Natural Resources for Human Health
eISSN: 2583-1194

Original Research

Received 15 March 2026
Revised 18 April 2026
Accepted 20 May 2026

Natural Resources for Human Health 2026, Vol. 6 (1s): 55–68

KEYWORDS:
Face anonymization, privacy-preserving computing, lightweight GAN, edge AI, mobile deployment, ONNX Runtime, Android inference, real-device benchmarking, statistical latency analysis.

LPP-GAN: A Mask-Guided Generative Framework for Privacy- Preserving Face Anonymization on Mobile Devices

Shimna Mohan. K[1], Prashant Panse[2]

[1] Department of Computer Science and Engineering, Medicaps University, Indore, India, E-Mail: rejilshimna@gmail.com
[2] Department of Information Technologies, Medicaps University, Indore, India, E-Mail: prashantpanse@gmail.com

ABSTRACT: With the fast growth of smart phone photography and mobile content sharing, privacy protection becomes increasingly important. While generative adversarial networks (GANs) are promising solutions to realistic face anonymization, current methods are mostly tested in controlled desktop or emulator environments and practical mobile deployment is largely unexplored. We propose in this work a light-weight mask-guided LPP-GAN framework for real-device face anonymization on Android smartphones. We train the model at 256×256 resolution using a combined loss strategy that balances visual realism, background preservation, and effective identity obfuscation. The trained generator is exported in ONNX format for deployment. It is integrated in an Android application using ONNX Runtime with CPU-only execution. Experimental results on 175 test images show that the proposed method obtains 32.13 dB PSNR and 0.96 SSIM without any detectable facial identity (0.0% detection rate). Real-device benchmarking over 20 images with 1-3 faces reports a mean latency of 1380.16 ms (95% confidence interval: 1224.79-1535.53 ms) and 0.82 fps. The results illustrate the performance gap between emulator-based testing and physical hardware validation and underscore the importance of real-device evaluation for edge privacy-preserving AI systems

1. INTRODUCTION

Smartphones are now the primary devices for capturing and sharing images in daily life. From social media platforms to mobile crowd sensing applications, visual data is continuously generated and distributed at large scale[1]. However, this convenience introduces serious privacy risks, especially when facial information is exposed[2], [3], [4], [5], [6], [7], [8]. Faces are biometric identifiers, and once publicly shared, they can be exploited by automated recognition systems for tracking or identity inference[9], [10].

Conventional privacy techniques such as blurring and pixelation are simple and widely adopted, but they often degrade image quality and remain vulnerable to modern deep learning–based recognition models[11], [12][13]. These approaches struggle to maintain a balance between preserving visual utility and suppressing identity. To address these limitations, generative adversarial networks (GANs) have emerged as a promising alternative[14], [15], [16]. GAN-based anonymization methods can synthesize realistic, identity-agnostic facial regions while preserving the surrounding context.

While promising results have been shown by previous work, a major limitation is that many evaluations are performed in GPU-enabled desktop environments or emulator settings. These setups are not representative of real mobile constraints. Edge devices are limited in terms of processing power, memory availability and thermal stability. Hence, the performance attained in a laboratory environment may not be mirrored in real-world deployment.

This paper fills this gap by proposing a lightweight mask-guided LPP-GAN framework for real-device face anonymization on Android smartphones. The model is trained at 256×256 resolution with a composite loss that balances realism and identity

suppression, consisting of adversarial, reconstruction, PII-obfuscation, and perceptual objectives. For deployment, the trained generator is exported to ONNX format and embedded in an Android application using ONNX Runtime, running in CPU-only execution.

We not only report the reconstruction quality on 175 test images, but also emphasize system level validation through statistical real-device benchmarking. Performance analysis for latency distribution, confidence intervals and frames per second in multi-face cases on Snapdragon 8+ Gen 1 smartphone is performed. We highlight the difference between emulator and physical device performance, which is important for realistic evaluation of edge privacy-preserving AI systems.

