GARP-Face: Balancing Privacy Protection and Utility Preservation in Face De-identification

Liang Du[1], Meng Yi[1], Erik Blasch[2], and Haibin Ling[1]
[1]Department of Computer and Information Science
Temple University
Philadelphia, PA, 19122
{liang.du,Mengyi,hbling}@temple.edu
[2]Air Force Research Lab
Rome, NY, 13441
erik.blasch@rl.af.mil

Abstract

Face de-identification, the process of preventing a person' identity from being connected with personal information, is an important privacy protection tool in multimedia data processing. With the advance of face detection algorithms, a natural solution is to blur or block facial regions in visual data so as to obscure identity information. Such solutions however often destroy privacy-insensitive information and hence limit the data utility, e.g., gender and age information. In this paper we address the de-identification problem by proposing a simple yet effective framework, named GARP-Face, that balances utility preservation in face de-identification. In particular, we use modern facial analysis technologies to determine the Gender, Age, and Race attributes of facial images, and Preserving these attributes by seeking corresponding representatives constructed through a gallery dataset. We evaluate the proposed approach using the MORPH dataset in comparison with several state-of-the-art face de-identification solutions. The results show that our method outperforms previous solutions in preserving data utility while achieving similar degree of privacy protection.

1. Introduction

With advances in digital imaging technologies, it has never been easier to capture and share visual data as it is today. We may take photos or videos conveniently using cell phones or other digital devices, and immediately share them on online platforms. On the other hand, we ourselves are also often under the lenses of surveillance cameras, or filmed by other people, sometimes unknownly. Accompanied with this digital convenience; however, is the potential privacy leakage for images to be used for identity theft.

An increasing amount of effort has been devoted towards addressing identity theft with imagery, from both academia and industry (See Section 1.1). Large, street-level image collections like Google Street View require automatic systems to detect and blur faces [3]. In television news, we see people whose faces are blocked or pixelated to protect their identities. We encounter a similar problem when distributing research datasets. Some medical face databases [33, 34] are not accessible to other groups, or require intensive manual post-processing for patient privacy.

A focus of previous studies is on reliably detecting facial regions[3]. Once a face is located, it will be either blurred (typically via Gaussian kernels) or blocked so as to obscure the identity. However, facial images contain rich information, such as gender, race and age, etc. This kind of identity-insensitive information is often the data utility that is desired to preserve in many applications involving visual understanding and data mining. Therefore, a successful face de-identification algorithm should balance protecting privacy and preserving utility.

Ad-hoc methods that simply blur the facial part in an image typically perform very well in terms of privacy protection, but cause serious loss of data utility as a side effect. This is obviously not desirable. Researchers have proposed some sophisticated methods for face de-identification [25, 27, 28, 7]. They can be roughly categorized into two categories, namely k-same based methods [25, 27, 28] and face replacement [7].

The *k-same* methods have made the pioneering attempt to borrow the k-anonymity concept from privacy research in data mining to de-identify facial images [25, 27, 28]. These methods investigate the k nearest neighbors of the query face in the face image set. In this way, the query

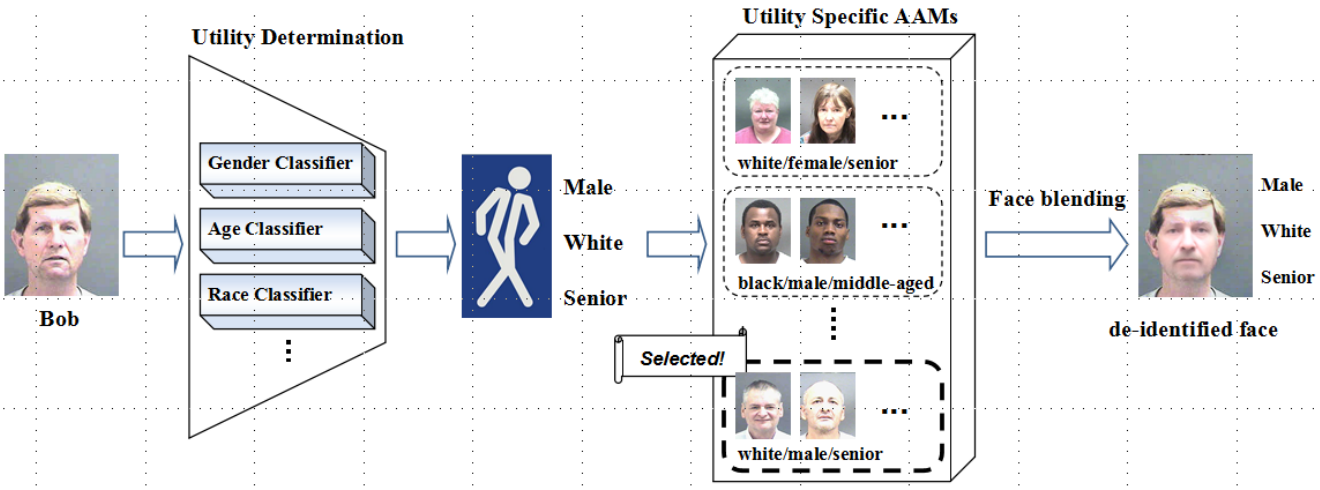

Figure 1. Pipeline of Proposed method. Given an input face image, first its utilities are determined by trained attribute classifiers. Then, it is modeled using corresponding utility specific AAM. Finally, de-identified face is formulated by blending the input and its closest neighboring superface in the selected sub-category.

