Demography-based Facial Retouching Detection using Subclass Supervised Sparse Autoencoder

Aparna Bharati[1], Mayank Vatsa[2], Richa Singh[2], Kevin W. Bowyer[1], Xin Tong[1]

[1]University of Notre Dame, [2]IIIT-Delhi

[1]{abharati, kwb, xtong1}@nd.edu, [2]{mayank, rsingh}@iiitd.ac.in

Abstract

Digital retouching of face images is becoming more widespread due to the introduction of software packages that automate the task. Several researchers have introduced algorithms to detect whether a face image is original or retouched. However, previous work on this topic has not considered whether or how accuracy of retouching detection varies with the demography of face images. In this paper, we introduce a new Multi-Demographic Retouched Faces (MDRF) dataset, which contains images belonging to two genders, male and female, and three ethnicities, Indian, Chinese, and Caucasian. Further, retouched images are created using two different retouching software packages. The second major contribution of this research is a novel semi-supervised autoencoder incorporating "subclass" information to improve classification. The proposed approach outperforms existing state-of-the-art detection algorithms for the task of generalized retouching detection. Experiments conducted with multiple combinations of ethnicities show that accuracy of retouching detection can vary greatly based on the demographics of the training and testing images.

1. Introduction

People are increasingly using software such as Adobe Photoshop, Pixlr and Fotor in order to improve the appearance of images contributed to media-sharing websites. Such processing when specifically applied to face images with the intent of improving attractiveness is what is considered digital facial retouching, as introduced in [13]. Some effects of the process include smooth skin, change in color and geometry of facial features and improved image lighting. Figure 1 presents few retouched versions of a face with some of these effects.

The resultant image from retouching is a plausible variant of the input face which, depending on the extent of retouching, may or may not match well to the original

Figure 1: Retouched samples of a face image using three different kinds of retouching tools. From top left to bottom right: original image, image retouched by *Fotor* [6], *PotraitPro Studio Max* [11] and *BeautyPlus* [1].

face [13]. When extensive retouching renders a face unrecognizable, this raises security concerns for face recognition systems [19]. There are scenarios where the gallery images for the Automatic Border Control (ABC) systems are collected through web uploading or mailing of printed photographs. In these cases, the quality or originality of the photographs cannot be strictly assured or checked. This demands for an automatic system that can raise an alarm in case of unacceptable level of alterations in submitted photographs.

Apart from the security concerns, there are also social concerns regarding the need for ideal appearance of an individual in the society [28]. A recent trend of retouched photos has been observed on online dating sites [4]. People retouch their photos to look younger and slimmer. Israel has announced the *'Photoshop Law'* [9] which regulates advertisements and magazine covers with photoshopped images to include declarations. The legislation was inspired from
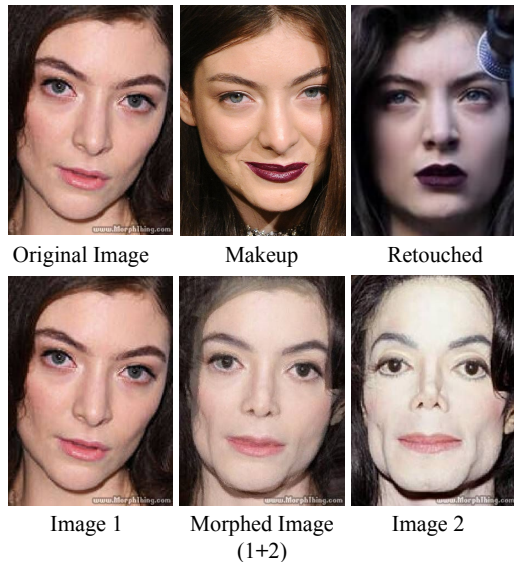
Figure 2: Examples of image with makeup and retouching (first row) and morphing (second row).

the correlation between increased rates of anorexia and bulimia in the young female population and media portrayals of the ideal size and figure. After Israel, the issue received attention in other countries such as France, UK and USA [7, 10, 12]. Therefore, an accurate automatic detector for retouching in images can help in better enforcement of these laws.

Facial retouching if performed to impersonate another identity, like morphing, has the potential to affect face recognition systems [19, 35]. Also, in a general scenario, the aim of retouching might not be to impersonate or spoof but to beautify the base face in a yet identifiable fashion [8]. For this kind of retouching, the facial features might change to some extent but the face still remains closer to the original identity than others in the face space. With spoofing, the aim is to push the face image away from the real identity and closer to another identity. Therefore, the methods used for detecting morphing and other kinds of spoofing might not be sufficient to detect retouching.

Another important related issue with face images is makeup. Even though makeup and retouching share the purpose of beautification and can easily be confused with each other, they are different processes. Real makeup is applied to a face before image acquisition while retouching and virtual makeup are applied after the image is acquired. Though retouching may include digital application of makeup, it is not limited to it. Retouching can yield various effects which is not possible through virtual makeup. For example, it can change facial shape which a virtual makeup tool cannot. Also, the lighting and ambience of the face in the picture cannot be altered through virtual makeup but such an effect can be achieved through retouching. Figure 2 shows an

example illustrating the case of morphing of two celebrity faces and real makeup versus retouching for a face.