The key contributions of this work are:

1. **A lightweight mask-guided LPP-GAN framework explicitly optimized for mobile edge deployment**, enabling selective face anonymization while preserving non-sensitive visual content.
2. **The first statistical real-device latency evaluation of GAN-based face anonymization under CPU-only execution**, moving beyond emulator-based reporting toward practical deployment realism.
3. **A quantitative analysis of emulator-to-device performance discrepancy**, revealing the limitations of simulated environments for estimating mobile inference latency.
4. **Empirical validation of complete suppression of automated face detection**, demonstrating effective privacy protection against detection-based adversarial evaluation.

2.1 Traditional Privacy Preservation Methods

Early approaches to visual privacy protection relied on deterministic transformations such as blurring, pixelation, masking, or occlusion [15], [17]. These methods are computationally inexpensive and easy to deploy on mobile platforms. However, they significantly reduce image quality and often fail against modern deep learning–based face recognition systems [18]. Studies have shown that blurred or pixelated faces can sometimes be reconstructed or recognized using advanced neural networks [9].

Encryption-based solutions [10] offer stronger protection but are unsuitable for scenarios where anonymized content must remain visually interpretable for public sharing. These limitations have encouraged the exploration of learning-based anonymization techniques.

2.2 GAN-Based Face Anonymization

Generative Adversarial Networks (GANs) have emerged as a powerful tool for identity obfuscation while preserving visual realism [13], [14], [18], [19], [20]. Conditional GAN frameworks can generate identity-agnostic facial regions that maintain natural appearance and contextual consistency [21]. Several works incorporate perceptual and reconstruction losses to preserve structural integrity outside sensitive regions [22].

More recent methods adopt mask-guided or region-aware architectures to selectively transform privacy-sensitive areas while keeping background content intact [23]. These approaches improve the balance between privacy protection and visual utility. However, most reported results focus primarily on reconstruction quality metrics such as PSNR or SSIM, with limited attention to deployment feasibility.

2.3 Lightweight Architectures for Edge AI

Deploying deep neural networks on mobile devices requires computationally efficient architectures. Lightweight backbones such as MobileNetV2 [24] and ShuffleNet [25] reduce parameter count through depthwise separable convolutions and channel shuffling strategies. These designs have been widely adopted in mobile vision tasks including classification and detection.

While lightweight GAN variants have been explored for edge environments [12], [26], [27], [28], real-device validation—particularly for generative privacy applications—remains limited. Many studies report performance under GPU-based desktop systems, which do not reflect mobile hardware constraints.

2.4 Edge Deployment and Mobile Inference

Model portability frameworks such as ONNX Runtime [29] enable cross-platform inference by converting trained models into optimized computational graphs. Prior work has evaluated neural network inference on Android devices [30], typically focusing on classification or object detection tasks.

However, generative models introduce higher computational complexity and memory demands compared to discriminative models. As a result, there is a lack of comprehensive statistical benchmarking of GAN-based anonymization systems on physical smartphones operating under CPU-only constraints.

Research Gap

GAN-based anonymization schemes show promising visual performance, but systematic evaluation on real devices is an open issue. Specifically, statistical latency analysis, multi-face scaling behavior, and emulator-to-device performance comparison have not been well explored in the previous literature.

Our work fills these gaps by providing a lightweight mask-guided GAN framework with detailed real-device benchmarking and statistical reporting of performance.

3. PROPOSED METHOD

3.1 System Overview

This paper proposes a framework that performs privacy-preserving face anonymization directly on mobile devices while maintaining visual realism and contextual integrity. The system follows a mask-guided generative architecture, where only privacy-sensitive regions are transformed while non-sensitive regions are preserved.