face is anonymized among at lease k candidates, namely k-anonymity. And it guarantees that after de-identification, face recognition accuracy is below $1/k$ [25]. Despite the guaranteed privacy gain and moderate consideration of data utility, the effectiveness of k-same methods in preserving data utility is questionable. Intuitively, the similarity of face images are correlated to consensus of utility. However, the face space is highly nonlinear, and affinity in a certain model space does not necessarily mean the fidelity in the semantic utility space. It is likely that faces similar in a general model space may have very different attributes, e.g. two faces may be close neighbors in a particular space, while one is male, and the other is female. Another potential weakness of k-same methods is that it is self de-identification done completely inside the original image set. It can be much easier to attack once the image set is leaked. We will address this security weakness by introducing a separate reference gallery.

Face replacement methods, such as the *Face Swapping* [7], replace a probe facial image with a "similar" face in a library $\mathcal{L}$. These methods benefit from the use of a library in a way that they can get a similar face with a different identity. However, there are two limitations: First, from a security point of view, this method can be attacked by duplicating the process and inferring the identity of the de-identified face with high confidence as the one returns the same top nearest neighbor. Second, the replacing method, as well as the k-same methods, may alter the attributes like gender of a face or create large artifacts.

In this paper we propose a new de-identification framework named *GARP-Face* that bridges the semantic gap between facial appearance and data utility. Unlike previous methods which implicitly use appearance similarity as a measurement for data utility, we select three essential utilities: *Gender, Age, and Race* information and preserve them *explicitly* in de-identification. More specifically, we design a structured utility hierarchy based on observation on real image sets, and build a utility specific Active Appearance Model (AAM) for each category. These models are pre-trained on an external image gallery. We also formulate superfaces by aggregation of similar faces in each category. Given an input face image, GARP-Face first determines its gender, age and race attributes using modern facial analysis techniques, then the de-identified face is generated by blending with the GAR representative superface, which is most similar to the original face and has consistent attributes. The pipeline of GARP-Face is summarized in Figure 1.

Compared with previous research, our contributions are mainly two folds:

- Unlike previous approaches which either ignore related utilities or implicitly work on them, our approach explicitly thus more effectively preserves facial image utilities involving gender, age and race information.
- The identity of query faces are diluted using an external gallery. This extension to the k-anonymity model further enhances visual privacy protection.

To validate the effectiveness of the proposed GARP-Face algorithm, a de-identification experiment is conducted using the MORPH database [14] involving state-of-the-art face de-identification algorithms. The results clearly demonstrate the superiority of our approach in utility preservation while achieving similar degree of privacy protection.

The rest of this paper is organized as follows. Related work are surveyed in Sec. 1.1. In Sec. 2 we first

discusses the formal definition and evaluation of face de-identification, then introduce our utility preserving GARP Face De-identifier. In Sec. 3 GARP is evaluated on the large MORPH public face dataset. Finally, we conclude this paper in Sec. 4.

1.1. Related Work

The problem of preserving privacy in data mining has been intensively studied [15, 31, 4, 30, 6]. [30] studied the problem of tradeoff between privacy and utility in data privacy protection problem. [4] studied the possible drawbacks of k-anonymity, showing that lack of diversity might lead to the fail of privacy protection. In [15], an attribute generalization supports as semantic hierarchy. For a survey of privacy protection, one can refer to [6]. Recently, there are increasing interests in visual privacy protection and gaining attraction through several initiatives, such as COST action IC-1206 [1]. However, due to the lack of direct semantic interpretation of visual information like images and videos, there's relatively few works on privacy protection in visual privacy protection. Privacy protection in visual analysis recently has increasing amount of research attention. A survey of privacy preserving video surveillance is given in [5]. Chen *et al*. [8] studied the privacy preserving problem in the context of health care related surveillance. Li *et al*. [10] uses coprime for privacy protected video communication. Du and Ling [17] studied the problem of preservative license plate de-identification. Chan *et al*. [1] proposed a method for counting people without explicit human detection. Upmanyu *et al*. [19] designed a secure video sharing system inspired by the Chinese Remainder Theorem which split each frame into a set of random images. Schiff *et al*. [11] uses markers in surveillance videos to detect persons whose identities are sensitive and thereafter hide such information by masking. In [12] an algorithm is presented to classify covert photos that often convey a privacy leakage.

