

Achieving Anonymity against Major Face Recognition Algorithms

Benedikt Driessen and Markus Dürmuth

Ruhr-University Bochum

{benedikt.driessen,markus.duermuth}@rub.de

Abstract. An ever-increasing number of personal photos is stored online. This trend can be problematic, because face recognition software can undermine user privacy in unexpected ways. Face de-identification aims to prevent automatic recognition of faces thus improving user privacy, but previous work alters the image in a way that makes them indistinguishable for both computers and humans, which prevents a wide-spread use.

We propose a method for de-identification of images that effectively prevents face recognition software (using the most popular and effective algorithms) from identifying people, but still allows human recognition. We evaluate our method experimentally by adapting the CSU framework and using the FERET database. We show that we are able to achieve strong de-identification while maintaining reasonable image quality.

Keywords: Face recognition, anonymity, de-identification.

1 Introduction

The number of personal photos that is available online has been rapidly increasing over the past years.¹ This development is driven by the wide availability of (stationary and mobile) high-speed Internet, cheap electronic storage, and ubiquitous digital cameras on the one hand, and a strong trend towards social networks and managing friends online on the other hand. Recently, some services publicly announced the deployment of face recognition software on the stored images (see, e.g., [12]). Face recognition software can be beneficial for the user, as it helps finding and tagging friends in pictures. However, it can also be used to gather additional information about friendship-relations (i.e., the social graph), even relations the user deliberately did not share with everybody.

People often try to separate some groups of people, e.g., personal friends and work colleagues and reserve an online profile for personal friends. Note that Facebook considers the social graph as public information that can even be queried via a special API [13], and in general one can easily imagine external services crawling the image database. A rather drastic example includes predicting a

¹ As an example, Facebook hosted 10 billion photos in Oct. 2008, receiving about 250 million new photos a day [11]. Flickr hosted 4 billion photos in Oct. 2009 [14].

persons sexual orientation from the social graph [21]. So a person might want to hide (parts of) their social graph in order to protect such information, but still might want to post images. (We stress that automated extraction of the social graph is a much bigger problem than manual extraction, because of the large-scale extraction of information that becomes possible.)

In the past few months, the criticism of automated face recognition, and in particular Facebook as the most prominent example, has increased, e.g., from the Electronic Privacy Information Center (EPIC), which considers Facebook’s handling of personal data a violation of European privacy law (see, e.g., [8,10]).

In this work we will demonstrate a system that effectively protects the anonymity, in particular the social graph, of a user by thwarting face recognition software, while still allowing humans to identify faces and keeping the visual changes to the image small. In particular, previous work [23] mapped several “similar” faces to the same “average”-face (see Section 2 for details), thus the resulting face-images are indistinguishable for both computers *and* humans. While their approach provides strong security guarantees, it constitutes a strong visual alteration of the face images, and in particular prevents humans from distinguishing persons. We believe that this is too intrusive for wide-spread use.

1.1 Our Contribution

We propose a system to anonymize face images in a way that retains more details of the original image than previous work, thus allowing a human to still identify the person from the image, and works against all major classes of face recognition algorithms.

We exploit the fact that face recognition algorithms reduce the dimensionality of the data of the face image in order to reduce noise and improve speed. We show how to manipulate these relevant parts of the image to fool all important face recognition algorithms. We aim at a slightly weaker form of anonymity than k -anonymity (see Section 4), where an attacker’s confidence in correct identification is small. This is well suited for the two main threats we are considering: first, we want to prevent extraction of the relationships (the social graph), e.g., by evaluating who is present on a sufficiently large number of pictures of one specific person; second, we want to prevent automated tagging of people on pictures.

Of course, when humans can recognize a face, then the face will not be anonymous in a strict sense. Studies show [2] that the “price” most users are willing to pay for privacy is pretty low, so we hope that by providing a reasonable image quality, this approach will find more acceptance by users, while still preventing automated extraction of information and thus providing a reasonable level of security.

2 Eigenface-Based Face Recognition

We introduce the basic terminology, face recognition algorithms based on Principal Component Analysis (PCA), and show the basic idea how we can manipulate

these algorithms. We consider the classical Eigenface-algorithm [32], which is interesting because it constituted a breakthrough in recognition performance at the time of its invention, and still provides very competitive performance for images taken in a moderately controlled environment. Also, it forms the basis for a wide range of algorithms, including Linear Discriminant Analysis (LDA), which can be applied after PCA, and the Bayesian classifier, both of which we present also.

2.1 Brief Introduction to Face Recognition Terminology

The task of face recognition can be described in simple terms: Given a set of images of a number of (known) persons (*gallery images*), and given an image for an unknown person (*probe image*), decide which person from the gallery is shown on the probe image.