1.1. Literature Review

In the existing literature, there are efficient algorithms designed specifically to predict the presence of retouching for images such as magazine covers. Kee *et al*. [24] propose a set of geometric and photometric features from face and body for retouching that correlates well with human perception. For profile pictures on social media and dating websites, the only information available can be obtained from face. The recent work by Bharati *et al*. [13] is relevant to facial retouching detection. It uses a deep learning framework and outperforms several handcrafted features. In this paper, the authors introduced the first retouched faces database - ND-IIITD Retouched Faces database, with images from 325 subjects (mostly Caucasian) and retouched images with one tool. They proposed a Supervised Deep Boltzmann Machine (SDBM) to extract features and a 2ν Support Vector Machine (SVM) for classifying the images into unaltered or retouched class.

Other detection algorithms relevant to our work are algorithms for makeup detection and morphing detection. For automatic detection of makeup, Chen *et al*. [14] propose extracting color, shape and texture features of three predefined facial regions which are classified using SVM with RBF kernel and Adaboost. They perform experiments using YouTube Makeup (YMU) database for training and Makeup in the wild (MIW) database for testing. Later, Kose *et al*. [25] proposed a more accurate makeup detection algorithm on the same datasets using texture and shape features.

A recent solution to detecting morphing for faces has been proposed by Raghavendra *et al*. [35]. The proposed algorithm uses Binarized Statistical Image Features (BSIF) and linear SVM. They report results on a dataset of 450 morphed face images from 110 individuals. These detectors have state-of-the-art performances in detecting makeup and morphing.

1.2. Research Contributions

Coetzee *et al*. [16] and Rhodes [36] found that the race of an individual affects their perception of beauty or facial attractiveness. They also suggest that the perception varies between males and females. These might lead to different desired retouching effects for groups of individuals, implying different subclass distributions within both original and retouched classes.

Previous research on retouching detection algorithms has not focused on detecting retouching in various demographics. Also, most of them only consider retouching effects provided by one tool which doesn't test the algorithm for tool biases. In this paper, we improve upon the existing literature by making the following key contributions:

- We introduce a Multi-Demographic Retouched Faces (MDRF) database which consists of retouched images from two retouching tools and three ethnicity variations.

- A novel subclass-based supervised sparse autoencoder (S^3A) formulation is proposed to impart specialized sparse encoding to samples. We claim that this helps generate subclass specific representations. A comparison of the proposed method with the existing state-of-the-art [13] retouching detection algorithm is provided for three different ethnicities.

- Through combined and cross-ethnicity performance analysis, we highlight the effects of different ethnicity, gender and tools on algorithms used for retouching detection.

2. Multi-Demographic Retouched Faces Dataset

Datasets used in previous research on retouching detection do not well represent different races and generally contain retouched images from a single software package. Considering these variants for retouched images is important for a real-world scenario where the image in question can belong to any ethnicity and can have specialized retouching effects pertaining to the preferences of the individual or group. In order to analyze the effects of retouching on faces from different demographics and test performance of retouching detection algorithms, we created a Multi-Demographic Retouched Faces (MDRF) database. The database has a total of 3600 images, where 1200 belong to original class and 2400 belong to retouched class; 1200 belong to each of the Chinese, Indian and Caucasian classes; and 1800 each to the male and female classes. Sample images from the database are presented in Figure 3. Further details about the tools and their retouching effects are provided next.

2.1. Retouching Tools

In order to capture the variance of retouching effects provided by various tools in the dataset, the retouched images were created using two tools. Both the tools have a different set of features or presets which lead to different effects.

BeautyPlus is a widely used (900 million users) [1] mobile phone application that helps users retouch pictures. It provides an array of features related to retouching such as skin smoothing, skin tone enhancement, acne removal, and face slimming [1]. For each original image, a set of effects to apply and the degree of application was randomly selected. This gives the liberty to the user to customize the

effect for each image but it acts as a limitation since processing each image by hand is a tedious task and hinders large-scale generation of retouched images.

Potraitpro Studio Max [11] is a paid service available as a stand-alone desktop application as well as plug-in for Adobe Photoshop. The tool detects the specific face features such as - eyebrows, eye corners, lips and nose centers and uses this information to provide different changes to each part. There are built-in presets to be used as filters on specific parts - eyes, lips, nose, or the entire face. These presets can change color of the eyes and lips, brighten the eyes, teeth and skin, sculpt (change the shape) of the jaw, forehead or the full face. The tool allows the creation of new presets using a mix and match of the existing ones. Once the effect has been determined, it can be applied to a set of faces using the batch processing mode. This helps us automatically create more than one retouched image for multiple individuals.

2.2. Distribution of Data

The images in the MDRF database belong to two categories - original and retouched, as defined below.