Recently, privacy-preserving biometric identification attracts interests from many researches [36, 37, 38, 39, 40, 41, 42]. Ross and Othman [23] explore the use of visual cryptography for imparting privacy to biometric data such as face images. In [35], it introduced a face identification system which reduces the privacy impact of camera based surveillance. Boult [24] explores privacy protection through invertible cryptographic obscuration.

Face de-identification is an important tool for visual privacy protection. Many face de-identification methods focus on face detection and simply blur or cover the detected faces, which often brings unpleasant artifacts to the data and damages data utility. Some researchers try to remedy this problem by applying a more "careful" blurring or masking. For example, a person de-identification method is proposed in [22] to de-identify a person but retain his/her action information. It implicitly uses the human action as the data utility. However, this utility is very limited and many important attributes (*e.g*., gender) are lost. In addition, how to conduct a "sufficient blur" itself is non-trivial [3].

The k-same de-identification, which is based on the k-anonymity framework introduced by Sweeney [18], guarantees that each de-identified facial image represents at least k faces in the gallery, therefore limiting face recognition performance to $1/k$. Its variants k-Same-Eigen, k-Same-M (Model) [25, 27] and multi-factor models [28] adopt different face modeling to produce results representative of the entire k gallery.

Replacing the facial image with a "similar" face in a library $\mathcal{L}$ is another way for face de-identification. Bitouk *et al*. [7] proposed an automatically face replacement method, *Face Swapping*, which replaces a target face by a similar face in a large pre-constructed library. It first detects all faces in the input image, and selects similar candidate faces from a face library. Then, it adjusts the input image to seamlessly blend in the top-ranked candidate faces.

As mentioned in the introduction, our work is closely related to both k-same and face replacement methods. The main difference between GARP and the k-same methods lies in the separation of probe images and gallery images; while GARP differs with face replacement by incorporating the k-anonymity concept. More detailed discussion is given in Sec. 2.

2. Utility Preserving Face De-identification

In this section, we first present the formal definition of face de-identification and its contradictory requirement to protect privacy while retaining utility. As well as setting up the problem, this discussion also inspires our utility preserving GARP Face De-identifier. Then, we present the GARP framework and which details on two key aspects: utility determination and utility specific AAMs.

2.1. Face De-identification

Here we provide the definition of face de-identification problem, and then introduce the evaluation criterion balancing privacy and utility concerns.

Given a set of probe faces $\mathcal{P}$ and a set of reference faces $\mathcal{R}$, face recognition is a function $f : \{\mathcal{P} \rightarrow \mathcal{R}\}$ which associates a probe face to a unique reference face. Face de-identification can be viewed as a transformation function δ from the original face image set $\mathcal{I} = \{I_1, I_2, ...I_l\}$ containing l face images to a set of de-identified face images $\hat{\mathcal{I}} = \{\hat{I_1}, \hat{I_2}, ..\hat{I_l}\}$, so that

$$\delta(I_i) = \hat{I}_i, i = 1, ..., l \tag{1}$$

The de-identification function δ intends to decrease recognition accuracy and protect privacy.

[1] http://costic1206.uvigo.es/

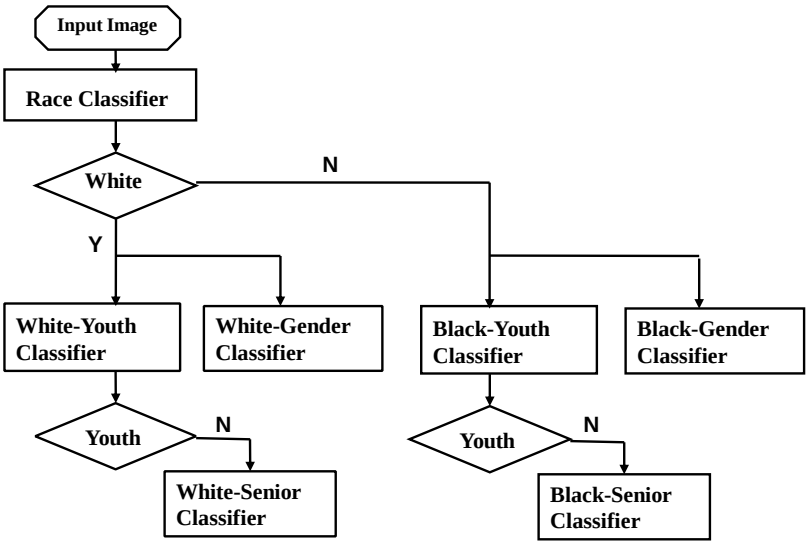

Figure 2. Utility Determination Hierarchy. The attributes for an input face is determined in a hierarchical manner. First, it goes through a race classifier. Then, in the corresponding race category, we classify its gender and whether it is youth. If it is not youth face, one more classifier is used to decide whether it is middle-aged or senior.