The process is usually divided into the following steps: First, one needs to find the approximate position of the face in the image, this is called *face-detection* [35] and a separate line of research. Most work on face recognition considers this job to be completed before². Second, images are *normalized*, which usually includes an affine transformation to align the eyes, histogram equalization, and sometimes masking of the background. Third, in *feature extraction* algorithm-dependent features the probe image are extracted. Representing an image by a set of features can be seen as a step in data reduction that aims at extracting a compact but discriminating description of the image. Ideally, the output of this step is at the same time robust against changes in posture, lightning, face expression, etc. Finally, the pre-processed probe image is *matched* against gallery images. Different distance functions, measuring the similarity of two images, can be used here, we will see Euclidean distance, weighted Euclidean distance, and angles in the sequel.

The output of a face recognition algorithm is a list of identifiers, where the algorithm estimates that the first identifier (e.g. name) is the most likely one, matching the subject on the probe image. The list is typically ordered by descending probability (or ascending distance, depending on the distance measure). The performance of face recognition algorithms is usually measured in *rank curves* (see Figure 1), where on the x -axis we plot ranks and on the y -axis recognition rates. When the curve passes point (x, y) this means that for a fraction of y probe images, the correct identifier was contained in the first x suggestions of the recognition algorithm.

2.2 Classical Eigenfaces

Given L training pictures $\mathbf{u}_1, \dots, \mathbf{u}_L$, written as vectors with p pixels each, we compute the (point-wise) average image $\mathbf{m} := \frac{1}{L} \sum_{k=1}^L \mathbf{u}_k$ and compute the mean-subtracted images $\mathbf{u}'_k := \mathbf{u}_k - \mathbf{m}$. Write $U = [\mathbf{u}'_1, \dots, \mathbf{u}'_L]^t$ for the matrix

² Commonly used face image databases such as the FERET database (see Section 5.2) annotate the images with the eye coordinates.

of images. Then the (empirical) covariance matrix is written as $C = U^t \cdot U$. The matrix U has size $p \times L$ and the matrix C has size $p \times p$. Let $\lambda_0, \dots, \lambda_{N-1}$ be the N largest Eigenvalues with associated Eigenvectors $\mathbf{e}_0, \dots, \mathbf{e}_{N-1}$. Using the orthonormal vectors $\mathbf{e}_0, \dots, \mathbf{e}_{N-1}$, we calculate a feature vector (in what is called *face-space*, the space spanned by Eigenvectors), by simply projecting an image \mathbf{u} on the Eigenvectors

$$s_k = \mathbf{e}_k^t \cdot (\mathbf{u} - \mathbf{m}) \quad \text{for } k = 0, \dots, N - 1, \quad (1)$$

where s_k is the k -th component of the projection of \mathbf{u} in face-space. Calculating a feature vector for every gallery image as well as the probe image, and using Euclidean distance to find the gallery image that is closest to the probe image, we get the original Eigenfaces face recognition algorithm.

Instead of Euclidean distance we can also use different distance measures, such as MahCosine, which measures the angle of the projections \mathbf{m}, \mathbf{n} of two vectors \mathbf{u}, \mathbf{v} into Mahalanobis space see [4] for details and more examples. In the following we use this measure because it outperformed other measures for purely PCA-based recognition methods [18] and in order to enhance our understanding of the effectiveness of de-identification methods in the presence of non-Euclidean distance measures.

2.3 Modifying Projections (in Image Space)

Our basic idea is to manipulate the face images in such a way that the projection onto the face-space changes, while hopefully making minimal changes.

Given the input image \mathbf{u} as a row-vector and a set of orthogonal and normalized vectors $\mathbf{e}_0, \dots, \mathbf{e}_{N-1}$ (the selected principal eigenvectors as computed by the PCA, spanning the face-space), consider Equation (1). By adding $\Delta_k \cdot \mathbf{e}_k$, a multiple of the k -th Eigenvector, to an image \mathbf{u} , we can arbitrarily change the k -th component of the projection; an easy calculation shows that

$$((\Delta_k \cdot \mathbf{e}_k^t + \mathbf{u}^t) - \mathbf{m}^t) \cdot \mathbf{e}_k = \Delta_k + s_k. \quad (2)$$

Another easy calculation shows that we can adjust several components by simply adding several Eigenfaces, because the $\mathbf{e}_1, \dots, \mathbf{e}_{N-1}$ are pairwise orthogonal.

Compared with previous work [23], the image quality of this approach is better, because we do not alter the information outside the Eigenspace.

2.4 PCA+LDA and Bayesian Face Recognition

LDA can be applied directly to the input data [1], but was found to be more effective when applied after a PCA transform [37]. While slightly different parameters are optimal for the PCA transform when used as preprocessing stage for LDA, one can see that the same techniques that we developed in Section 2.3 is applicable.

The Bayesian face recognition algorithm [19,31] is different from most other algorithms in that it breaks down face recognition into a series of classification

problems: in order to recognize a face, the algorithm iterates over all stored *persons* (not faces), and for each decides if it is the correct person or not. The central idea is that it tries to decide if the difference of two faces is in one of two classes, either *inter-personal* (Ω_I) or *extra-personal* (Ω_E), where in preprocessing the algorithm learns what are “typical variations” for the difference of two images of the same face, and for two images of different faces. In this training, again, PCA is applied to the input to reduce dimensionality, with the difference that here it is applied to the (point-wise) difference of two images. This yields Eigenfaces different from the above, but a closer look reveals that again methods similar to those from Section 2.3 can be applied.

