger than that obtained by the unbalanced MOON network, *this is an artifact of the significant imbalance in the original test set*; the error measure in Eq. (7) has not been adapted to the target domain. A fair comparison would measure the balanced classification error E_i^B that weights the positive and negative classes according to the target distribution:

$$E_i^B(\mathbf{X}, \mathbf{Y}) = \sum_{j=1}^{N_{test}} \begin{cases} \frac{e_i(X_j, Y_{ji}) T_i^+}{N_i^+} & \text{if } Y_{ji} = +1 \\ \frac{e_i(X_j, Y_{ji}) T_i^-}{N_i^-} & \text{if } Y_{ji} = -1, \end{cases} \quad (9)$$

where N_i^+ and N_i^- are the respective numbers of positive and negative examples of attribute i in the test set. With $T_i^+ = T_i^- = \frac{1}{2}$, this error is effectively identical to the *equal error rate* (EER) between errors made with positive and negative target values. When computing classification error of the re-balanced MOON network example with $T_i^+ = T_i^- = \frac{1}{2}$, we obtain an average E_i^B error of 12.98%.

Note that the unbalanced MOON network, which is not trained to follow the target distribution, obtains an E_i^B error of 21.41%. This is precisely what we would expect of a domain adaptation system: A classifier adapted to the target distribution does better than a classifier that is not.

5 Discussion

5.1 Handling Mis-aligned Images

In our experiments in Sect. 4, we used aligned images to train and test the networks. To show that MOON is able to deal with badly aligned images, we conducted an additional experiment in which we used perturbed test images. To perturb the images, we applied a random rotation within $\pm 10^\circ$, a random scaling with a scale factor in $[0.9, 1.1]$, and a random translation of up to 10 pixels in either direction to the pre-aligned faces in the CelebA dataset. We selected these parameters to be well outside of the error range of a reasonable (frontal face) eye detector. Alignment errors of these magnitudes have been shown to *highly* influence the performance of many traditional face recognition algorithms [34].

When running this perturbed test set through our (unbalanced) MOON network, which was trained purely on aligned faces, we obtain a classification error of 11.62%, which is higher than the 9.06% obtained with aligned test images, but still better than the current state of the art in [13]. We assume that we can improve the network stability against mis-alignment by incorporating augmented (e.g., misaligned perturbations) training data into the training process, since this has shown to improve the performance of DCNNs [35].

Some preliminary experiments seem to verify this claim: When training with mis-aligned and horizontally mirrored images (in total 10 copies for each training image), we were able to decrease the classification error on the mis-aligned test images to 9.50%. Unfortunately, this also caused a slight performance degradation when evaluating on purely aligned images, causing classification error to increase from the 9.06% to 9.23%. Hence, in principle, the MOON architecture is able to work with aligned and mis-aligned images, as long as the conditions during training and testing are similar. These tests further highlight the need to select data augmentation methods appropriate to the respective quality of the actual alignment algorithms used in real end-to-end systems.

5.2 Face Verification on LFW

One application of facial attributes is to enhance other recognition algorithms. In order to evaluate our attribute classifiers on another dataset and to examine the effectiveness of our attributes for a particular application, we conducted the same View 2 LFW verification evaluation as Kumar et al. in [22], using the 40 attributes extracted from MOON under both balanced and unbalanced networks. We also tested the extracted attributes with respect to the features of Face Tracer (we downloaded the attribute vectors from [22] provided on the LFW

web page <http://vis-www.cs.umass.edu/lfw>), using the approach detailed in [22]. After optimizing RBF SVM parameters for each feature type separately using View 1 protocol of the labeled faces in the wild (LFW) dataset, the final classification accuracies that we obtained were $83.43\% \pm 2.22$ for Kumar’s attributes, and $85.05\% \pm 1.57$ for the re-balanced MOON network. Hence, our 40 MOON attributes provide better face recognition capabilities than the 73 attributes defined by Kumar et al. [22], though they are far from the current state-of-the-art on LFW. This result, consistent with intuition, suggests that the accuracy of the attribute classification is important to providing noticeably better recognition results. With $84.73\% \pm 1.99$, the verification accuracy for the unbalanced MOON network is only slightly lower than that of the re-balanced MOON network, but the stability is decreased. We assume that training on a target distribution that better reflects the distribution of facial attributes in LFW will result in further increased accuracy/stability. See the supplement to this paper for additional qualitative analysis of our LFW evaluation.