The performance of a de-identification algorithm can be measured in two aspects: *privacy gain (PG)* and *utility loss (UL)* [31]. Privacy gain, also viewed as loss of the identify information from an attacker's point of view, can be expressed as:

$$PG(\delta, \mathcal{I}) = \sum_{i=1}^{l} (P(i|\hat{I}_i) - P(i|I_i)) \,, \tag{2}$$

where, with slight overload of notation, we use i to denote the identity of face image I_i for the sake of conciseness; and $P(i|I_i)$ denotes the probability of finding the true identity i given face image I_i.

Utility of a dataset, whether de-identified of not, is innately tied to the computation that one may perform on it [13]. Here, we evaluate the utility of the de-identified probe set $\mathcal{I}$ in terms of count querying, which has been widely used as a measurement of data utility [31, 13]. Let Q denote the count querying operation, then the utility loss can be defined as:

$$UL(\delta, \mathcal{I}) = \sum_{i=1}^{l} (Q(I_i) - Q(\hat{I}_i)) \,. \tag{3}$$

2.2. GARP De-identification

The proposed GARP de-identification mainly consists of the following components: 1) Utility determination, 2) Utility-specific AAM models, and 3) a diverse face gallery. For utility, classifiers for selected attributes are trained to determine which utility specific model should be applied to the query face. Since the faces and utilities cannot be well captured by a single general model, we propose to build an attribute specific face model using AAM (Active Appearance Model) [29]. A large and diverse face gallery $\mathcal{G}$ is used for both training AAM models and attribute classifiers. Furthermore, the superfaces are generated from $\mathcal{G}$ according to the utility class of the face to-be de-identified.

Figure 1 illustrates the pipeline of our de-identification procedure. Input is a face I containing explicit identity information: e.g. "Bob", which is private; and underlying descriptive attributes: e.g. white, male, senior, which are useful and privacy insensitive. First, our pre-trained utility classifiers will extract the attributes of the face image. Next, based on the extracted attributes, we refer to the utility specific AAM model that associates with the particular attributes, and parameterizes the input image in that model space. Last, we refer to the superfaces that associate with the particular attributes, and blend the input face with the closet superface to form the de-identified face. We will elaborate on utility determination in Sec. 2.3; and the utility specific AAMs in Sec. 2.4.

Our strategy of forming a superface by aggregation can be interpreted as a k-anonymity approach in the utility space. If an attribute-specific sub-gallery is divided into m clusters, it means that the whole attribute class of the whole population is represented with m super-faces. Then every de-identified face will be undistinguished with k' faces,

$$k' = \frac{|C|}{m},$$

Where $|C|$ is the cardinality of utility of sub-gallery C. Typ-

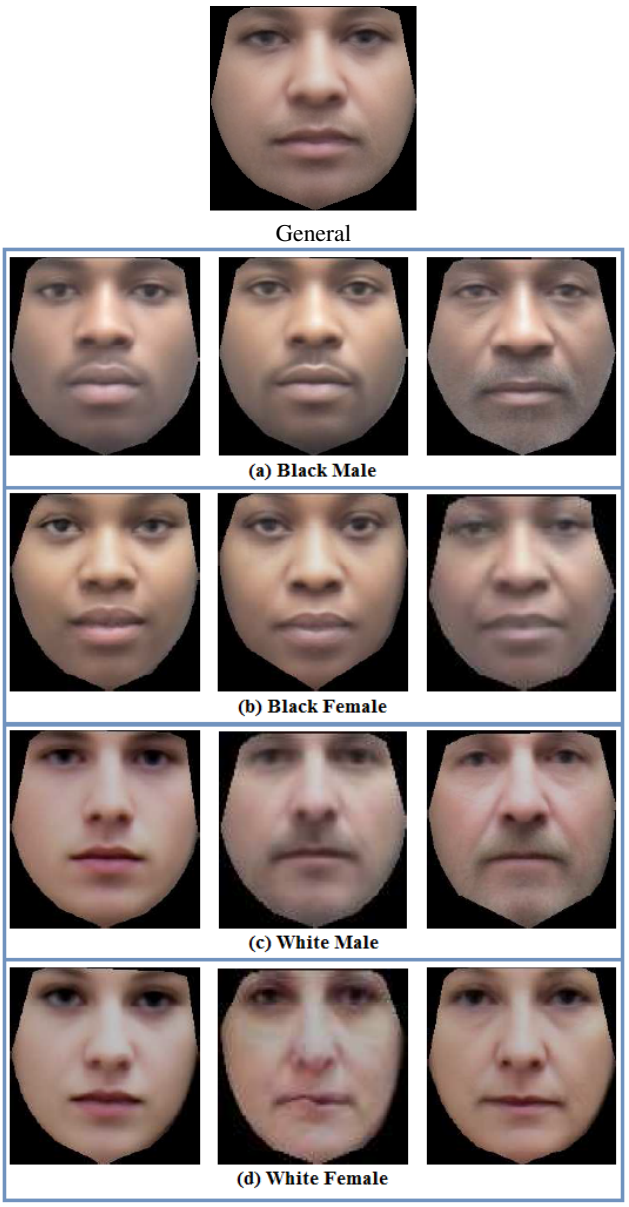

Figure 3. Mean face of the general AAM model and utility specific AAM models. On the first row is mean face of the general model. From row 2 to row 5 are mean faces for black male, black female, white male and white female respectively. In each row, from left to right are faces of youth, middle-aged and senior respectively.