3 Elastic Bunch Graph Mapping Face Recognition

Elastic bunch graph mapping [33] is another algorithm for face recognition that fared very well in the FERET tests [28,24]. What makes this algorithm particularly interesting is that it is fundamentally different from the previous algorithms: it is not based on PCA and is commonly classified as feature-based instead of holistic, i.e., it bases its decision on particular local features (eyes, mouth, ...) instead of a holistic view of the face.

3.1 Algorithm Description

We give a brief overview of the algorithm, for more details we refer the reader to [33,5]. A central tool for EBGGM are *Gabor wavelets*, convolution kernels which are plane waves bounded by a Gaussian envelope function. Let ψ_j be a Gabor wavelet, then the convolution at point x with the image \mathbf{u} is given by

$$J_j(x) = \int \mathbf{u}(x')\psi_j(x' - x)dx', \quad (3)$$

where $J_j(x)$ is a complex value and the index j ranges over 40 values for 8 orientations and 5 frequencies. Convolution of a fixed point of an image with Gabor wavelets of different orientation and frequency is called a *jet*; intuitively, a jet contains a reduced description of the surrounding of that point. Gabor wavelets are robust against a number of variations and are motivated by human vision research.

For the faces, one defines a set of *fiducial points*, such as pupils, corners of the eyes or mouth, and top and bottom of the ears. The nodes of these graphs are labeled with a jet. Initially, for a small set of faces, the fiducial points are extracted by hand, and the jets are computed. When presented with a new face, the information extracted above is used to automatically fit the above graph to a new face: First, the rough position of the face is determined by matching the average of all above graphs onto the probe image. Then the graph that fits best is selected, allowing for small displacements and scaling of the graph, followed by successively relaxing the graph geometry and adapting the points individually.

The graph is fitted on every image in the gallery, and the resulting vector of jets is stored. For a probe image, the closest match with a gallery image is computed as the mean of the individual jet similarities, where jet similarities are computed as “normalized vector product”.

3.2 Modifying Jets (in Image Space)

Our basic idea is that by adding appropriate multiples of a wavelet $\alpha \cdot \psi_j$ to an image \mathbf{u} at position x , we can change the value of the convolution with this particular wavelet at the specific position. This can be verified by a simple calculation.

One difference to the situation for PCA is that these changes are not independent of each other: modifying one jet value also changes other values for this jet, and several jets are close enough that other jets are influenced as well. For this reason we proceed iteratively as follows:

1. Do 150 times, over all jets and wavelets:
 - (a) Find the maximum difference between the current and target value
 - (b) Add Gabor wavelet to bridge 1/5-th of the distance
2. Over all jets and wavelets (wavelets with large radius first):
 - (a) Add Gabor wavelet to bridge 1/20-th of the distance

We established these parameters empirically and found them to work well. As for PCA-based techniques, it is not necessary to set the image to be equal to the target image, because probe and gallery images have a certain distance anyway.

4 Achieving Anonymity

Next, we describe how we utilize the approaches from the previous sections to anonymize face images.

4.1 k -Anonymity

An established definition of security against identification is *k-anonymity* [29,30], see also the notion of *anonymity sets* [25]. For face recognition, a person remains *k*-anonymous if the face recognition algorithms cannot narrow the person down to a set of less than *k* persons.

For our envisioned targets, weaker forms of anonymity are sufficient. There are two scenarios we would like to protect against: First, automated tagging of persons on uploaded pictures (note that we are not targeting the automated proposal of persons to tag, because this still contains human interaction and thus is only making tagging easier...); second, automated derivation of friendship-relations from a large set of pictures. For both scenarios, weaker privacy guarantees suffice. Along with this weaker privacy guarantee comes a large improvement in image quality (as perceived by a human), so we hope that our system will lead to more wide-spread usage of privacy-protecting systems.

4.2 Anonymizing Face Images

Here we describe our approach to face anonymization, which builds on the methods we have developed above.

1. We select a partition of the involved persons such that each set has at least k members. We choose them by picking a random face image and selecting the $k - 1$ nearest images (of distinct persons) according to a suitable distance measure (we will elaborate on this in Section 5) and we call this set a *cluster*.
2. For every cluster we project each of the k images and compute the average projection (wrt. PCA).
3. All images in a cluster are modified (as described in Section 2.3) to have the same, averaged projection. However, we may also choose to adjust the projection by a fraction $\sigma \in \mathbb{R}$ with $0 < \sigma < 1$ of the difference between an images actual projection and the average projection. The parameter σ was determined experimentally, see Section 5.4.
4. Next, all images in a cluster are modified wrt. EBGM (as described in Section 3.2) to resemble the average face for that cluster. We apply EBGM-modifications after PCA-modifications, because EBGM-modifications are local, whereas PCA-changes influence the entire image (see classification of EBGM as a local feature-approach, as opposed to PCA being a holistic approach).