- **Original** - The database includes two frontal images for each of the 600 subjects from same session as original images. The original images have been acquired under controlled conditions and are known to have not been retouched. The database has balanced distribution of images for two genders (100 males and 100 females subjects for each ethnicity) and three ethnicities (200 subjects each).

  The Caucasian subjects were selected from the Collection B of Notre Dame database [20]. The original images for Chinese subjects were sampled from the visible spectrum images provided by the CASIA-NIR-vs-VIS (version 2) dataset [27]. The images for Indian subjects were selected from the CSCRV database [37].

- **Retouched** - The second image of each subject is used to create retouched images. Two tools were used to retouch images of all subjects - *BeautyPlus* and *Potraitpro Studio Max* (v12). For female subjects, we also created images with virtual makeup applied using an application similar to BeautyPlus, called MakeupPlus. For male subjects, two retouched images (with different combination of retouching options) were obtained from each of the tools while for female subjects, 2 images were obtained from BeautyPlus/MakeupPlus and 2 from Potrait Pro. This resulted in 2400 retouched images in total (4 images per subject).

3. Proposed Algorithm

The effect of race has also been a widely studied topic in the field of machine-based face recognition [21, 32, 34].
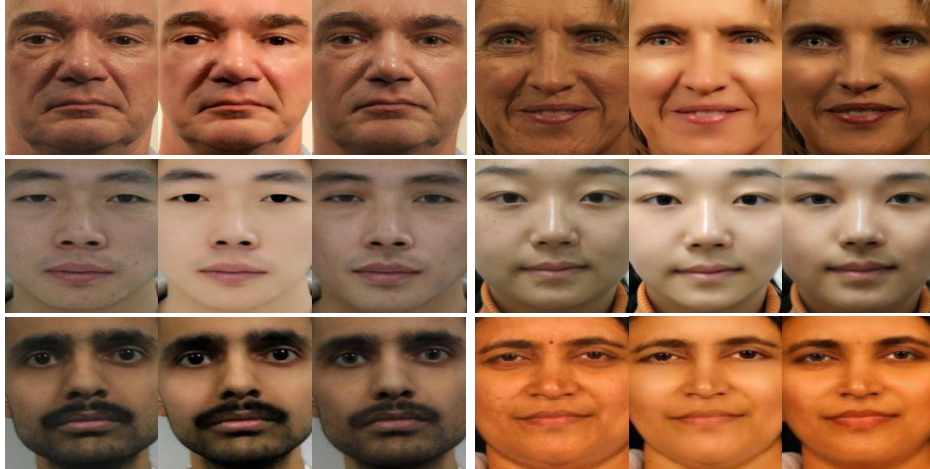
Figure 3: Sample images from the MDRF database. The first row consists images of Caucasian subjects, second row consists images of Chinese subjects and the third row consists of Indian subjects. For each of the six subjects, the first image is an original image, the second image is a retouched version of the first image using *BeautyPlus*, and the third image is a retouched version of the first image using *PotraitPro Studio Max*. The male images are on the left half and female images are on the right.

The existing literature has registered ethnicity as one of the important co-variates such as pose, illumination and expression. Notably, ethnicity co-exists with every other variate as it is a permanent attribute of a face unlike pose, illumination and expression which can be controlled. As for other face-related problems [18, 31], ethnicity of a face also becomes an element of consideration for facial retouching detection. The demographics also affect the perception of beauty in humans [5, 36]. Due to this, the desired effects might be different in different demographics [39], making it difficult for a detector built for subjects from one demographic to work with subjects from another. One method to make the system generalizable is to train the algorithm with 'enough' data from every ethnicity [22, 26] but the data is difficult to obtain given the different regulatory constraints across countries. Another possibility is to enable the algorithm to learn these variations properly and utilize the knowledge to learn better representations. To achieve this for our algorithm, we utilize a classic idea from pattern recognition of utilizing class hierarchy to learn better data distributions [23, 40].

The problem of retouching detection can be framed as a two-class classification problem where the two classes are original and retouched. In this paper, we extend the notion of the two classes (original and retouched) to be denoted by a hierarchy - classes (original and retouched) and subclasses (based on ethnicity and gender of the base face). The class hierarchy concept has been utilized in [23] to estimate a better (more discriminative) distribution of the underlying data using a mixture of Gaussians. This helps in capturing the variance of each of the sub-classes to further reduce the within class scatter and increase the between class variance. The work of Zhu and Martinez [40] also adopts the same

idea and allows the modeling of multiple distributions (depending on the number of subclasses) using the same formulation. It extends the prior work by proposing criteria to determine the number of subclasses in order to define the model.

For our problem, the original and retouched are the two major classes while Caucasian, Chinese and Indian are the sub-classes within the major classes. The branching in class hierarchy can depend on the type of variations being captured by the data. The idea is to capture the correct distribution of the two classes based on demographics and use this estimate to obtain features that can better distinguish between the two major classes. To learn a discriminative representation between the two major classes that considers the variance of distribution between subclasses, we inherit the subclass distribution modeling into a sparse autoencoder, termed as Subclass Supervised Sparse Autoencoder (S^3A).

