International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Conference Proceedings: Volume: 05 ACME – 2026 DOI: 10.55041/ISJEM.ACME131
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

Face Identification from Obfuscated Images in Deep Learning using Feature Compensation

MR. K. Perumal[1] *Assistant professor, Department of CSE, Annamacharya Institute of* Technology and sciences, Tirupati-517520,A.P, India perumalinfo@gmail.com

M Gowthami[2] *UG Student, Department of CSE, Annamacharya Institute of* Technology and sciences, Tirupati-517520,A.P, India. mallegowthami8@gmail.com

k Kavya[3] *UG Student,Department of CSE, Annamacharya Institute of* Technology and sciences, Tirupati-517520,AP, India. Kunanikavya906@ gmail.com

K Bhanu Prakash[4] *UG Student, Department of CSE, Annamacharya Institute of* Technology and sciences, Tirupati-517520,A.P, India. bhanuprakash02062001@gmail.com

B TharunKumar Reddy[5] *UG StudentDepartment of CSE, Annamacharya Institute of* Technology and sciences, Tirupati-517520,AP,India tharunkumarreddy12.06@gmail.com

***Abstract*—**The technology of face recognition has had immense positive changes in aspects like the security of the people and convenience to the user. But the fact that it is widely used has also generated serious issues on pri- vacy since it is quite simple to gather facial images and misuse them. It poses a big dilemma as, although the requirement of the use of a reliable face recognition system is high, people are becoming more and more unwilling to provide their original facial information.To deal with this problem, we suggest that we use is the PRO- Face C that is a privacy-sensitive face recognition system, able to maintain the required balance between privacy and accuracy. The suggested approach is based on the client-server approach in which obfuscated face image is sent over the server only by the client of the solution. Identity detection is then done with a pre- trained model along with privacy- free complementary features so that original face appearance is never actually obtained, yet the obfuscated images still provide good preview of an image and a mechanism of identity-guided feature compensation mechanism is used to further increase the accuracy of recog- nition. Moreover, there are a number of privacy-related measures, which are implemented to reinforce data security even more. A comprehensive test on various face recognition data sets indicates that the proposed system has very high recognition performance at minimal user privacy at different contexts.

Index Terms—Face recognition, privacy preservation, obfus- cated images, client–server architecture, feature compensation

I. INTRODUCTION

The most common applications of the face recognition technology are in authentication, access control and security of the public, which offer a huge enhancement of safety and user conve- nience. Its massive usage, however, has caused significant problems with regards to privacy, primarily because of the vast gathering, saving, and abuse of facial information. This makes it create one essential dilemma, and that is the fact that, as there is the growing demand of highly accurate face recognition systems, people are ever reluctant to provide their original biometric data to anyone[H].

Fig. 1. visual privacy for high accuracy

In a standard face recognition system, it is really very simple; a picture of the face is scanned as an input and is compared to the facial images that have been stored on a server to give a true or false result as to whether the identity is already present in the database or not. In case of a match, the system authenticates the identity, otherwise it grants a non- match. Though a very useful and widespread method of approach, the serious question of privacy of this approach poses serious issues since this method relies on the gathering, relaying and storage of real face images. Homefaces are very sensitive data of the person and their disclosure especially in remote systems over the cloud or through third party may result in unauthorized unsatisfied access, identity theft, or eventual misuse of the biometries. To mitigate such risks, privacy- conscious face recognition proposals are more often being created to utilize secured facial representations, as opposed to raw images. When using these systems it is the face image that is blurred or obfuscated at the source and remains

Identify applicable funding agency here. If none, delete this.

undisclosed to the end user providing that the identity of the individual being targeted cannot be easily shared. Even with a successful system an identity that is verified is only available in its obfuscated form and this greatly limits the chance of visual identity leakage. Notably, such secured images are not entirely disfigured, they do not lose as much structural and semantic information as they can be recognized reliably and perform simple visual tasks, e.g. preview image or image quality. This solution will allow a more responsible exploitation of face recognition technology by moving the focus away form raw pictures to some privacy- preserving representations. It enables the systems to be highly identification-accurate at the same time and maintain the privacy of their users, which is especially important in real-world applications where trust, data security, and regulatory compliance are gaining ascen- dancy. This leads to the fact that privacy and recognition performance are not played as conflicting aims anymore, but rather a compromise in the framework of a single and practical approach is made.