A central observation is that we do not need to change the projections of PCA in face-space to the actual average, but it's sufficient to go some way in that direction. The reason is that probe image and gallery image of the same person are already quite some distant apart, so moving partially in the correct direction suffices. We will show experiments substantiating this claim.

5 Experiments

In this section we present extensive experiments that substantiate our privacy claims. We used the CSU framework of face recognition algorithms to test our results on a subset of 1000 images of the FERET database. We performed de-identification experiments for all three classes of algorithms (Eigenface-based, Bayesian, EBGM) and finally realized a synthesis of our results. The FERET database contains images of faces only. Thus our experiments omit some pre-processing steps that would be required to apply this idea in a real-world scenario, but these are well-known and add little insight to our goal of de-identification.

5.1 The CSU Framework

The CSU framework [6,4] was created at Colorado State University to facilitate the comparison of different algorithms, and is available for free for research. The framework runs on UNIX/Linux systems, the source code is available and therefore easily adaptable. The current version 5.1, published July 2010, supports the

following face recognition algorithms: (i) Classical Eigenfaces (i.e., PCA) with different distance measures, e.g. Euclidean, MahCosine, etc., (ii) LDA+PCA, also with different measures, (iii) Bayesian classification with the MAP and ML classifier, and (iv) Elastic Bunch Graph Mapping. Details about the specific implementations can be found in a series of papers, most notably [31,5,34,3]. The framework utilizes the FERET dataset and allows to easily measure the recognition performance. Furthermore, due to its modularity, the framework can be extended by new algorithms in order to benchmark these against already known methods. Although not intended, the framework can easily be adapted to allow benchmarking the de-identification performance of our algorithms.

5.2 The FERET Database

The FERET program [26,27] started 1993 and ran until 1997. It was sponsored by the Department of Defense Counterdrug Technology Development Program through the Defense Advanced Research Products Agency. Its primary mission was to develop automatic face recognition algorithms that could be employed to assist security, intelligence and law enforcement personnel. The FERET dataset was assembled to support government monitored testing and evaluation of face recognition algorithms using standardized tests and procedures. The final set of images has 3300 images from 1200 persons, with varying mimical expressions, from different dates, under semi-controlled conditions. The dataset is available for research related to face recognition.

5.3 Scope and Conduct of Experiments

For our experiments we used the FA and FB subsets of the FERET database, each containing one facial expression of 1195 subjects. We ran the experiments on a random subsets of 500 subjects each. The FA set served as our gallery, face recognition performance figures are given in terms of successively matching subjects from the probe set FB to their alternate image in FA. In our experiments, we apply our de-identification methods against both sets of images which implies that, although all gallery images have been de-identified, the face recognition system still has identifiers associated to them. This is not only a necessary prerequisite for identification, but also realistic when considering that most social networks encourage users to manually identify persons on photos, thus adding images to the gallery. Several experiments were performed to validate our de-identification approaches and identify a reasonable set of parameters. First, we tested our de-identification method for each recognition algorithm individually; the results are shown in Section 5.4 and Section 5.5. Then we combined these preliminary findings, applying both techniques at the same time, for our final results in Section 5.6.

All experiments were conducted with the default configuration of the CSU framework, the only exception concerns the normalization step. EBGM uses a pre-processing procedure which is different from the normalization required by all other recognition methods. Our experiments were conducted entirely on the

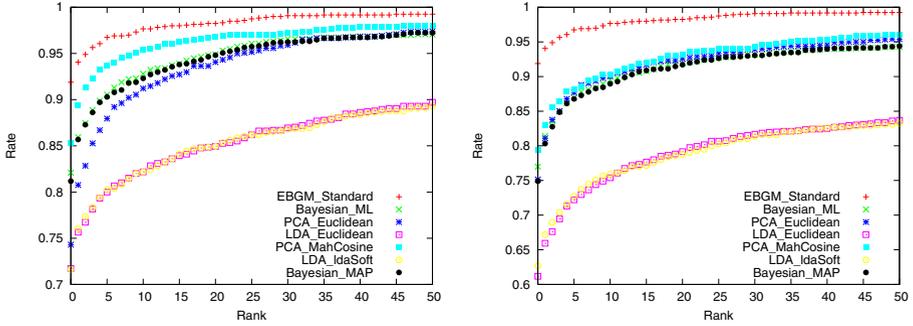


Fig. 1. Baseline performances of the algorithms (prior to de-identification), in the first image with the original normalization parameters, in the second image with identical normalization for all algorithms

data-set generated by EBGM pre-processing. The reason is our two-fold strategy for de-identification. Ideally, all PCA-related modifications are applied to images normalized for PCA-based methods, then the result is transformed back to the original image which is subsequently pre-processed for EBGM. Finally, the EBGM modifications are applied. However, this is tedious work that adds little insight to the interesting questions, so we chose to use the EBGM data-set for all methods. The impact of this strategy on recognition rates is shown in Figure 1. It displays the rank curves for all of the algorithms targeted by us. The left graph shows recognition performance (prior to de-identification) in the default configuration, i.e. where each class of algorithms operated on specifically normalized images. The right graph shows a slight degradation in recognition performance, which is due to non-optimal normalization, which seems to affect LDA-based methods most. The performance results of our de-identification methods, which are expressed as rank curves as well, are to be understood in relation to the right graph.