3.1. Subclass Supervised Sparse Autoencoder

An autoencoder is an unsupervised method to learn features that well represent the input data. For this purpose, the autoencoder has two parts – the encoder and decoder. The encoder maps the input vector x to a hidden representation $h = \phi(Wx)$ while the decoder maps the hidden representation to the reconstructed output $\hat{x} = W'\phi(Wx)$. Here, W represents the weights between the input layer and the hidden layer and $\phi(.)$ is the activation function, usually a non-linear mapping function such as sigmoid. W' are the weights between the hidden layer and reconstruction output layer. The weights W and W' are optimized based on the following cost function:
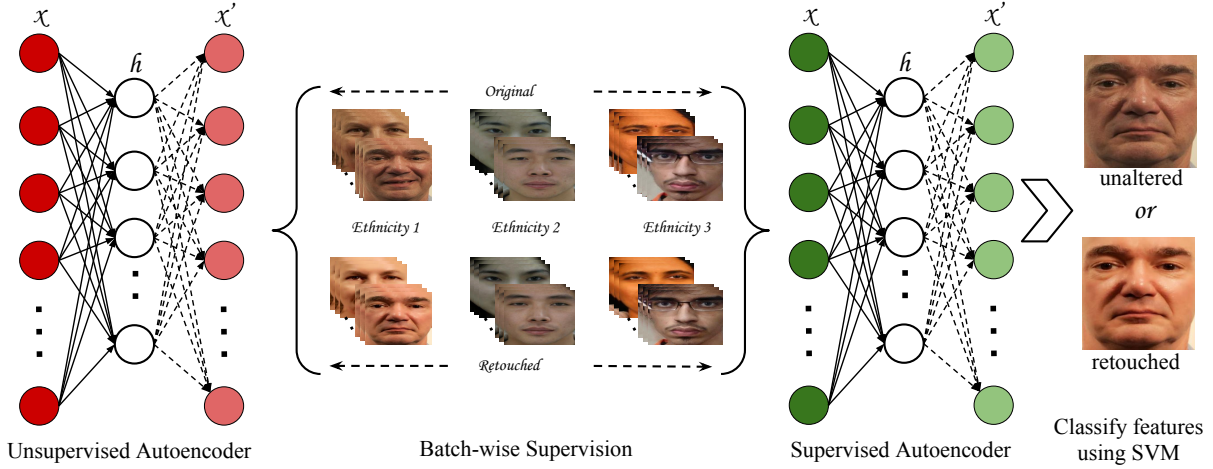
Figure 4: A diagram representing the steps involved in the proposed semi-supervised retouching detection algorithm.

$$\underset{W,W'}{\arg\min} ||X - W'\phi(WX)||_2^2 + \lambda R \quad (1)$$

The cost function is the squared form of the reconstruction error (first term in equation 1) over the entire training set regularized with a function R. The training set is denoted by X where $X = \{x_1, x_2, x_3, ...x_n\}$ and x_i is an individual training instance. For the autoencoder to learn sparse representations, we can add the sparsity constraint through the regularization function. Regularization is the technique used to avoid overfitting in networks and make the networks more generalizable. It is based on adding terms to the objective function to penalize large weights. The term includes a regularization parameter λ which regulates the amount of overfitting and a function $R(W, X)$ which together with λ adjusts the cost function depending on W. The two popular regularization functions are the l_2 norm and the l_1 norm. Due to the nature of l_1 norm, most of the weights are suppressed to 0 and only a small subset of them remain nonzero. This is how the learned representations are sparse and robust to noise.

To improve discriminability in the learned features, Majumdar *et al.* proposed Class Sparsity based Supervised Encoding (CSSE) [30]. The aim of this method is to encode samples from the same class with same sparsity signature. This requires the class information of the training samples, thus making the autoencoder supervised. CSSE is achieved by organizing the training data as $X = \{x_{11}, ...x_{1n_1}, x_{21}, ...x_{2n_2}, ...x_{c1}, ...x_{cn_c}\}$ where $x_{i1}...x_{in_i}$ are samples from the ith class and replacing the l_1 norm regularization function with a $l_{2,1}$ norm regularizer. The inner l_2 norm performs an aggregation within feature values while the outer l_1 norm encodes sparsity i.e., which features are enhanced for this particular set of samples. The main idea is to divide the data class-wise and then update the weights according to the sparsity constraint on each partition separately. This ensures the same level of sparsity in features belonging to the same class. The cost function of the CSSE is written as:

$$\underset{W,W'}{\arg\min} ||X - W'\phi(WX)||^2 + \lambda \sum_{c=1}^{C} ||WX_c||_{2,1} \quad (2)$$

This method of class sparsity encoding treats all the samples of a class in a similar fashion. For problems that have a defined class hierarchy, we extend the idea in [30] to include subclass information while doing supervised encoding. This is achieved by regularizing the cost function based on the batches for each subclass. Hence, the updated cost function can be written as follows.

$$\underset{W,W'}{\arg\min} ||X - W'\phi(WX)||^2 + \lambda \left[\sum_{i=1}^{c} \sum_{j=1}^{s_i} ||WX_{ij}||_{2,1} \right] \quad (3)$$