II. LITERATURE SURVEY

In this section, we review two major research directions closely related to our work and discuss their limitations, which motivate the design of our proposed approach.

A. Deep Learning–Based Face Identification from Privacy-Protected Images

To address the privacy concerns in face recognition sys- tems, extensive research has been done on the topic of face recognition with privacy- assured or obscured faces using deep learning technologies. The aim of these approaches is to conduct identity recognition in a manner such that no explicit revelation of the original facial appearance is made, commonly with the aid of transformed or auxiliary informa- tion to maintain recognition ability. The current approaches to this direction can be discussed in two broad groups. 1) Cryptographic Protection: These methods perform face recog- nition in real time in put-encrypted spaces. Secure crypto- graphic algorithms are used in important processes like feature extraction and similarity calculation. Examples of representa- tive techniques are garbled circuits, homomorphic encryption, secure multiparty compu- tation, functional encryption and stable hashing. The recognition accuracy of these methods is usually high with good theoretical security guarantees since they do not consist of any lossful operations. Nonetheless, encryption makes facial images conceptually incomprehensi- ble, doing away with visualization utility. In addition, these techniques are frequently associated with both computation and communication overheads, and are closely coupled with a particular recognition model, requiring restricted practical use.
2) Representation Transformation-Based Recognition:This is another stream of work that does face recognition using trans- formed representations of face image. Very little attention has been given to frequency-domain learning, in which identity-related information is stored using selective preservation of frequency components, and visual details are suppressed. A

Fig. 2. client server Architecture

number of different studies use discrete cosine transform (DCT) components, differential privacy, or client-server ar- chitectures to turn a privacy-recognition trade-off. Generative methods have been developed as well, which train deep models to identify faces on masked, synthesized, or reconstructed faces. Despite the fact that these methods bring more pri- vacy over direct image sharing, any substantiations likely to transform the image are bound to draw in information loss. Consequently, the recognition accuracy can be compromised and it has a tendency to limit visualization utility so such methods are inappropriate when the application needs image preview.
Overall, existing methods for face identification from privacy-protected images tend to prioritize recognition per- formance while sacrificing visual interpretability and system usability. B. Identifiable Face Anonymization Anonymiza- tion of faces has been extensively investigated as a way of safeguarding identity through changing the look of the face. Classical methods involve image blurring and face generation whereas more current techniques are able to use generative models to anonymize faces in a more aesthetically real way. One of these sets of works is devoted to the conservation of the recognizable features in anonymized pictures or the identification based on privacy-preserving faces. The existing research indicates that deep learning models were capable of identifying faces even in extreme degradation or occlusion, which prompts anonymization techniques that selectively re- tain machine-recognizable features. A few of them separate fa- cial identity and appearance, substitute sensitive properties or create virtual identities to afford secrecy and allow recognition at the same time. Visualization as a feature compared to the features of a protected-domain recognition meth- ods, anonymization-based methods have a clear benefit in that the resulting anonymized data can be hashed into a form that human beings can easily visualize and which is meaningful and presentable. Nonetheless, since anonymization is incom- patible with the identity discrimination process in question, these methods are generally less recognized. The use of face recognition is common as a secondary goal which results in a trade-off whereby visual privacy is conserved at the expense of recognition utility.

III. PROPOSED FRAMEWORK

A. Overall Framework Description

This paper introduces a deep learning model on face recog- nition of faces in visually-obfuscated faces through a feature compensation approach. The first goal is to maintain facial privacy, and at the same time, achieve good results in identity

recognition. This framework accepts a collaborative client server system. Client device does privacy protection on the face image that was captured and server is doing the iden- inference using privacy-preserved representation based only on the representations. In this process, the original face image does not get access to the server at any point and the privacy is highly guaranteed.