ically, $|C|$ is a very large number, for example, the whole population of middle-aged white males. Thus, the k' here, which is equivalent to the k in k-anonymity algorithms, is very large. Thus the probability of finding the true identity, which is less than $1/k$, is very small. This is another merit over the k-same algorithms. In principal, dividing the sub-gallery into clusters is not necessary. However, by doing so we may prevent seeing many similar faces in the de-identified image. The larger m is, the more realistic and visually meaningful the de-identified images are. On the other hand, a larger m leads to a longer computation time. Thus m should be chosen according to the requirement of specific privacy application. Generally, $|C|$ is very large (*e.g.*, in our experiments > 1000), GARP de-identification can achieve near-optimal privacy gain.

The gallery $\mathcal{G}$ plays the role of modeling the facial image space. From $\mathcal{G}$ we gain knowledge about the attribute space. Thus, $\mathcal{G}$ should be diverse enough so that it can represent various types of data utility. Meanwhile, $\mathcal{G}$ should be readily to be downgraded to a privacy preserved level in the underlying facial information hierarchy.

2.3. Utility Determination

The utility of a facial image is the informative yet privacy insensitive attributes. Here we select three attributes of common interest: gender, race and age. One can also retain other attributes, or use more sophisticated classifiers to get better accuracy, under the similar framework.

The three attributes are organized in a structured manner. It is noticed that race could affect facial appearance to a very large extent, thus we first determine the race of a face, and then train race-specific gender classifiers and and race-specific age classifiers [32]. The implementation of age classifiers is also done in a hierarchical way: two binary classifiers are trained to classifier the three classes: *youth*, *middle-aged* and *senior*. We first apply one classifier to determine if a photo is youth or not. If not youth, we apply the second classifier to distinguish between middle-aged and senior. The attribute hierarchy is shown in Figure. 2.

The attribute-level classification problem has been recently studied for face verification and many other computer vision problems [2, 20]. In this paper, the gender and race classifiers are trained using adaboost classifiers and the Haar features. Age classifiers are trained using adaboost classifiers and Gabor features. These attribute classifiers produce satisfactory accuracies. Results are summarized in Table 1.

2.4. Utility Specific AAM Model

AAM model is a generative parametric model which have been successfully utilized in many face modeling and face tracking applications. It not only seeks to matches the shape of the model but also match the representation of texture over the object [29]. Here, we propose to use utility specific AAMs, in order to explicitly preserve data utility. For each of the 12 attribute classes, we manually label facial shape points for images in the reference gallery $\mathcal{G}$, and train 12 separate AAM models. Then the parametrization of a face image is to minimize the *mean square error* (MSE) of the difference between a utility specific AAM model and the input image. Figure 3 demonstrates mean face of each

Table 1. Accuracy of structured attribute classifiers. Description in the parenthesis is the condition assumed to be true. E.g. Middle-aged (White, not Youth) denotes the accuracy of Middle-aged classifier, on the subset of white and not youth faces.

Attributes	Error rate
Race	0.0572
Gender (White)	0.0545
Gender (Black)	0.0552
Youth (White)	0.1275
Youth (Black)	0.2270
Middle-aged (White, not Youth)	0.0836
Middle-aged (Black, not Youth)	0.0846

AAM space corresponding to the 12 attribute classes.

3. Experiments

Our GARP de-identifier is evaluated against state-of-the-art methods using the publicly available MORPH database [14], which is a large dataset containing 55,000 unique images of more than 13,000 individuals, with diverse demographic information e.g., age, gender and race. Each image is associated with its attribute information. We randomly divide the dataset into a test set and a gallery set. The attributes we test in this experiment are age, race and gender. Age is divided in to three groups: youth, middle-aged and senior. Race contains two categories: black and white. Gender contains male and female.

The gallery $\mathcal{G}$ consists of 13620 individual facial images covering various facial attributes. The set of images to be de-identified is a subset $\mathcal{I}$ of 1200 facial images randomly sampled from the MORPH database which is not overlapping with $\mathcal{G}$.

We compare the proposed GARP de-identifier to two face de-identification algorithms. One is the k-same algorithm as described in [27]. This method works solely on the input image set, and de-identify each image using the affine combination of its k-nearest neighbors. We also implemented an model-based de-identifier which uses one general AAM model (See Procedure 1). And the procedure of GARP is shown in Procedure 2. Both general model-based de-identifier and GARP utilize an external gallery $\mathcal{G}$, thus the comparison between these two can demonstrate the advantage of utility specific AAMs over general AAM.