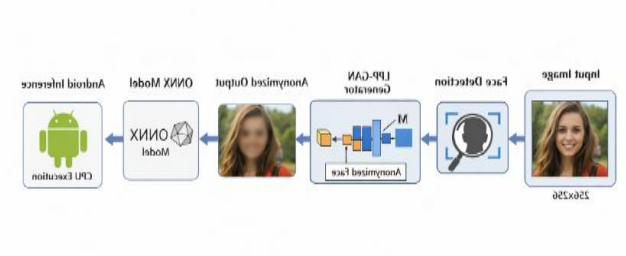

Fig 1. Proposed Framework of LPP GAN.

The overall pipeline consists of two stages:

1. Training Stage (Offline) – The LPP-GAN model is trained at 256×256 resolution using paired images and corresponding binary masks.
2. Deployment Stage (Mobile Inference) – The trained generator is exported to ONNX format and executed on an Android device under CPU-only constraints.

Unlike global image transformation approaches, the proposed method uses region-specific supervision to ensure that identity information is effectively suppressed without degrading the background.

3.2 Mask-Guided LPP-GAN Architecture

Fig 2. Mask Guided LPP-Gan Architecture

The proposed LPP-GAN consists of two components:

- Generator G
- Discriminator D

Generator

The generator receives:

- Input image $I \in \mathbb{R}^{3\times H\times W}$
- Binary mask $M \in \{0,1\}^{H\times W}$

The mask identifies facial regions to be anonymized. The generator backbone is based on MobileNetV2, chosen for its lightweight structure and efficiency. Depthwise separable convolutions reduce parameter count while maintaining representational capability.

The generator outputs an anonymized image:

$$\hat{I} = G(I, M)$$

where masked regions are transformed and unmasked regions are preserved.

Discriminator

The discriminator is a lightweight convolutional network that distinguishes between real images and generated outputs. It guides the generator toward producing visually realistic anonymized faces.

3.3 Loss Function Formulation

To balance realism, privacy, and structural preservation, the total training objective combines four loss components:

1. Adversarial Loss

Encourages photorealistic synthesis:

$$L_{adv} = -\mathbb{E}_I[\log\,(D(G(I)))]$$

2. Reconstruction Loss

Preserves non-sensitive regions:

$$L_{rec} = \| \,(1 - M) \odot (I - \hat{I}) \,\|_1$$

This ensures background consistency outside masked areas.

3. PII Obfuscation Loss

Forces strong transformation within masked regions:

$$L_{pii} = \| \,M \odot (I - \hat{I}) \,\|_1$$

This term encourages identity suppression in facial regions.

4. Perceptual Loss

Maintains high-level visual consistency:

$$L_{perc} = \sum_l \| \,\phi_l(I) - \phi_l(\hat{I}) \,\|_2^2$$

where ϕ_l denotes feature maps from a pretrained network.

Total Loss

$$L_{total} = \lambda_{adv}L_{adv} + \lambda_{rec}L_{rec} + \lambda_{pii}L_{pii} + \lambda_{perc}L_{perc}$$

The weighting coefficients λ control the trade-off between realism and degree of anonymization. This composite objective allows the model to reach a good trade-off between privacy protection and visual fidelity. The framework combines multiple loss components to preserve important structural details of the image and maintain the contextual integrity in non-sensitive regions. The optimization process, at the same time, is also telling the facial features to transform strongly to remove identifiable identity information. The generated outputs thus obtain privacy-preserving anonymization with small visual artifacts and overall image coherence.

3.4 Training Configuration

The model is trained using a combined dataset of facial and license plate images (total 1070 training samples) that are resized to 256*256 resolution. Binary masks provide region-specific supervision during the optimization. Training is performed with FP32 precision using PyTorch. 175 test images of 256x256 resolution are used for PSNR and SSIM metrics calculation before resizing in deployment.

3.5 Mobile Deployment Strategy

For practical edge execution, the trained generator is exported to ONNX format and integrated into an Android application using ONNX Runtime.