B. System Assumptions and Threat Considerations

The proposed system operates under the following assumptions:

- The architecture includes a client server paradigm in which the client is a natural resource-constrained object like a camera or mobile phone and the server has adequate computational capabilities.
- The original facial data is securely managed and is fully trusted by the client.
- Accidental that it is a server, the server is claimed to be honest-but-curious i.e. he/she can correctly perform the recognition process on the received data, but he/she can make the inferences about the visual data.

In this case the structure permits the protection of the facial identity despite possible inspection efforts at the server level.

C. Privacy-Aware Face Identification Mechanism

In order to achieve identity recognition without exposing sensitive visual data, the client breaks down the input face image into two privacy saving parts:

- An obfuscated image (that covers the information that could be related to identity, yet can still be used in image preview).
- A collection of deep incompatibilities that represent identity-relevant information without revealing face ap- pearance.

These two parts are combined in a pre- trained deep recognition model on a server and will also enable the server to be involved in face iden- tification and avoid reconstruction of the face of the original face.

D. Region-Adaptive Image Obfuscation

Visual privacy is achieved using a region-adaptive obfusca- tion strategy that combines multiple filtering techniques, such as Gaussian blurring, median filtering, and pixelation. These filters are applied across different spatial regions of the image, creating heterogeneous obfuscation patterns. This strategy offers two key benefits:

- Enhanced resistance against visual inspection and recon- struction attacks.
- Improved robustness and generalization during model training.

E. Extraction of Compensatory Feature Representations

Obfuscation eliminates important identity information so that the client derives compensatory feature representations of the difference between the original and the obfuscated images. It uses a lightweight deep network that efficiently computes multi- scale features hence the method is ideal when deployed in the real- world with resource-limited devices and is also able to provide precise inferences on the server side.

F. Patch-Level Feature Randomization

In order to stop further leakage of privacy to shallow fea- tures a patch level feature randomization approach is proposed. This process randomizes channels in features on a local patch of space, and greatly impairs visual interpretability. Moreover, this operation is also a key-dependent trans- formation so that it can be correctly identified only in case the same configuration is used both in training and inference.

G. Identity-Guided Feature Integration

The compensatory features on the server side are dynami- cally combined with intermediate representations obtained of the obfuscated image. This identity guided integration allows the recovery network to retrieve discriminative information that has been obfuscated. The training process is facilitated by:

- Discriminative embedding loss, to enforce discriminative embeddings.
- A feature alignment loss which promotes consistency with embeddings of original images.

H. Model Compatibility and Deployment Efficiency

The suggested framework will be compatible with the available deep face recognition models com- phatively. The client-side feature extractor and fusion modules only need to be trained, thus the full recognition backbone does not have to be retrained. Compensatory features are quantised or transmitted selectively to reduce the communication overhead to allow efficient data transfer with only slight impairment in identification performance.

EXPERIMENTAL EVALUATION AND ANALYSIS

In this section, a thorough analysis of the suggested feature-compensated face identification framework is made in privacy-conservant environments. The experiments both test reliability of identity recognition and visual privacy strength. Regular face verification procedures are utilized to perform recogni- tion analysis whereas the privacy protection is considered by quantitative measures and qualitative visual observations.

I. Experimental Protocol and Setup

Datasets and Evaluation Benchmarks: In order to develop and test the suggested framework, we are using the CelebA album comprised of about 202K face pictures of over 10K identities. The data has a great diversity in regards to fa- cial features and imaging conditions thus favoring learn- ing of strong representations with obfuscation privacy. To be evaluated, several commonly accepted face verification benchmarks are embraced in order to compare them fairly and extensively with current methodologies. These are LFW, CFP- FP, AgeDB, CPLFW, CALFW, IJB-B and IJB-C on the

conventional evaluation procedure. Most of the LFW-based datasets are claimed to have average verification accuracy, and the more difficult IJB benchmarks employ the true accept rate at a false accept rate of 0.01In order to match identifiable face anonymization (IDFA) lit- erature, further analyses are performed on CelebA, LFW and VGGFace2 test splits. All the verification experiments are used with the face.evoLVe framework, which gives the standardized pipelines and trained face recognition schemes.