Equation 3 implies a two-level sparse encoding. The first level is for each subclass and the second level of sparsity imparts uniformity to the subclasses within a class. $l_{2,1}$ norm is non-differentiable but can be solved via approximation using Iterative Reweighted Least Squares (IRLS) approach. This uses the version of cost function in equation 4.

$$\underset{W,W'}{\arg\min} ||X - W'\phi(WX)||^2 + \lambda \left[\sum_{i=1}^{c} \sum_{j=1}^{s_i} ||\beta_{ij}WX_{ij}||_2 \right] \quad (4)$$

Here, β_{ij} is the set of parameters associated with each subclass. The weight update for each batch of data (samples within each subclass) is performed variably based on β_{ij}.

The value of β_{ij} is learned during training using M-Focuss algorithm [17]. M-Focuss is a version of FOCal Underdetermined System Solver (Focuss) for multiple measurement vectors (MMV). It is a method to form sparse solutions to problems having solutions that share sparsity by reducing diversity. Thus, for a particular subclass, the optimization method will try to find solutions in the same neighborhood of parameter space (i.e. similar weights).

3.2. Retouching Detection Using S^3A

For two-class classification problem of retouching detection {*original/unaltered*, *retouched*}, the cost function can be formulated as.

$$\underset{W,W'}{\operatorname{argmin}} ||X - W'\phi(WX)||^2 \\ + \lambda \left[\sum_{i=1}^{s_1} ||WX_i||_{2,1} + \sum_{j=1}^{s_2} ||WX_j||_{2,1} \right] \quad (5)$$

where, first term is a standard autoencoder and the second term encodes the subclass variations belonging to two classes. Here s_1 and s_2 are the number of subclasses in two classes, respectively.

As shown in Figure 4, the proposed algorithm is trained to determine if a given image is unaltered or retouched. For input of size MN, first, a two-hidden-layer network of size $[\frac{2}{3}MN, \frac{1}{2}MN]$ is pre-trained using normal face images from CMU Multi-PIE database [3] without using the subclass information. The network is fine-tuned using the training partition of MDRF database. For subclass based training, one experiment is performed with three ethnicity labels (Caucasian, Chinese, Indian) as three subclasses and another experiment is performed with gender (male, female) labels as two subclasses within each class. Note that, in this application, the subclasses are known, but if the subclasses are unknown then we can utilize subclass estimation techniques to label the training data with respect to subclasses. Once the model is trained, it is used to extract features for each image. Using the features in the training set, a 2ν-SVM classifier [15] is trained to perform two-class classification with the original class being considered positive and the retouched class as negative.

4. Experimental Results

Face detection and normalization are performed on the data to be used in all the experiments. For the purpose of face detection, the detector from Viola and Jones [38] is used. All the faces are tightly cropped and resized to 256×256. Once the images are pre-processed, they are used for analysis based on ethnicity of the trained retouching detection model and the analysis based on combined ethnicity trained models. Based on the different experimental protocols used, our experiments can be categorized into two parts.

4.1. Experimental Protocols

- **Combined Evaluation** – To compare the performances of various retouching detection algorithms, the dataset is divided into two subject disjoint sets with equal number of male and female subjects belonging to each of the three ethnicities. Therefore, there are images from 300 subjects (150 males and 150 females) in each set. One of the sets is used for training and the other is used for testing.

- **Cross Ethnicity Evaluation** – In order to analyze the differences of retouching among different ethnicities, we setup same-ethnicity training and testing, and other-ethnicity training and testing experiments. For training, 2 original images and 2 retouched images (one from Beauty Plus and one from Potrait Pro) are used for each subject in the set while all 6 images of remaining subjects are used for testing. The experiment uses 5-fold cross validation for same ethnicity and a 5×5-fold cross validation for different ethnicity. This implies that 5 trained models (trained with class balanced data) are obtained for each ethnicity using different combinations of the five folds. Each of the other remaining folds (full) for all three ethnicities are tested. In this way, we obtain 5 values of performance for same ethnicity and 25 for other ethnicities. An average of these performance values are reported along with standard deviation.

Comparison with state-of-the-art detection algorithms is performed on the MDRF database using the combined evaluation protocol. The specific algorithms used for comparison are explained below.

Makeup Detection – In order to analyze the difference between the makeup and retouching, we compare the performance of an existing state-of-the-art makeup detection algorithm [25] over the MDRF database. The algorithm uses Local Binary Pattern (LBP) and Histogram of Gradients (HOG) as features. The features upon normalization and concatenation are provided to two types of classifiers – SVM (linear and with RBF kernel) and Alligator (a combination of classifiers) for classification. We use our implementation of the method based on the explanation provided by the authors and report classification accuracy of the method in detecting retouching.