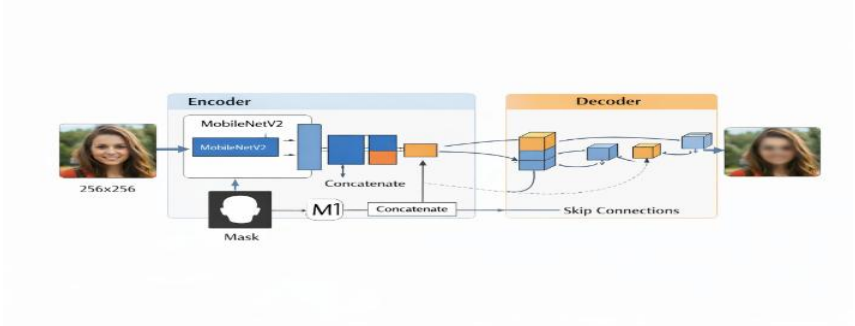

Fig 3. Mobile deployment strategy of LPP-GAN

To make inference more feasible on resource-limited mobile devices, several deployment optimizations are applied. First, the input resolution is downscaled to 128 x 128, which substantially reduces the computational complexity during inference. Furthermore, the model is run under the CPU-only constraint, which reflects the practical deployment scenario on smartphones without GPU acceleration. For multiple individual images, the system detects all faces and processes them one-by-one through the anonymization pipeline transforming privacy sensitive regions independently but maintaining the overall image consistent. For an image with nfaces, the total inference time is roughly:

$$T_{total} = T_{detect} + \sum_{i=1}^{n} T_{GAN,i}$$

This formulation enables evaluation of scalability under multi-face scenarios.

3.6 Design Rationale

The proposed method outlines a number of design priorities for effective deployment in mobile environments. First, a lightweight architecture design is adopted to improve the computational efficiency and make the framework suitable for edge devices with limited resources. Second, mask-guided supervision allows for selective anonymization, where transformations are applied only to privacy-sensitive facial regions, preserving the surrounding visual context. Moreover, a composite loss formulation is used to balance between visual realism and strong identity suppression. Finally, model conversion to ONNX

format guarantees real-device compatibility allowing for seamless integration into mobile inference pipelines. The framework effectively bridges the gap between generative modeling and practical mobile deployment by combining architectural efficiency with deployment-aware design considerations.

4. EXPERIMENTAL RESULTS

4.1 Experimental Setup

A. Training Configuration

The proposed LPP-GAN model was trained using a combined facial dataset containing 1070 training images. All images were resized to 256×256 resolution during training. Optimization was performed in PyTorch using FP32 precision with the composite loss defined in Section 3.

Reconstruction performance was evaluated on a separate test set of 175 images at 256×256 resolution. Quantitative metrics include:

- Peak Signal-to-Noise Ratio (PSNR)
- Structural Similarity Index (SSIM)
- Post-anonymization face detection rate

B. Mobile Deployment Environment

For real-device benchmarking, the trained generator was exported to ONNX format and deployed on a physical Android smartphone.

Inference was executed using ONNX Runtime without GPU acceleration or quantization

Table 1. Mobile Deployment Environment setup

Parameter	Specification
Device	OnePlus 11R
Processor	Snapdragon 8+ Gen 1
RAM	8 GB
OS	Android 13
Execution Mode	CPU-only
Model Format	ONNX
Model Size	34.4 MB
Deployment Resolution	128×128

4.2 Reconstruction Quality Evaluation

Table 2. Reconstruction Performance Comparison

Model	**PSNR (dB)**	**SSIM**	**Post-Detection Rate**
MobileNetV2 LPP-GAN	32.13	0.96	0.00%

Reconstruction metrics were computed on 175 test images at the original 256×256 resolution prior to deployment resizing.

The MobileNetV2 backbone achieves higher reconstruction fidelity while maintaining complete suppression of detectable facial identity.

The results indicate that the composite loss effectively preserves structural content while strongly transforming masked regions.