1) Privacy Obfuscation Strategies: There are also 4 privacy-preserving transformations that are discussed in this paper: Gaus- sian blurring, median filtering, pixelation, and a hybrid obfus- cation method that integrates all three underlying filters. The hybrid obfuscation is only used during training, to make it more robust. Obfuscation strength parameter M is randomly selected amongst the set of the number 2, 3, 4, 5 and image patches are randomly selected among the varying filter outputs. The filter settings can be concluded as follows:

• Gaussian blur: fixed kernel size of 31 with the standard deviation sampled uniformly from the range [2, 8]. •Median filtering: kernel size randomly selected from odd integers between 7 and 19. •Pixelation: block size randomly chosen from integer values between 4 and 10. During inference, a moderate hybrid obfuscation setting is adopted to balance privacy preservation and recognition accu- racy. In addition, each base filter is evaluated independently to assess the generalization capability of the trained model. 2) Network Architecture and Training Details:Face images are all made to a resolution of 112 x 112. The recognition backbone on the server-side uses a ResNet-100 model pre-trained on AdaFace making it have high baseline recognition capability. The feature extractor adopted in the client side is MobileFaceNet which is efficient and consumes low-resource devices. The stochastic gradient descent (SGD) with a momentum and weight decay is used to train. Following hyperpa- rameter tuning, the trade-off coefficient is set to 1.0 which provides the best validation result. An effective patch-wise channel permutation (PCP) configuration of 2 x 2 is chosen and it offers good privacy improvement with minimum distortion to recognition capability. All the experiments are implemented based on PyTorch on the NVIDIA RTX A5000.

J. Face Recognition Performance

1) Comparison with Privacy-Preserving Recognition Meth- ods: The proposed approach is first compared with state-of- the- art privacy-preserving face recognition (PPFR) methods across standard benchmarks. While a marginal accuracy drop is observed compared to unprotected face recognition mod- els, the degradation remains below 0.5Notably, the model maintains strong performance even when tested with indi- vidual base obfuscations that were not explicitly used during training, demonstrating robustness and adaptability to unseen privacy transformations. 2) Comparison with Identifiable Face Anonymization Tech- niques: To ensure fair comparison with identifiable face anonymization (IDFA) methods, evaluations are conducted under identical experimental protocols. The

Fig. 3. Face Identification

results indicate that the proposed framework consistently out-performs exist- ing anonymization-based approaches in terms of verification accuracy. Unlike prior methods that preserve identifiable visual cues for recognition, the proposed approach achieves superior performance through complementary feature compensation with- out reintroducing perceptible facial de- tails. Visual comparisons further confirm the enhanced privacy preservation capability of the proposed method. 3) Cross- Domain Verification Capability: In real-world de- ployments, face verification often involves matching privacy- protected query images against unprotected enrollment tem- plates. To address this scenario, cross-domain verification experiments are conducted with the support of an embedding- level loss. The results demonstrate that the proposed method signifi- cantly outperforms existing approaches, confirming its ability to bridge the domain gap between protected and unprotected facial representations effectively.