As mentioned in Sec. 2.2, GARP, as well as the other two approaches, fall into the k-anonymity framework. Thus, when choosing the same k, the recognition accuracy of de-identified faces after apply any of these algorithms share the same upper bound $1/k$. We can achieve desirable privacy protection level by setting proper k. In other words, the privacy gain of these methods are controllable and comparable. Therefore, here we set k to 30, and focus on the utility side. The utility loss is measured following equation 3, with a small twist. We normalize it using the size of input image set $\mathcal{I}$, so that *UL* is within $[0, 1]$ thus more interpretable.

Procedure 1 General model-based De-identification with Gallery

Input: Probe face set $\mathcal{I} = \{I_1, I_2, ...I_l\}$;
Reference gallery $\mathcal{G}$;
Output: De-identified face set $\mathcal{I}_d$;
1: Initialize an empty set $\mathcal{I}_d$;
2: Train a general AAM model $\mathcal{M}$ on $\mathcal{G}$;
3: Generate superfaces $\mathcal{S} = \{S_1, S_2, ...S_n\}$ on $\mathcal{G}$;
4: **for** $i = 1$ to l **do**
5: Represent I_i using AAM model: $p_i = \mathcal{M}(I_i)$;
6: Find the superface S_k which is the closest neighbor of p_i, $k \in \{1, 2, ..., n\}$;
7: Blend S_k into p_i and add the de-identified face $\hat{I}_i$ to $\mathcal{I}_d$;
8: **end for**

Procedure 2 GARP-Face De-identification

Input: Probe face set $\mathcal{I} = \{I_1, I_2, ...I_l\}$;
Reference gallery $\mathcal{G}$;
Output: De-identified face set $\mathcal{I}_d$;
1: Initialize an empty set $\mathcal{I}_d$;
2: Train attribute classifiers $\mathcal{C}_s$ on $\mathcal{G}$;
3: Divide $\mathcal{G}$ into sub-galleries $\mathcal{G}_1$,$\mathcal{G}_2$,...,$\mathcal{G}_m$ according to attributes;
4: Train a utility specific AAM $\mathcal{M}_j$ on each $\mathcal{G}_j$, $j \in \{1, 2, ..., m\}$;
5: Generate superfaces $\mathcal{S}_j = \{S_1^j, S_2^j, ...S_n^j\}$ on each $\mathcal{G}_j$
6: **for** $i = 1$ to l **do**
7: Determine the attributes of I_i using $\mathcal{C}_s$, find the corresponding sub-category j;
8: Represent I_i using AAM model $\mathcal{M}_j$: $p_i = \mathcal{M}_j(I_i)$;
9: Find the superface S_k^j which is the closest neighbor of p_i, $j \in \{1, 2, ..., n\}$;
10: Blend S_k^j into p_i and add the de-identified face $\hat{I}_i$ to $\mathcal{I}_d$;
11: **end for**

Table 2 shows the de-identification results. We evaluate utility loss on all aspects combined, and on each utility (age, gender, race) separately. From Table 2 and Figure 5, we can see that GARP ensures significantly lower data utility loss in all categories after de-identification, even with near optimal privacy gain achieved. Visual results are demonstrated in Figure 4.

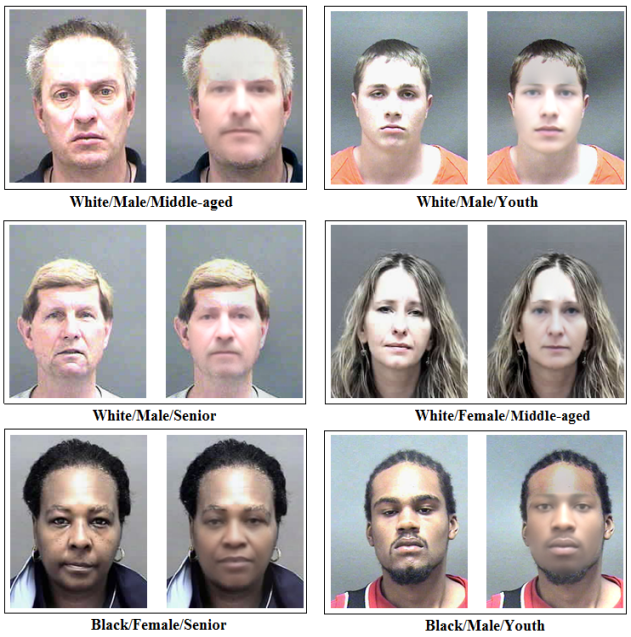

Figure 4. Sample De-identification results of GARP-Face: In each rectangle, the left face is the original one, and the right face is the de-identified face. Utilities are shown below each pair.