Photorealism Detection – A huge amount of manipulation of the original photograph can make the photo look unrealistic which may be utilized to detect retouching. One of the seminal works in detecting photorealism has been from

Lyu and Farid [29]. The authors use first-order and higher-order statistical characteristics of wavelets. Specifically, the wavelet decomposition of images is carried out using quadrature mirror filters. For each sub-band, at each scale and orientation, the mean, variance, skewness, and kurtosis are computed. These form the first half of the feature vector. The second half is comprised of statistics derived from the errors of linear predictor model that predicts magnitude of neighbouring coefficients. Once the features have been extracted from images, they are classified using SVM classifier. In this comparison, we consider the original class as photographic and the retouched class as photorealistic.

Spoofing Detection – Morphing is a common method of spoofing. A recent proposed solution to accurately detect morphed images uses BSIF and a linear SVM [35]. For the purposes of comparison, we use the same implementation [2] of BSIF as the authors and linear SVM from the libSVM package.

Facial Retouching Detection – The detection algorithms mentioned so far use hand-crafted features. In order to compare with algorithms that use deep learned features, we evaluate the features from VGG-Face - a pre-trained convolutional neural network (CNN) model for faces [33] with a SVM classifier. Another deep learning framework used for comparison is the facial retouching detection algorithm using Supervised Deep Boltzmann Machine (SDBM) and SVM by Bharati *et al.* [13]. They use 25,000 face images (original and retouched versions of images from Multi-PIE dataset) for training. For comparison purposes, the features are extracted using a trained SDBM model and a SVM is trained with features from the combined ethnicity training set of MDRF database.

4.2. Observations for Combined Evaluation

The results of the proposed algorithm for retouching detection along with other detection algorithms have been demonstrated in Figure 5. Some key observations based on the performance curves are:

- The specialized retouching detection algorithms perform better at retouching detection tasks than other detection algorithms. This shows that the other detection algorithms considered in this paper are not sufficient for generalized retouching detection.

- The proposed algorithm outperforms the existing retouching detection algorithm by Bharati *et al*. This demonstrates that incorporating subclass information can help estimate a better data distribution while learning representations. In this way, more discriminative (between classes) but generalizable (among subclasses) representations can be learned for each class using the proposed S^3A algorithm.
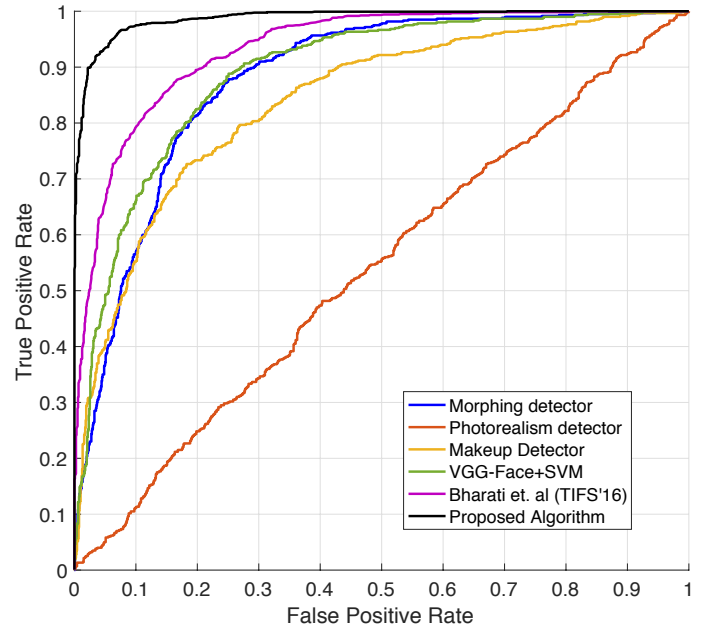

Figure 5: Performance curves of detection algorithms with combined evaluation protocol on MDRF database.

Table 1: Performance breakdown of retouching detection algorithms according to males and females and the two retouching tools.

	Female		Male	
Algorithm	Tool 1	Tool 2	Tool 1	Tool 2
Bharati *et al.* [13]	84.1%	81.3%	90.3%	85.9%
Proposed Algorithm	94.2%	92.8%	95.9%	94.3%

- In our results, descriptors learned from a supervised deep face recognition framework trained with VGG-Face are not as effective as the proposed algorithm in detecting retouching. As a part of the future work, this observation can be tested across other datasets and networks.

To provide a more comprehensive understanding of the performance of the proposed algorithm with respect to existing state-of-the-art retouching detection, the breakdown of accuracies based on gender and tool has been provided in Table 1. Tool 1 refers to BeautyPlus/MakeupPlus whereas Tool 2 refers to Potrait Pro Studio Max. The accuracy for male subjects is higher than the accuracy for female subjects. It is possible that, for females, the original face image might have pre-applied makeup which can lead to confusion between original and retouched class. According to the results, retouching for Tool 2 is slightly more difficult to detect than Tool 1.

Table 2: Results of cross ethnicity evaluation using VGG-Face+SVM, Bharati *et al.* [13] and the proposed algorithm.