K. Privacy Protection Assessment

1) Exposure Analysis on the Server Side: Visual inspection of all of intermediate data presented to the server, such as obfuscated images, complementary feature maps and fused representations, bear testify to the inability to recover any detail whatsoever of the faces. Direct recognition that operates with such representations produces catastrophic performance degradation which confirms the usefulness of the privacy protection mechanism suggested. 2) Patch-wise Channel Per- mutation Effectiveness: Patch-wise channel permutation (PCP) mechanism is an important mechanism in improving privacy. The feature visualizations indicate that PCP causes significant disturbance to the spatial coherence, rendering it impossible to visually interpret it. In addition, the failure of recognitions on mismatched PCP seeds during training and inference prove the secret-key-like functionality of the PCP mechanism. 3) Attack against Reconstruction: In order to measure the strength of the system to reconstruction attacks, high-quality image restoration models are tested with different defended repre- sentations. Blurred and median-filtered images are partially recoverable, but obfuscation by hybrid, as well as by PCP-processed features is much stronger, constructing images with lower quality, preserving privacy, yet removing information at the same time. 4) Utility Privacy Trade-off Analysis: Tests done with different levels of obfuscation strength display managed trade off between recognition accuracy and privacy preservation. The recognition accuracy is kept at over 96% even in the most challenging of the obfuscation settings, which also proves the feasibility of practicality of the proposed approach.

L.Generalization and Computational Analysis

1) Compatibility with Different Recognition Backbones: The framework is tested on several server-side recognition backbones such as ResNet-18 and ResNet-50. Regardless of the models, the recognition performance remains consistently high, and this proves that the proposed framework will use the established recognition systems as plug-and-play components. 2)Optimization of Communication overheads: Various compression and feature-selection strategies are studied in order to minimize the cost of data transmission. Close redundancies in data representation granted by feature quantization and selective dropout of select stages and minimal effects on recognition accuracy reveal the effective management of communication overheads. 3)Model Size and Inference Efficiency: Runtime The server side recognition network is the most dominant computational cost, despite its smaller size. Other elements, such as obfuscation, PCP, and feature compensation, create zero overheads and this makes the system applicable in real time implementations.

M. Ablation and Component Contribution Study

An extensive ablation study is conducted to quantify the contribution of each system component. Removing either the obfuscated images or the complementary features leads to severe drops in verification accuracy, confirming the neces- sity of feature compensation. Additional experiments validate the effectiveness of hybrid obfuscation, embedding loss, and weighted feature fusion. Overall, the ablation results strongly support the final design choices and demonstrate the robustness and effectiveness of the proposed framework.

IV. CONCLUSION

In this paper, we have discussed Face Identification of Obfuscated Images in Deep Learning Using Feature Comdisposition, a new framework that seeks to deal with the privacy- utility dilemma in face recognition and facial capturing systems. The proposed solution will show images that are face-based on these devices owned by a client, like a smartphone or a webcam, but are filtered through an intuitively applied privacy filter, which will give the images a visually preserved representation. Such obfuscated images are then augmented with a feature compensation mechanism which incorporates privacy preserving feature maps in the recognition mechanism. This allows the proper identification of faces and at the same time retains the original visual information intact. Most importantly, the obscurant images retain valuable visual information needed in low- level processes, including photo preview, but is not completely randomized. In addition, recognition performance is high with only a pre-trained recognition backbone at the server, providing no retraining or fine-tuning and making it applicable to a high number of diverse systems. A large set of experiments that has been performed on a variety of datasets, including a diversity of recognition and privacy conditions, prove that the suggested approach substantially exceeds the currently available privacy-sensitive face recognition (FFPR) and identifiable face anonymization (IDFA) systems and methods. A number of limitations are, however, still there. Although the hybrid obfuscation solution offers better protection of the way of traditional filters, it is more or less susceptible to reconstruction attacks, an unsurprising claim that is that additional sophisticated techniques of ob- fuscation might increase further privacy. Also, communicative and storage overhead is added due to the complementary fea- ture maps, which highlights the necessity of a more efficient representation. The available framework is mostly general and lacks the specifics of extreme or problematic samples, which may serve as a chance in future research. Lastly, the work focuses on visual privacy protection in the process of infer- ence but it does not take into account the possibilities of threat to privacy related to face embeddings or training data that should be regarded in the wider face recognition systems.