Table 2. Utility losses. Utility losses of k-same, general AAM model-based de-identification and proposed GARP-Face are measured using normalized count querying. The last row is the utility loss of three attributes combined together.

	k-same	Gen. AAM	GARP-Face
Race	0.4818	0.3727	**0.0897**
Gender	0.1469	0.3139	**0.1372**
Age	0.3606	0.4056	**0.0878**
Combined	0.4897	0.5106	**0.1173**

4. Conclusion

Face de-identification is an important component in visual privacy protection. In this paper, we studied the objective of face de-identification, and developed a novel face de-identifier which explicitly addresses utility concern. We demonstrated that our GARP de-identifier outperforms other state-of-the-art methods, following the evaluation criterion combining privacy protection and utility preservation. In future research, we plan to apply the methodology to other kinds of visual information, including soft biometrics traits which could be auxiliary information for identification. Additionally, a more accurate attribute classifier can be developed to improve the quality of this de-identifier.

Acknowledgement

The authors would like to thank anonymous reviewers for valuable suggestions to improve the paper. The work is supported in part by the NSF Grants IIS-1218156 and CAREER Award IIS-1350521.

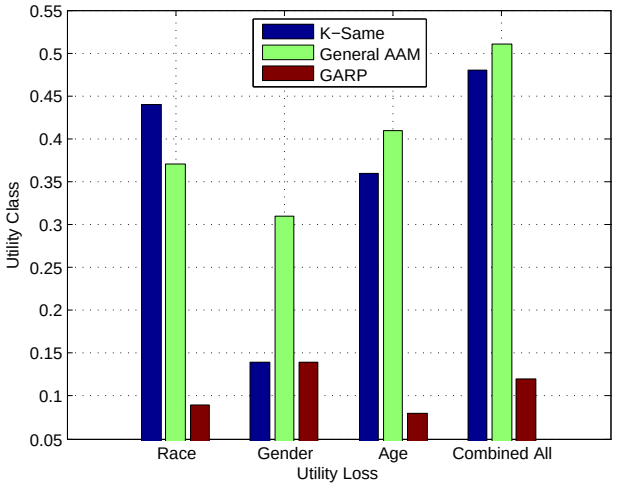

Figure 5. The comparison of utility losses of different methods. Blue bar is k-same, green bar is general AAM, crimson is for GARP-Face.

References

[1] A. Chan, Z.-S. Liang, and N. Vasconcelos. Privacy preserving crowd monitoring: Counting people without people models or tracking. In *CVPR*, 2008.

[2] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *CVPR*, 2009.

[3] A. Frome, G. Cheung, A. Abdulkader, M. Zennaro, B. Wu, A. Bissacco, H. Adam, H. Neven, and L. Vincent. Large-scale privacy protection in google street view. In *ICCV*, 2009.

[4] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkitasubramaniam. L-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data*, 1, March 2007.

[5] A. Senior. *Protecting Privacy in Video Surveillance*, chapter Privacy Protection in a Video Surveillance System, pages 35– 47. 2009.

[6] B. Fung, K. Wang, R. Chen, and P. Yu. Privacy-preserving data publishing: A survey of recent developments. *ACM Comput. Surv.*, 42(4), 2010.

[7] D. Bitouk, N. Kumar, S. Dhillon, P. N. Belhumeur, and S. K. Nayar. Face Swapping: Automatically Replacing Faces in Photographs. *ACM Trans. on Graphics (also Proc. of ACM SIGGRAPH)*, Aug 2008.

[8] D. Chen, Y. Chang, R. Yan, and J. Yang. Tools for protecting the privacy of specific individuals in video. *EURASIP J. Appl. Signal Process.*, 2007:107–107, January 2007.

[9] E. Newton, L. Sweeney, and B. Malin. Preserving privacy by de-identifying facial images. *IEEE Transactions on Knowledge and Data Engineering*, 17(2):232-243, 2005.

[10] F. Li, Z. Li, D. Saunders, and J. Yu. A theory of coprime blurred pairs. In *ICCV*, 2011.

[11] J. Schiff, M. Meingast, D. Mulligan, S. Sastry, and K. Goldberg. Respectful cameras: detecting visual markers in real-time to address privacy concerns. In *International Conference on Intelligent Robots and Systems*, 2007.

[12] H. Lang and H. Ling Classifying Covert Photographs. In *CVPR*, 2012.

[13] J. Brickell and V. Shmatikov. The cost of privacy: destruction of data-mining utility in anonymized data publishing. In *SIGKDD*, 2008.

[14] K. Ricanek and T. Tesafaye. Morph: a longitudinal image database of normal adult age-progression. In *FG*, 2006.

[15] K. Wang, P. S. Yu, and S. Chakraborty. Bottom-up generalization: A data mining solution to privacy protection. In *ICDM*, 2004.