			Testing Set		
		Algorithm	Caucasian	Chinese	Indian
Training Set	Caucasian	VGG+SVM	79.3% (± 3.1%)	81.1% (± 3.2%)	70.7% (± 2.5%)
		Bharati *et al.*	84.2% (± 2.9%)	80.9% (± 3.6%)	76.3% (± 4.1%)
		Proposed	94.3% (± 1.1%)	91.9% (± 1.8%)	92.2% (± 2.1%)
	Chinese	VGG+SVM	73.8% (± 2.5%)	87.0% (± 1.4%)	70.9% (± 4.6%)
		Bharati *et al.*	79.4% (± 3.2%)	84.7% (± 2.2%)	80.1% (± 3.3%)
		Proposed	91.7% (± 2.3%)	97.5% (± 1.2%)	92.8% (± 2.2%)
	Indian	VGG+SVM	77.5% (± 2.8%)	81.4% (± 4.8%)	73.0% (± 4.1%)
		Bharati *et al.*	78.4% (± 3.7%)	79.3% (± 4.1%)	85.2% (± 2.9%)
		Proposed	92.1% (± 2.3%)	93.6% (± 1.9%)	96.2% (± 1.4%)

4.3. Observations for Cross-Ethnicity Evaluation

In order to analyze the performance for models trained on data from one ethnicity and tested on data from same ethnicity and different ethnicities, we performed experiments using the cross ethnicity evaluation protocol. Since we have 3 ethnicities in the MDRF database, we have 9 scenarios. The results for each of the scenarios are reported in Table 2. The performance is reported for the 3 best performing algorithms in the combined evaluation. For each scenario and algorithm, the mean and standard deviation for classification accuracies of all trials have been reported.

According to the results of various scenarios, the following can be observed:

- As expected, the performance when the model is trained and tested on the same ethnicity (cells along the diagonal in table 2) is generally higher than when the model is trained and tested on different ethnicities.

- For all the scenarios, the proposed algorithm performs better than VGG+SVM and Bharati *et al.* [13]. This shows that subclass sparsity encoding of gender helps in distinguishing between original and retouched images within a particular ethnicity.

- The standard deviation of accuracies across all folds and across all testing scenarios for each trained model is less for the proposed algorithm. This implies that the proposed technique generalizes better for multi-demographic retouching detection.

- Of all the cross ethnicity evaluations (non-diagonal cells in Table 2), retouching is easier to detect in Chinese faces while it is more difficult to detect in Indians.

The results presented in the paper show that the proposed algorithm improves upon the existing state-of-the-art by improving generalizability across demographics. The results also explain the relative difficulty in retouching detection among the three ethnicities - Caucasian, Chinese and Indian.

5. Conclusion & Future Work

The paper presents the limitations of state-of-the-art algorithms in detecting retouching for multiple demographics. To achieve this, a novel dataset with subjects from three ethnicities and retouching from two tools is introduced. Combined ethnicity and cross-ethnicity evaluations were performed to obtain a detailed understanding of performance in scenarios with one or more ethnicities available for training. Finally, a novel semi-supervised framework with Subclass Supervised Sparse Autoencoder (S^3A) is proposed to improve detection of retouching across ethnicity and gender.

Some possible avenues of future research include analyzing differences among tools from various demographic area and to create generalized algorithms. Also, quantifying how much retouching is acceptable and beyond what levels is retouching considered spoofing can be beneficial for ensuring standards for images.

References

[1] BeautyPlus. http://www.beautyplus.com/EN/. Accessed: 2017-03-30.

[2] BSIF code. http://www.ee.oulu.fi/ jkannala/bsif/bsif.html. Accessed: 2017-04-30.

[3] CMU's Multi-PIE database. http://www.cs.cmu.edu/afs/cs/project/PIE/MultiPie/Multi-Pie/Home.html. Accessed: 2017-04-15.

[4] Dating Websites Reviews. http://www.datingsitesreviews.com/article.php?story=new-study-reveals-trend-in-photo-retouching-among-online-daters. Accessed: 2017-04-15.

[5] Esther Honig's Before & After project. http://time.com/2934434/this-is-what-the-same-woman-looks-like-photoshopped-in-different-countries/. Accessed: 2017-04-20.

[6] Fotor. http://www.fotor.com/. Accessed: 2017-03-30.

[7] French law on photoshopped images. https://mic.com/articles/130789/france-passes-law-requiring-companies-to-admit-when-models-have-been-photoshopped. Accessed: 2017-04-15.

[8] Hany Farid's Interview. http://www.glamour.com/story/retouching-how-much-is-too-much. Accessed: 2017-04-20.

[9] Israel's Photoshop Law. http://www.ibtimes.com/supermodels-without-photoshop-israels-photoshop-law-puts-focus-digitally-altered-images-photos. Accessed: 2017-04-15.

[10] Photoshop Laws. https://www.pixelz.com/blog/photoshop-models-laws/. Accessed: 2017-04-15.

[11] PotraitPro Studio Max v12. http://www.portraitprofessional.com/#. Accessed: 2017-04-15.