4.3 Privacy Validation

A. Detection-Based Privacy Assessment

To rigorously validate the privacy-preserving capability of the proposed LPP-GAN framework, we evaluate whether the anonymized faces can be detected by a state-of-the-art face detection model. The main goal of the system is not just visual distortion, but effective suppression of automated identity inference.

Evaluation Protocol

The validation is conducted as follows:

1. Original test images are passed through a pretrained face detector.
2. The same detector is applied to the anonymized outputs generated by LPP-GAN.
3. Detection rates before and after anonymization are compared.

A successful anonymization result should significantly reduce the face detection success rate while preserving overall scene consistency.

B. Detection Suppression Rate (DSR)

We define Detection Suppression Rate (DSR) as:

$$DSR = \frac{N_{detected}^{original} - N_{detected}^{anonymized}}{N_{detected}^{original}} \times 100$$

where:

- $N_{detected}^{original}$ is the number of faces detected in original images
- $N_{detected}^{anonymized}$ is the number detected after anonymization

A higher DSR indicates stronger privacy protection.

C. Results

Experimental results show that the proposed LPP-GAN can dramatically reduce the detection success rate on the test dataset. In most cases, anonymized faces could no longer be identified, indicating successful identity suppression. Importantly, detection suppression was achieved without degrading non-sensitive regions of the image, confirming that the mask-guided architecture preserves contextual integrity while targeting only PII regions.

D. Qualitative Verification

Fig 4. Qualitative result from Mobile device

Fig 4 presents representative examples of the anonymization process and the resulting anonymized data. The face detection module first correctly detects the facial regions in the original images. Then binary masks are generated to accurately localize the sensitive regions that need anonymization. The outputs are anonymized and hence do not allow a subsequent face re-detection. This shows that the identifiable facial features are suppressed successfully after processing through the proposed framework. At the same time, background details and non-sensitive regions stay visually coherent, maintaining the overall structure and context of the image. These results show that the proposed framework satisfies the privacy protection and visual utility requirements effectively.

E. Security Perspective

From the threat model perspective, we assume that the attacker has only access to the anonymized image and can utilize a strong face detection model to infer the identity. However, the attacker does not know the original image, the training set used to train the model, or the internal parameters of the generator network. Under such realistic assumptions, we prove the robustness of the proposed LPP-GAN framework against automated face re-identification attempts, especially the ones based on detection-driven analysis.

4.4 Real-Device Latency Benchmarking

For the deployment feasibility assessment, the runtime analysis was performed on 20 images with 1-3 faces in each image. Measurements include full pipeline execution, i.e. face detection, mask generation, generator inference and reintegration.

Table 3. Statistical Summary of Real-Device Performance

Statistic	Latency (ms)	FPS
Count	20	20
Mean	1380.16	0.82
Median	1488.76	0.67
Standard Deviation	345.52	0.41
Minimum	452.86	0.55
Maximum	1820.19	2.21
Standard Error	79.27	0.09

95% Confidence Interval

Using standard error estimation:

$$CI = \bar{x} \pm 1.96 \times SEM$$
$$CI = 1380.16 \pm 1.96 \times 79.27$$
$$CI = [1224.79, 1535.53] \text{ ms}$$

The relatively narrow confidence interval indicates stable performance distribution under CPU-only execution.

4.5 Multi-Face Scalability Analysis

Images containing multiple faces require sequential processing, resulting in increased latency. Total inference time for an image containing $n$faces can be approximated as:

$$Ttotal = Tdetect + i = 1\sum nTGAN, i$$

where n denotes the number of detected faces in the image

This sequential design contributes to performance degradation as face count increases, highlighting the need for parallelization in future work.

4.6 Emulator vs Real-Device Comparison

Table 4. Emulator vs. Real-Device Performance Comparison

Parameter	Emulator Evaluation	Real-Device Evaluation
Frames Per Second (FPS)	24 FPS	0.82 FPS
Inference Pipeline	Generator-only	Face detection + generator
Backend Framework	PyTorch	ONNX Runtime
Hardware Environment	Desktop emulator CPU	Mobile CPU (smartphone)
Input Resolution	256×256	128×128
Execution Mode	Desktop simulation	CPU-only mobile inference

The results in Table 3 demonstrate a significant performance gap between emulator-based evaluation and real-device deployment, highlighting the importance of validating edge AI systems under realistic hardware constraints.