5.4 Experiments for Eigenface-Based Face Recognition

Experiments to determine the effectiveness of our de-identification method against Eigenface-based methods are parameterized by k , which is the size of the anonymity clusters, and $0 < \sigma < 1$ which is a factor weighting the addition of modifications. In a first series of experiments, presented in this section, we tested and validated our de-identification method against face recognition using PCA and LDA+PCA. For both, we tested different distance measures, which gives us assurance that the proposed method is robust against changes in this metric.

In a first series of experiments, we targeted each of these methods individually by building clusters of persons according to the same face recognition method, because we wanted to learn how sensitive our method is to how exactly the clusters are chosen. For our first experiment we consider de-identification with

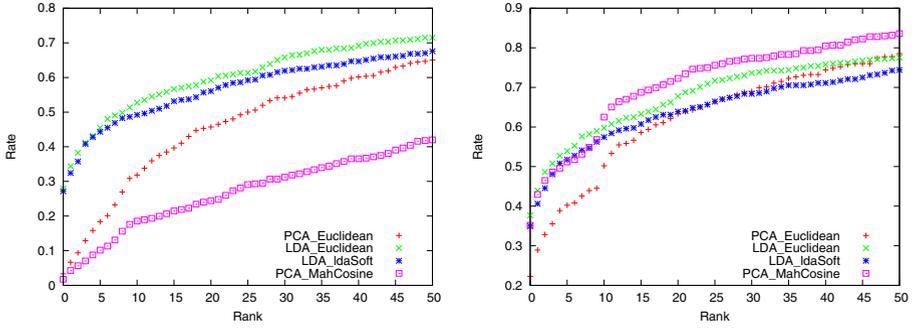


Fig. 2. Eigenface-based recognition after de-identification with $k = 10$ and $\sigma \in \{1, 0.85\}$ clusters obtained from the same measure



Fig. 3. Comparison of original image (left) with image after de-identification (right) with parameters $k = 10, \sigma = 0.85$

$\sigma = 1, k = 10$ and measure the recognition rate of all four methods, again. Comparing the left graph in Figure 2 with the original results from Figure 1, we see that the two PCA algorithms perform quite similar with an extremely low rank-0 recognition rate, and the the PCA+LDA algorithms perform better, yet still much lower than without de-identification. Also, we can see that the performance of the MahCosine distance measure, which is very accurate without de-identification, decreases disproportionately strong.

In a second series of experiments, we determined a suitable weighting factor σ . The lower we choose σ , the less an image is actually altered (thus not completely bridging the distance to the cluster’s target image) which consequently yields a better image quality. We found that de-identification with $k = 10$ and $\sigma = 0.85$ works well, and additionally this parameter balances the recognition rate for the four algorithms, as can be seen in the right graph in Figure 2. Figure 3 shows the visual effects of de-identification for $k = 10, \sigma = 0.85$ for all four recognition methods in comparison with the original image. In both cases the person is clearly recognizable. The strongest effect on the pictures can be seen at the line between the person’s hair and the background. This effect hardly affects the recognizability of a face, and can most likely be avoided by restricting the Eigenfaces to the actual face, as in the usual preprocessing for PCA. Also, the modified images look somewhat lighter than the original images.

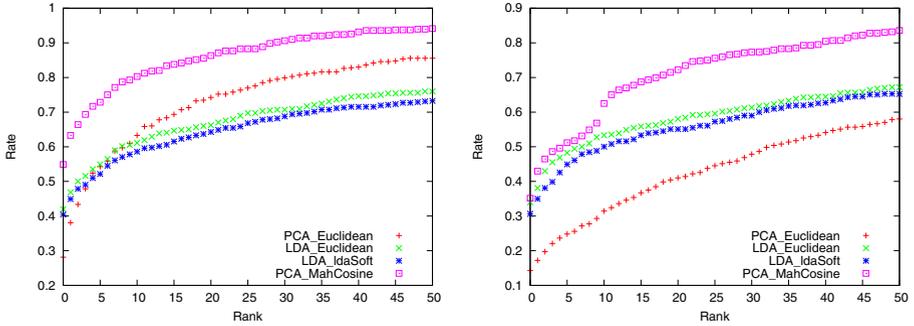


Fig. 4. Eigenface-based recognition after de-identification of the same cluster with $k = 10$, $\sigma = 0.85$, using averaging (top) and MahCosine (bottom).

In the previous experiments, we have selected the image clusters using the same distance measure as in the recognition task. This is not a realistic option in practice, so in a third series of experiments, we determined the best approach to de-identify images using a single clustering. We tested two approaches:

- Compute the four de-identified images for one subject, each grouped by one of the Eigenface-based measures, and average these pixel-wise.
- Compute clusters using a single distance measure, here we used the MahCosine measure (which performed best for $\sigma = 0.85$).