[12] Self Esteem Act. http://www.dailymail.co.uk/femail/article-2048375/Self-Esteem-Act-US-parents-push-anti-Photoshop-laws-advertising.html. Accessed: 2017-04-15.

[13] A. Bharati, R. Singh, M. Vatsa, and K. W. Bowyer. Detecting facial retouching using supervised deep learning. *IEEE Transactions on Information Forensics and Security*, 11(9):1903–1913, 2016.

[14] C. Chen, A. Dantcheva, and A. Ross. Automatic facial makeup detection with application in face recognition. In *International Conference on Biometrics (ICB)*, pages 1–8. IEEE, 2013.

[15] H.-G. Chew, C.-C. Lim, and R. E. Bogner. An implementation of training dual-nu support vector machines. In *Optimization and Control With Applications*, pages 157–182. Springer, 2005.

[16] V. Coetzee, J. M. Greeff, I. D. Stephen, and D. I. Perrett. Cross-cultural agreement in facial attractiveness preferences: The role of ethnicity and gender. *PloS one*, 9(7):e99629, 2014.

[17] S. F. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado. Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Transactions on Signal Processing*, 53(7):2477–2488, 2005.

[18] H. Dehon and S. Brédart. An other-race effect in age estimation from faces. *Perception*, 30(9):1107–1113, 2001.

[19] M. Ferrara, A. Franco, D. Maltoni, and Y. Sun. On the impact of alterations on face photo recognition accuracy. In *International Conference on Image Analysis and Processing*, pages 743–751. Springer, 2013.

[20] P. Flynn, K. Bowyer, and P. Phillips. Assessment of time dependency in face recognition: An initial study. In *Audio-and Video-Based Biometric Person Authentication*, pages 1057–1057. Springer, 2003.

[21] L. Ge, H. Zhang, Z. Wang, P. C. Quinn, O. Pascalis, D. Kelly, A. Slater, J. Tian, and K. Lee. Two faces of the other-race effect: Recognition and categorisation of caucasian and chinese faces. *Perception*, 38(8):1199–1210, 2009.

[22] A. G. Goldstein and J. E. Chance. Effects of training on japanese face recognition: Reduction of the other-race effect. *Bulletin of the Psychonomic Society*, 23(3):211–214, 1985.

[23] T. Hastie and R. Tibshirani. Discriminant analysis by gaussian mixtures. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 155–176, 1996.

[24] E. Kee and H. Farid. A perceptual metric for photo retouching. *Proceedings of the National Academy of Sciences*, 108(50):19907–19912, 2011.

[25] N. Kose, L. Apvrille, and J.-L. Dugelay. Facial makeup detection technique based on texture and shape analysis. In *11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 1, pages 1–7. IEEE, 2015.

[26] S. Lebrecht, L. J. Pierce, M. J. Tarr, and J. W. Tanaka. Perceptual other-race training reduces implicit racial bias. *PLoS one*, 4(1):e4215, 2009.

[27] S. Li, D. Yi, Z. Lei, and S. Liao. The CASIA NIR-VIS 2.0 face database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 348–353, 2013.

[28] A. C. Little, B. C. Jones, and L. M. DeBruine. Facial attractiveness: evolutionary based research. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1571):1638–1659, 2011.

[29] S. Lyu and H. Farid. How realistic is photorealistic? *IEEE Transactions on Signal Processing*, 53(2):845–850, 2005.

[30] A. Majumdar, R. Singh, and M. Vatsa. Face recognition via class sparsity based supervised encoding. *IEEE transactions on pattern analysis and machine intelligence*, 2016.

[31] A. J. O'Toole, J. Peterson, and K. A. Deffenbacher. An other-race effect for categorizing faces by sex. *Perception*, 25(6):669–676, 1996.

[32] A. J. Otoole, K. A. Deffenbacher, D. Valentin, and H. Abdi. Structural aspects of face recognition and the other-race effect. *Memory & Cognition*, 22(2):208–224, 1994.

[33] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *BMVC*, volume 1, page 6, 2015.

[34] P. J. Phillips, F. Jiang, A. Narvekar, J. Ayyad, and A. J. O'Toole. An other-race effect for face recognition algorithms. *ACM Transactions on Applied Perception (TAP)*, 8(2):14, 2011.

[35] R. Raghavendra, K. B. Raja, and C. Busch. Detecting morphed face images. In *International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–7. IEEE, 2016.

[36] G. Rhodes. The evolutionary psychology of facial beauty. *Annu. Rev. Psychol.*, 57:199–226, 2006.

[37] M. Singh, S. Nagpal, N. Gupta, S. Gupta, S. Ghosh, R. Singh, and M. Vatsa. Cross-spectral cross-resolution video database for face recognition. In *8th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–7. IEEE, 2016.

[38] P. Viola and M. J. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 57(2):137–154, 2004.

[39] L. A. Zebrowitz, J. M. Montepare, and H. K. Lee. They don't all look alike: Individual impressions of other racial groups. *Journal of personality and social psychology*, 65(1):85, 1993.

[40] M. Zhu and A. M. Martinez. Subclass discriminant analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(8):1274–1286, 2006.