6 Conclusion

The MOON architecture achieves an accurate, computationally efficient, and compact representation which advances the state of the art on the CelebA dataset. Unlike competing approaches, our experiments did not rely on any datasets external to CelebA to train our network. We also investigated dataset bias in CelebA and proposed domain adaptation methods for training to a different target distribution without requiring training samples from that population. Combining domain adaptive methods and multiple-task objectives into one mixed objective function, we conducted evaluations on a novel re-balanced version of CelebA (the *CelebAB* dataset) and the LFW dataset that demonstrate the effectiveness of our approach.

Our work raises a philosophical question about the mathematics of attribute recognition: How *should* the attribute recognition problem be treated? Contrary to previous work, in which attribute labels are independently learnt, our approach implicitly leverages attribute correlations by explicitly forcing hidden layers in the network to incorporate information from multiple labels while simultaneously enforcing specified balance constraints via a domain adaptive loss. While CelebA labels are binary, MOON’s weighted Euclidean loss also offers the capacity to learn labels along a continuous range, which is perhaps a more suitable representation for some attributes (e.g., **Big Nose**, **Young**). Matching output score distributions to perceptual continuity and incorporating different types of attribute labels are interesting topics which we leave for future research.

Acknowledgments. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 2014-14071600012. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is

authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

References

1. Zhou, Z.H., Zhang, M.L.: Multi-instance multi-label learning with application to scene classification. In: *Advances in Neural Information Processing Systems*, pp. 1609–1616 (2007)
2. Quattoni, A., Collins, M., Darrell, T.: Transfer learning for image classification with sparse prototype representations. In: *Computer Vision and Pattern Recognition*, pp. 1–8. IEEE (2008)
3. Huang, Y., Wang, W., Wang, L., Tan, T.: Multi-task deep neural network for multi-label learning. In: *International Conference on Image Processing*, pp. 2897–2900. IEEE (2013)
4. Wu, F., Wang, Z., Zhang, Z., Yang, Y., Luo, J., Zhu, W., Zhuang, Y.: Weakly semi-supervised deep learning for multi-label image annotation. *Trans. Big Data* **1**(3), 109–122 (2015)
5. Zhang, T., Ghanem, B., Liu, S., Ahuja, N.: Robust visual tracking via multi-task sparse learning. In: *Computer Vision and Pattern Recognition*, pp. 2042–2049. IEEE (2012)
6. Zhang, Z., Luo, P., Loy, C.C., Tang, X.: Facial landmark detection by deep multi-task learning. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014, Part VI. LNCS*, vol. 8694, pp. 94–108. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10599-4_7](https://doi.org/10.1007/978-3-319-10599-4_7)
7. Zhang, C., Zhang, Z.: Improving multiview face detection with multi-task deep convolutional neural networks. In: *Winter Conference on Applications of Computer Vision*, pp. 1036–1041. IEEE (2014)
8. Wang, X., Zhang, C., Zhang, Z.: Boosted multi-task learning for face verification with applications to web image and video search. In: *Computer Vision and Pattern Recognition*, pp. 142–149. IEEE (2009)
9. Yan, Y., Ricci, E., Subramanian, R., Lanz, O., Sebe, N.: No matter where you are: flexible graph-guided multi-task learning for multi-view head pose classification under target motion. In: *International Conference on Computer Vision*, pp. 1177–1184. IEEE (2013)
10. Ouyang, W., Chu, X., Wang, X.: Multi-source deep learning for human pose estimation. In: *Computer Vision and Pattern Recognition*, pp. 2329–2336. IEEE (2014)
11. Yim, J., Jung, H., Yoo, B., Choi, C., Park, D., Kim, J.: Rotating your face using multi-task deep neural network. In: *Computer Vision and Pattern Recognition*, pp. 676–684. IEEE (2015)
12. Kumar, N., Belhumeur, P., Nayar, S.: FaceTracer: a search engine for large collections of images with faces. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008, Part IV. LNCS*, vol. 5305, pp. 340–353. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-88693-8_25](https://doi.org/10.1007/978-3-540-88693-8_25)
13. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: *International Conference on Computer Vision*, pp. 3730–3738. IEEE (2015)
14. Patel, V.M., Gopalan, R., Li, R., Chellappa, R.: Visual domain adaptation: a survey of recent advances. *Sig. Process. Mag.* **32**(3), 53–69 (2015)
15. Wei, Y., Xia, W., Huang, J., Ni, B., Dong, J., Zhao, Y., Yan, S.: CNN: single-label to multi-label (2014). arXiv preprint: [arXiv:1406.5726](https://arxiv.org/abs/1406.5726)