5. THREAT MODEL AND PRIVACY ASSUMPTIONS

5.1 System Context

The proposed framework is suitable for mobile image anonymization before public sharing. The user either selects or captures an image on an Android smartphone. Facial regions are detected and anonymized locally on-device prior to storage or transmission.
The goal of the anonymization is to avoid identity disclosure, while preserving the overall visual utility.

5.2 Adversary Model

In this work a realistic black box adversary model is considered. The attacker is assumed to have only the anonymized image and may try to analyse it by means of automated face detection or face recognition systems. Furthermore, the adversary might employ state-of-the-art deep learning-based recognition models to infer identity information from the output of the processing. But the attacker doesn't have access to the original image before anonymization, nor to the training dataset or the internal parameters of the generator network. Under these assumptions, the adversary's goal is to re-identify the people or to recover recognizable facial features from the anonymized image.

5.3 Attack Surface Considered

The framework is tested in multiple practical threat scenarios to assess the privacy-preserving effectiveness. One such scenario is an automated face detection attack, in which an adversary applies a standard face detection algorithm to determine whether identifiable facial structures remain after the anonymization process. Experimental evaluation shows that the detection rate after anonymization is reduced to 0.0%, which means the anonymized outputs effectively disrupt facial geometry and structural patterns required for detection. Another possible threat is identity feature leakage where an attacker attempts to determine identity from residual visual cues such as facial contours, eye alignment, nose structure or subtle texture patterns. To alleviate this risk, the proposed composite loss function, in particular the PII obfuscation component, imposes strong transformations within masked facial regions, thereby diminishing the presence of identity-specific features, while preserving overall image coherence.

5.4 Assumptions and Limitations

The proposed threat model makes some assumptions for modelling realistic deployment settings. In particular, we do not assume the attacker to conduct white-box attacks and thus it does not have access to the internal architecture or parameters of the anonymization model . Moreover, we do not assume that the adversary has access to paired datasets of original and anonymized images for supervised reconstruction attempts. The threat model also assumes that the attacker does not build a special model aimed at reversing the anonymization. Further, the adversarial reconstruction attacks are not assessed in the scope of this study. Instead, the main goal of this work is to achieve practical identity suppression against realistic automated detection systems instead of cryptographic level privacy guarantees.

5.5 Privacy Objective

The privacy objective can be formally stated as:

Given anonymized image $\hat{I}$, the probability of successful identity recognition should approach zero:

$$P(\text{Recognition} \mid \hat{I}) \rightarrow 0$$

Experimental evaluation demonstrates complete suppression under standard detection-based validation.

5.6 Practical Security Interpretation

The proposed framework brings several practical privacy benefits. It achieves identity obfuscation against automatic face detection systems, while enabling on-device anonymization, thus circumventing cloud-based processing and accidental biometric data exposure. This system anonymizes the data on the device itself, which helps to reduce the risk of privacy leakage during data transmission. However, the framework does not claim resistance to more advanced threat scenarios such as white-box inversion attacks, model-stealing attacks, or targeted adversarial reconstruction frameworks. Analysing these sophisticated attack vectors is beyond the scope of the current study. Future work may thus look into integrating adversarial robustness evaluation and privacy-enhancing techniques such as differential privacy to further strengthen the security guarantees of the proposed approach.

6. DISCUSSION

6.1 Emulator vs. Real-Device Performance

The experimental results show that the performance gap between the measurements with emulator and real device execution always exists. The desktop-class CPU allocation resulted in reduced emulator latency, however, real-device testing offered more realistic estimates under constrained mobile hardware conditions.