The results for both approaches are shown in Figure 4. The graph on the left shows that averaging over the four de-identified images per subject yields rather bad results, all algorithms have a rank-0 recognition rate of 40%-55%. The likely reason for this is that the clusters used to compute each of these projections were different, and the right image lies in the intersection between these. The right graph shows much better de-identification and is our preferred method. The recognition rate of the classical Eigenfaces method is worse than in the case were we *specifically* targeted this method, see Figure 2. The curves of all other three algorithms very much resemble the already known results as in the afore mentioned graph. What is more important, they are still worse than the performance of the MahCosine distance measure, which serves as our benchmark in this case.

We conclude that clustering wrt. MahCosine achieves a high degree of de-identification among all tested Eigenface-based methods while only minimally impacting identification by humans. We expect that other distance measures exhibit similar performance.

5.5 Experiments for Bayesian Face Recognition

For the Bayesian face recognition, there are two (related) classifiers: MAP operates on both intra- and extra-personal spaces, while the simpler ML classifier bases its estimate on the intra-personal space only. For most applications, the

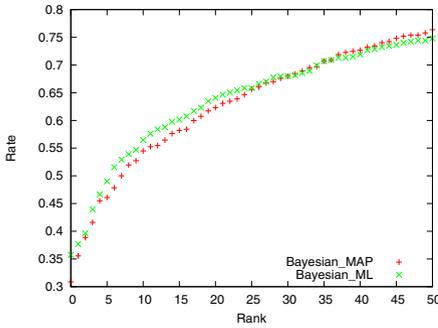


Fig. 5. Bayesian recognition after de-identification with $\sigma = 0.85$ and $k = 10$

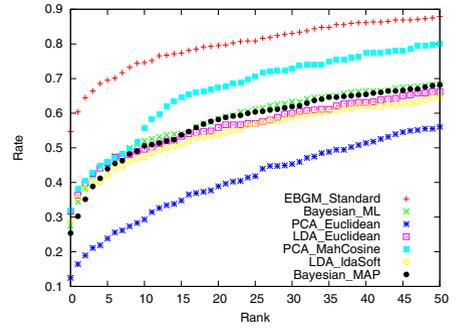


Fig. 6. Recognition rates of all algorithms after de-identification with $\sigma = 0.85$ and $k = 10$



Fig. 7. Comparison of image quality after de-identification with $k = 10$, $\sigma = 0.85$ for clusters grouped by MAP and ML

two variants provide very similar results, but our first set of experiments targets each method individually.

For the Bayesian method we have again performed de-identification as described previously. We have grouped $k = 10$ subjects into clusters, determined by their closeness according to the classification by MAP and ML. Then we have de-identified the clusters with $\sigma = 0.85$. The performance of the recognition algorithms is shown in Figure 5 and is in line with the results expected due to prior experiments for Eigenface-based algorithms.

Figure 7 shows the visual effects of the de-identification. Again we see that both subjects are clearly recognizable by humans, although the outline of the heads is blurred to a certain degree, which we attribute to variations in the pre-processing/normalization of the EBGM-method as we explained before.

5.6 Experiments for Combined Face Recognition

Finally, we combined the method of modifying PCA-based projections with our algorithm to alter EBGM-specific features. Since clustering subjects by their closeness according to MahCosine proved to be effective against Eigenface-based approaches, we decided to do the same for this experiment. We have de-identified the probe- and gallery-set with $k = 10$ and $\sigma = 0.85$ in the first step and then

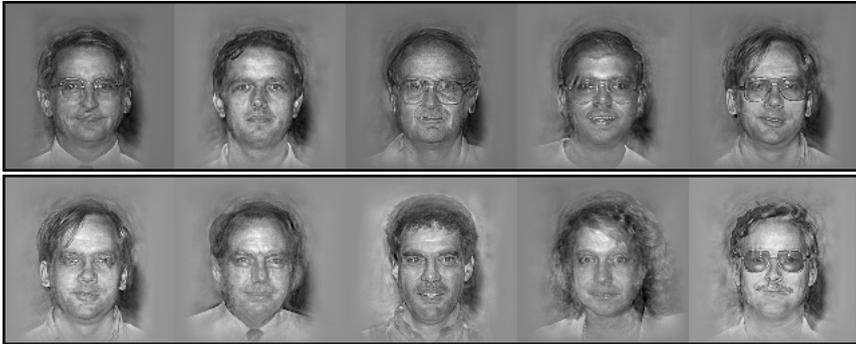


Fig. 8. Image quality after de-identification for all algorithms with $k = 10, \sigma = 0.85$. The images shown represent one cluster grouped according to MahCosine, i.e., previous work maps these to the *same* average image.

modified the jets of the resulting images to resemble the average of the same set of clusters (still grouped by MahCosine).