16. Jaderberg, M., Simonyan, K., Vedaldi, A., Zisserman, A.: Deep structured output learning for unconstrained text recognition. In: International Conference on Learning Representations (2015)
17. Huang, Y., Wang, W., Wang, L.: Unconstrained multimodal multi-label learning. *Trans. Multimedia* **17**(11), 1923–1935 (2015)
18. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. *Trans. Pattern Anal. Mach. Intell.* **23**(6), 681–685 (2001)
19. Blanz, V., Vetter, T.: Face recognition based on fitting a 3D morphable model. *Trans. Pattern Anal. Mach. Intell.* **25**(9), 1063–1074 (2003)
20. Kumar, N., Berg, A.C., Belhumeur, P.N., Nayar, S.K.: Describable visual attributes for face verification and image search. *Trans. Pattern Anal. Mach. Intell.* **33**(10), 1962–1977 (2011)
21. Scheirer, W.J., Kumar, N., Belhumeur, P.N., Boult, T.E.: Multi-attribute spaces: Calibration for attribute fusion and similarity search. In: *Computer Vision and Pattern Recognition*, pp. 2933–2940. IEEE (2012)
22. Kumar, N., Berg, A.C., Belhumeur, P.N., Nayar, S.K.: Attribute and simile classifiers for face verification. In: *International Conference on Computer Vision*, pp. 365–372. IEEE (2009)
23. Parikh, D., Grauman, K.: Interactively building a discriminative vocabulary of nameable attributes. In: *Computer Vision and Pattern Recognition*, pp. 1681–1688. IEEE (2011)
24. Zhang, Z., Luo, P., Loy, C.C., Tang, X.: Learning social relation traits from face images. In: *International Conference on Computer Vision*, pp. 3631–3639. IEEE (2015)
25. Wilber, M.J., Rudd, E., Heflin, B., Lui, Y.M., Boult, T.E.: Exemplar codes for facial attributes and tattoo recognition. In: *Winter Conference on Applications of Computer Vision*, pp. 205–212. IEEE (2014)
26. Kang, S., Lee, D., Yoo, C.D.: Face attribute classification using attribute-aware correlation map and gated convolutional neural networks. In: *International Conference on Image Processing*, pp. 4922–4926. IEEE (2015)
27. Escorcia, V., Niebles, J.C., Ghanem, B.: On the relationship between visual attributes and convolutional networks. In: *Computer Vision and Pattern Recognition*, pp. 1256–1264. IEEE (2015)
28. Zhong, Y., Sullivan, J., Li, H.: Leveraging mid-level deep representations for prediction face attributes in the wild. In: *International Conference on Image Processing*. IEEE (2016)
29. Ehrlich, M., Shields, T.J., Almaev, T., Amer, M.R.: Facial attributes classification using multi-task representation learning. In: *Computer Vision and Pattern Recognition Workshops*, pp. 47–55 (2016)
30. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: *International Conference on Computer Vision*, pp. 4068–4076. IEEE (2015)
31. Wang, J., Cheng, Y., Feris, R.S.: Walk and learn: Facial attribute representation learning from egocentric video and contextual data. In: *Computer Vision and Pattern Recognition*. IEEE (2016)
32. Jayaraman, D., Sha, F., Grauman, K.: Decorrelating semantic visual attributes by resisting the urge to share. In: *Computer Vision and Pattern Recognition*, pp. 1629–1636. IEEE (2014)
33. Parkhi, O.M., Vedaldi, A., Zisserman, A.: Deep face recognition. In: *British Machine Vision Conference*, vol. 1(3), p. 6 (2015)

34. Dutta, A., Günther, M., El Shafey, L., Marcel, S., Veldhuis, R., Spreuwers, L.: Impact of eye detection error on face recognition performance. *IET Biometrics* **4**, 137–150 (2015)
35. Simard, P.Y., Steinkraus, D., Platt, J.C.: Best practices for convolutional neural networks applied to visual document analysis. In: *International Conference on Document Analysis and Recognition*, pp. 958–963. IEEE (2003)