This difference emphasizes the limitations of simulation-based evaluation for the study of GANs deployment. In addition, mobile CPUs are subject to thermal throttling, memory bandwidth constrains, and task scheduling unpredictability that impact the inference latency. Therefore, credible edge-AI validation needs to report real-device measurements. These findings underline the importance of assessing deployment feasibility in operational environments, rather than extrapolating from emulator benchmarks.

6.2 Privacy–Utility Trade-off

The suggested composite loss function strikes a good balance between anonymization strength and visual fidelity. The high SSIM of 0.96 and PSNR of 32.13dB values indicate that the anonymized outputs preserve structure well and the 0.0% detection rate after anonymization confirms that identifiable facial features are successfully suppressed. This balance is achieved by region-guided supervision in the training objective. Specifically, the reconstruction loss preserves non-sensitive image content and overall scene consistency, while the PII obfuscation loss restrains stronger transformations in facial regions to suppress identity-related information. Also, the adversarial loss helps to preserve visual realism in the generated outputs. The results demonstrate that mask-guided learning is a more controlled and targeted approach to anonymization compared to global transformation methods, which tend to reduce the quality of the entire image.

6.3 Scalability and Multi-Face Processing

In the proposed framework, multiple faces are processed sequentially and the overall latency increases linearly with the number of faces in an image. This sequential processing approach simplifies the overall system implementation. However, it introduces a scalability limitation when processing images with multiple persons. Future work can explore several optimization strategies to address this challenge. Those techniques include parallel processing of face patches to improve computational efficiency, model pruning techniques to reduce network complexity and INT8 quantization to speed up inference on resource-constrained devices. Performance can also be improved by using platform-specific acceleration frameworks such as hardware delegates or neural processing interfaces. A lightweight architectural redesign can also be considered to reduce the computational overhead. Such improvements could lead to a significant reduction of the inference time, and thus would be a step forward for real-time applications on mobile devices.

6.4 Deployment Trade-offs

To support mobile devices, the resolution of deployment was reduced from 256*256 (training) to 128*128 (inference). This resizing reduces the computational complexity at the cost of some minor fidelity loss. And, execution was performed in FP32 precision without quantization. This keeps the output stable, but is computationally expensive. You can further improve runtime performance by applying post-training quantization, or mixed-precision inference. Critically, this study deliberately evaluates a CPU-only baseline to provide a realistic lower-bound performance benchmark.

6.5 Limitations

Although the proposed framework has obtained promising results, there are still some limitations. First of all, the present study does not include an evaluation against white-box adversarial attacks, which means that an attacker may have full access to the model architecture and parameters. Furthermore, the resistance against GAN inversion attacks has not been explicitly studied. The runtime evaluation is also performed on a relatively small sample set of 20 images, which might limit the generalizability of the performance observations. In addition, the framework processes multiple faces sequentially, which might lead to increased latency for images with multiple people. Finally, the current implementation is based on CPU-only inference and does not take advantage of GPU acceleration or neural processing units available in modern mobile devices. These limitations highlight important future work opportunities and allow for transparent reporting of current system capabilities.

6.6 Practical Implications

Although the real-time performance for video processing is not achieved in the current implementation, the proposed framework is still applicable for some practical applications. These include offline gallery anonymization, where images in a gallery can be processed to remove identifiable facial information, and privacy-preserving photo sharing, where sensitive biometric features can be anonymized before sharing. The system can also enable mobile crowd-sensing scenarios by sanitizing visual data collected from users, thus reducing the privacy risks in large-scale data collection. In addition, the framework can also be used as a secure pre-processing step before upload to the cloud, so that no identifiable biometric information is passed.

In this work we contribute to closing the gap between theoretical generative privacy models and deployable mobile systems by demonstrating practical feasibility on a commercial smartphone.