Figure 6 shows the recognition performance of all algorithms when operating on the same set of images. We see that the EBGM algorithm fares better in recognizing de-identified subjects than the other methods, but still only gets a 55% rank-0 recognition rate, see the discussion in the next section. Interestingly, the curves for Bayesian MAP, ML and Eigenface-based methods are very close to each other. Figure 8 shows the visual outcome of the de-identification procedure. The images are all taken from the same cluster (of 10 images total). Note that previous work (e.g., [16]) would have assigned the *same* average image to all of these images, i.e., they would be indistinguishable for humans. However, the images produced by our method are clearly distinguishable, which was the main goal of our work.

6 Discussion of Results

Face recognition algorithms work very well for images taken in a controlled environment (i.e., with regard to background, illumination, tilt, etc.). For example, the CSU implementation of the EBGM algorithm achieves a rank-0 recognition rate of more than 90% (c.f. Figure 6), and the original EBGM implementation fared even better. Our modifications reduce the rank-0 recognition rate to 55% for EBGM, and below 30% for the other algorithms, which is a big improvement.

The EBGM algorithm is, according to our experiments, harder to fool than other algorithms. Possible reasons are that EBGM is a feature-based algorithm (i.e., it works on small patches of the face, not a holistic view of the face), and it has a very good recognition rate in general. In practice, images stored on image sharing sites are not taken in a controlled environment, and we expect that in a real environment the recognition rates will drop substantially compared to the above experiments.

Often, a corporation’s interest in collecting data and a user’s interest in privacy are diametrical. That said, using anonymization techniques will probably cause a reaction by the corporations deploying face recognition algorithms, which could eventually lead to an arms-race between both sides. This and finding a solution to the remaining issues will most likely inspire future research.

7 Related Work

Systematic research on face de-identification started with work by Newton, Sweeney, and Malin (e.g., [22,23]). Their goal was to achieve *perfect anonymity*, which makes identifying an individual from the anonymized photo provably impossible for machines (and humans). Their first approach, called “k-same”-method, computed the de-identified image either averaging the closest k images pixel-wise (“k-same-pixel”-method) or in Eigenface (“k-same-eigen”-method). While achieving strong security guarantees, the resulting image quality of both approaches was mediocre. Consequently, follow-up work improves image quality [16] and defines a more rigorous framework for the different notions of privacy protection models [15].

More recent work by Do et al. focuses on systematic de-identification methods for feature-based systems which are used in forensics [9]. In this work, the authors delude a class of image recognition algorithms (as opposed to face recognition), which can be understood as a superset of the class of algorithms we are concerned with.

A couple of ad-hoc methods such as masking parts of the face (e.g. the eyes) or blurring or pixelation of faces [7,17,20,36] were eventually tested. However, these methods are visually intrusive and target human and algorithmic recognition alike. Even worse, it was shown that their effectiveness is very limited [23], so they are not a good option.

8 Conclusion

We have shown a reliable way to de-identify face images against a wide range of currently available recognition algorithms. While not achieving the very strong notion of k -anonymity, we achieved a level of anonymity which is sufficient to counter the two most pressing problems that face recognition software poses for users of social networks: first, automated extraction of the social graph, i.e., friendship relations; second, automated tagging of people in images. At the same time we get an image quality which still allows humans to identify persons in images.

Acknowledgment. Portions of the research in this paper use the FERET database of facial images collected under the FERET program, sponsored by the DOD Counterdrug Technology Development Program Office [26,27].

References

1. Belhumeur, P., Hespanha, J., Kriegman, D.: Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection. In: Buxton, B.F., Cipolla, R. (eds.) ECCV 1996. LNCS, vol. 1064, pp. 45–58. Springer, Heidelberg (1996)
2. Beresford, A.R., Kübler, D., Preibusch, S.: Unwillingness to pay for privacy: A field experiment. IZA Discussion Paper No. 5017 (2010), <http://ftp.iza.org/dp5017.pdf>
3. Beveridge, J., She, K.: Fall 2001 update to the CSU PCA versus PCA+LDA comparison (2001), <http://www.cs.colostate.edu/evalfacerec/>
4. Beveridge, R., Bolme, D., Teixeira, M., Draper, B.: The CSU face identification evaluation system user’s guide: Version 5.0 (2003), <http://www.cs.colostate.edu/evalfacerec/>
5. Bolme, D.S.: Elastic bunch graph matching. Master’s thesis, CSU Computer Science Department (2003)
6. Bolme, D.S., Beveridge, J.R., Teixeira, M., Draper, B.A.: The CSU face identification evaluation system: Its purpose, features and structure. In: Proc. International Conference on Vision Systems, pp. 304–311 (2003)
7. Boyle, M., Edwards, C., Greenberg, S.: The effects of filtered video on awareness and privacy. In: Proc. of the 2000 ACM Conference on Computer Supported Cooperative Work, CSCW 2000, pp. 1–10. ACM (2000)
8. Deutsche Welle. Facebook facial recognition raises eyebrows in Germany, EU, <http://www.dw-world.de/dw/article/0,,15144128,00.html>
9. Do, T.-T., Kijak, E., Furon, T., Amsaleg, L.: Deluding image recognition in sift-based cbir systems. In: Proc. 2nd ACM Workshop on Multimedia in Forensics, Security and Intelligence, pp. 7–12. ACM (2010)
10. EPIC.org. EPIC files complaint, urges investigation of Facebook’s facial recognition techniques (2011), <http://epic.org/2011/06/epic-files-complaint-urges-inv.html>
11. Facebook Announcement: 10 billion photos, http://www.facebook.com/note.php?note_id=30695603919
12. Facebook Announcement: Making Photo Tagging Easier, <http://www.facebook.com/blog.php?post=467145887130>
13. Facebook Inc., <http://developers.facebook.com/docs/opengraph/>
14. Flickr Blog: 4 billion photos, <http://blog.flickr.net/en/2009/10/12/4000000000/>
15. Gross, R., Sweeney, L.: Towards real-world face de-identification. In: First IEEE International Conference on Biometrics: Theory, Applications, and Systems, BTAS 2007, pp. 1–8 (2007)
16. Gross, R., Sweeney, L., de la Torre, F., Baker, S.: Model-based face de-identification. In: Proc. 2006 Conference on Computer Vision and Pattern Recognition Workshop, CVPRW 2006. IEEE Computer Society (2006)
17. Hudson, S.E., Smith, I.: Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In: Proc. 1996 ACM Conference on Computer Supported Cooperative Work, CSCW 1996, pp. 248–257. ACM (1996)
18. Miller, P., Lyle, J.: The effect of distance measures on the recognition rates of PCA and LDA based facial recognition. Tech. rep., Clemson University (2008)
19. Moghaddam, B., Jebara, T., Pentland, A.: Bayesian face recognition. *Pattern Recognition* 33(11), 1771–1782 (2000)