[16] K. LeFevre, D. J. DeWitt, and R. Ramakrishnan. Incognito: efficient full-domain k-anonymity. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, 2005.

[17] L. Du and H. Ling. Preservative license plate de-identification for privacy protection. In *ICDAR*, 2001.

[18] L. Sweeney. K-anonymity: a model for protecting privacy. *Int'l J. on Uncertainty, Fuzziness, and Knowledge-Based Systems*, 10(5):557–570, 2002.

[19] M. Upmanyu, A. Namboodiri, K. Srinathan, and C. Jawahar. Efficient privacy preserving video surveillance. In *ICCV*, 2009.

[20] N. Kumar, A. Berg, P. Belhumeur, and S. Nayar. Describable visual attributes for face verification and image search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(10):1962 –1977, 2011.

[21] P. Samarati. Protecting respondents identities in microdata release. *IEEE Transactions on Knowledge and Data Engineering*, 13(6):1010 –1027, 2001.

[22] P. Agrawal and P. Narayanan. Person de-identification in videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 21(3):299 –310, March 2011.

[23] A. Ross and A. Othman. Visual Cryptography for Biometric Privacy. *IEEE Transactions on Informaiton Forensics and Security*, 6(1):70–81, 2011.

[24] T. E. Boult. PICO: Privacy through invertible cryptographic obscuration. In *IEEE Computer Vision for Interactive and Intelligent Environment*, 2005.

[25] R. Gross, E. Airoldi, B. Malin, and L. Sweeney. Integrating utility into face de-identification. In G. Danezis and D. Martin, editors, *Privacy Enhancing Technologies*, volume 3856 of *Lecture Notes in Computer Science*, pages 227–242. Springer, 2005.

[26] R. Gross, I. Matthews, and S. Baker. Generic vs. person specific active appearance models. *Image and Vision Computing*, 23(1):1080–1093, November 2005.

[27] R. Gross, L. Sweeney, F. de la Torre, and S. Baker. Model-based face de-identification. In *Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop*, 2006.

[28] R. Gross, L. Sweeney, F. D. la Torre, and S. Baker. Semi-supervised learning of multi-factor models for face de-identification. In *CVPR*, 2008.

[29] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):681–685, 2001.

[30] T. Li and N. Li. On the tradeoff between privacy and utility in data publishing. In *SIGKDD*, 2009.

[31] V. Rastogi, D. Suciu, and S. Hong. The boundary between privacy and utility in data publishing. In *Proceedings of the 33rd international conference on Very large data bases*, VLDB '07, pages 531–542. VLDB Endowment, 2007.

[32] W. Gao and H. Ai. Face gender classification on consumer images in a multiethnic environment. In *Proceedings of the Third International Conference on Advances in Biometrics*, 2009.

[33] A. Ashraf, S. Lucey, J. Cohn, T. Chen, Z. Ambadar, K. Prkachin, and P. Solomon. The painful face - pain expression recognition using active apprearance model. In *ICML*, 2009.

[34] G. Abowd, A. Bobick, I. Essa, E. Mynatt, and W. Rogers. The Aware Home: Developing Technologies for Successful Aging In *Proceedings of the AAAI Workshop on Automation as a Caregiver*, 2002.

[35] M. Osadchy,B. Pinkas,A. Jarrous,and B. Moskovich. SCiFI A System for Secure Face Identification. *IEEE Symposium Security and Privacy (SP)*,2010.

[36] Y. Huang, L. Malka, D. Evans, and J. Katz. Efficient Privacy-Preserving Biometric Identification. 18th Network and Distributed System Security Symposium (NDSS 2011), 2011.

[37] J. Bringer, H. Chabanne, M. Favre, A. Patey, T. Schneider, and Michael Zohner. GSHADE: faster privacy-preserving distance computation and biometric identification. In *Proceedings of the 2nd ACM workshop on Information hiding and multimedia security*, 2014.

[38] J. Bringer, H. Chabanne, and A. Patey. Privacy-Preserving Biometric Identification Using Secure Multiparty Computation: An Overview and Recent Trends. *IEEE Signal Processing Magazine*, 30(2):42–52, 2013.

[39] K. Simoens, J. Bringer, H. Chabanne, and S. Seys. A Framework for Analyzing Template Security and Privacy in Biometric Authentication Systems. *IEEE Transactions on Information Forensics and Security*, 7(2):833–841, 2012.

[40] B. Yang, L. Rajbhandari, C. Busch and X. Zhou. Privacy Implications of Identity References in Biometrics Databases. In *Proceedings of the IEEE Eighth International Conference on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP), Piraeus*, 2012.

[41] Privacy and Security in Biometrics, Editor: Patrizio Campisi, Springer, July 2013.

[42] Z. Jin, A. B.J. Teoh, T. S. Ong and C. Tee. Fingerprint Template Protection with Minutiae-based Bit-string for Security and Privacy Preserving. Expert Systems with Applications, 39(6): 61576167, 2012.