REFERENCES

[1] [1] A. Satariano, “Police Use of Facial Recognition is Accepted by British Court,” The New York Times, 2019.
[2] [2] K. Hill, “The Secretive Company That Might End Privacy as We Know It,” The New York Times, 2020.
[3] [3] T. Winkler and B. Rinner, “TrustCAM: Security and privacy- protection for an embedded smart camera based on trusted computing,” in Proc. IEEE Int. Conf. Advanced Video and Signal Based Surveillance, 2010, pp. 593–600.
[4] [4] H. Wu et al., “PECAM: Privacy-enhanced video streaming and analytics via securely-reversible transformation,” in Proc. ACM Int. Conf. Mobile Computing and Networking, 2021, pp. 229–241.
[5] [5] S. Shan, E. Wenger, J. Zhang, H. Li, H. Zheng, and B. Y. Zhao, “Fawkes: Protecting privacy against unauthorized deep learning models,” in Proc. USENIX Security Symposium, 2020, pp. 1589–1604.
[6] [6] V. Mirjalili, S. Raschka, and A. Ross, “PrivacyNet: Semi-adversarial networks for multi-attribute face privacy,” IEEE Transactions on Image Processing, vol. 29, pp. 9400–9412, 2020.
[7] [7] X. Yang et al., “Towards face encryption by generating adversarial identity masks,” in Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV), 2021, pp. 3877–3887.
[8] [8] S. Hu et al., “Protecting facial privacy via adversarial identity masks using style-robust makeup transfer,” in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2022, pp. 14994–15003.
[9] [9] Y. Zhong and W. Deng, “OPOM: Customized invisible cloak towards face privacy protection,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 45, no. 3, pp. 3590–3603, 2023.
[10] [10] L. Yuan, P. Korshunov, and T. Ebrahimi, “Privacy-preserving photo sharing based on a secure JPEG,” in Proc. IEEE Conf. Computer Communications Workshops, 2015, pp. 185–190.
[11] [11] J. Zhou et al., “Face template protection through residual learning- based error-correcting codes,” in Proc. Int. Conf. Control, Computing and Vision, 2021, pp. 112–118.
[12] [12] G. Mai, K. Cao, X. Lan, and P. C. Yuen, “SecureFace: Face template protection,” IEEE Transactions on Information Forensics and Security, vol. 16, pp. 262–277, 2021.
[13] [13] V. M. Patel, N. K. Ratha, and R. Chellappa, “Cancelable biometrics: A review,” IEEE Signal Processing Magazine, vol. 32, no. 5, pp. 54–65, 2015.
[14] [14] H. Hukkelas, R. Mester, and F. Lindseth, “DeepPrivacy: A gener- ative adversarial network for face anonymization,” in Proc. Int. Symp. Visual Computing, 2019, pp. 565–578.
[15] [15] M. Maximov, I. Elezi, and L. Leal-Taixe´, “CIAGAN: Conditional identity anonymization generative adversarial networks,” in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2020, pp. 5446–5455.
[16] [16] D. Osorio-Roig et al., “Stable hash generation for efficient privacy-preserving face identification,” IEEE Transactions on Biometrics, Be- havior, and Identity Science, vol. 4, no. 3, pp. 333–348, 2022.
[17] [17] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, “VGGFace2: A dataset for recognising faces across pose and age,” in Proc. IEEE Int. Conf. Automatic Face and Gesture Recognition, 2018,
pp. 67–74.

[18] [18] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2014, pp. 1701–1708.
[19] [19] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2015, pp. 815–823.
[20] [20] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angu- lar margin loss for deep face recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2019, pp. 4690–4609.
[21] [21] H. Wang et al., "CosFace: Large margin cosine loss for deep face recog- nition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2018, pp. 5265–5274.
[22] [22] Z. Zhao et al., "Feature disentanglement for privacy-preserving face recognition," Pattern Recognition, vol. 108, 2020.
[23] [23] H. Zhang et al., "Privacy-aware face recognition using complemen- tary feature learning," Neurocomputing, vol. 470, pp. 64–76, 2022.