20. Neustaedter, C., Greenberg, S., Boyle, M.: Blur filtration fails to preserve privacy for home-based video conferencing. *ACM Trans. Comput.-Hum. Interact.* 13, 1–36 (2006)
21. New York Times, How privacy vanishes online (2010), <http://www.nytimes.com/2010/03/17/technology/17privacy.html>
22. Newton, E., Sweeney, L., Malin, B.: Preserving privacy by de-identifying facial images. Tech. Rep. CMU-CS-03-119, Carnegie Mellon University, School of Computer Science (2003)
23. Newton, E., Sweeney, L., Malin, B.: Preserving privacy by de-identifying facial images. *IEEE Transactions on Knowledge and Data Engineering* 17(2), 232–243 (2005)
24. Okada, K., Steffens, J., Maurer, T., Hong, H., Elagin, E., Neven, H., von der Malsburg, C.: The Bochum/USC face recognition system and how it fared in the FERET phase III test. In: *Face Recognition: From Theory to Applications* (1998)
25. Pfitzmann, A., Köhntopp, M.: Anonymity, unobservability, and pseudonymity – a proposal for terminology. In: *Workshop on Design Issues in Anonymity and Unobservability*, pp. 1–9 (2000)
26. Phillips, J.P., Moon, H., Rizvi, S.A., Rauss, P.J.: The FERET evaluation methodology for face-recognition algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(10), 1090–1104 (2000)
27. Phillips, P., Wechsler, H., Huang, J., Rauss, P.J.: The FERET database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing* 16(5), 295–306 (1998)
28. Phillips, P.J., Rauss, P.J., Der, S.Z.: FERET (face recognition technology) recognition algorithm development and test results. Tech. Rep. ARL-TR-995, Army Research Laboratory (1996)
29. Samarati, P., Sweeney, L.: Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. In: *IEEE Symposium on Security and Privacy*. IEEE Computer Society (1998)
30. Sweeney, L.: k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.* 10(5), 557–570 (2002)
31. Teixeira, M.L.: The Bayesian intrapersonal/extrapersonal classifier. Master’s thesis, Colorado State University (2003)
32. Turk, M., Pentland, A.: Eigenfaces for recognition. *Journal of Cognitive Neuroscience* 3(1), 71–86 (1991)
33. Wiskott, L., Fellous, J.-M., Kruger, N., von der Mals, C.: Malsburg, C. V. Face recognition by elastic bunch graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19(7), 775–779 (1997)
34. Yambor, W., Draper, B., Beveridge, R.: Analyzing PCA-based face recognition algorithms: Eigenvector selection and distance measures. In: *Empirical Evaluation Methods in Computer Vision* (2002)
35. Yang, M.-H., Kriegman, D., Ahuja, N.: Detecting faces in images: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24(1), 34–58 (2002)
36. Zhao, Q.A., Stasko, J.T.: Evaluating image filtering based techniques in media space applications. In: *Proc. of the 1998 ACM Conference on Computer Supported Cooperative Work, CSCW 1998*, pp. 11–18. ACM (1998)
37. Zhao, W., Krishnaswamy, A., Chellappa, R., Sweets, D., Weng, J.: Discriminant analysis of principal components for face recognition. In: *Proc. 3rd IEEE Int. Conference on Face and Gesture Recognition*, pp. 336–341 (1998)