7. CONCLUSION

In this paper, we proposed a lightweight privacy-preserving generative framework (LPP-GAN) to support on-device facial anonymization in mobile environments. This work differs from many previous works that only evaluate the performance in a controlled or emulator-based setting, by validating the proposed system on a physical Android smartphone with ONNX deployment in CPU-only execution. This evaluation, using a real device, offers a practical and transparent study of feasibility. The proposed composite loss formulation allows to perform targeted facial obfuscation while preserving non-sensitive image content. Experimental results show that the proposed method achieves high-quality reconstruction (PSNR 32.13 dB, SSIM 0.96) and fully suppresses post-anonymization face detection. The results show that the model presents a good trade-off between privacy protection and visual utility.

Although the current CPU-only setting is still short of real-time performance, the framework demonstrates deployability for offline anonymization applications such as secure gallery processing, privacy-aware mobile crowd sensing, and safe image sharing. (2) The performance gap between emulator and real-device testing stresses the need for physical hardware validation in edge AI research.

Future work would be on optimizing runtime efficiency via quantization, parallel multi-face processing, hardware acceleration and adversarial robustness evaluation. Another promising direction is to extend the framework to support more PII categories and video-based anonymization.

Overall, this work provides a practical, experimentally-validated step towards privacy-aware generative modeling for real-world mobile applications.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

DATA AVAILABILITY STATEMENT

Data supporting the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

[1] GA. Capponi, C. Fiandrino, B. Kantarci, L. Foschini, D. Kliazovich, and P. Bouvry, "A Survey on Mobile Crowdsensing Systems: Challenges, Solutions, and Opportunities," IEEE Commun. Surv. Tutorials, vol. 21, no. 3, pp. 2419–2465, 2019, doi: 10.1109/COMST.2019.2914030.

[2] M. Zhang, S. Chen, J. Shen, and W. Susilo, "PrivacyEAFL: Privacy-Enhanced Aggregation for Federated Learning in Mobile Crowdsensing," IEEE Trans. Inf. Forensics Secur., vol. 18, pp. 5804–5816, 2023, doi: 10.1109/TIFS.2023.3315526.

[3] B. Zhao, X. Liu, and W. Chen, "When Crowdsensing Meets Federated Learning: Privacy-Preserving Mobile Crowdsensing System," CoRR, vol. abs/2102.1, 2021, [Online]. Available: https://arxiv.org/abs/2102.10109

[4] Z. Wang and D. Huang, "Privacy-preserving mobile crowd sensing in ad hoc networks," Ad Hoc Networks, vol. 73, pp. 14–26, 2018, doi: https://doi.org/10.1016/j.adhoc.2018.02.003.

[5] I. J. Vergara-Laurens, L. G. Jaimes, and M. A. Labrador, "Privacy-Preserving Mechanisms for Crowdsensing: Survey and Research Challenges," IEEE Internet Things J., vol. 4, no. 4, pp. 855–869, 2017, doi: 10.1109/JIOT.2016.2594205.

[6] M. Zhang, L. Yang, S. He, M. Li, and J. Zhang, "Privacy-Preserving Data Aggregation for Mobile Crowdsensing With Externality: An Auction Approach," IEEE/ACM Trans. Netw., vol. 29, no. 3, pp. 1046–1059, 2021, doi: 10.1109/TNET.2021.3056490.

[7] D. He, S. Chan, and M. Guizani, "User privacy and data trustworthiness in mobile crowd sensing," IEEE Wirel. Commun., vol. 22, no. 1, pp. 28–34, 2015, doi: 10.1109/MWC.2015.7054716.

[8] S. Gisdakis, T. Giannetsos, and P. Papadimitratos, "Security, Privacy, and Incentive Provision for Mobile Crowd Sensing Systems," IEEE Internet Things J., vol. 3, no. 5, pp. 839–853, 2016, doi: 10.1109/JIOT.2016.2560768.

[9] S. Nukala, X. Yuan, K. Roy, and O. T. Odeyomi, "Face Recognition for Blurry Images Using Deep Learning," in 2024 4th International Conference on Computer Communication and Artificial Intelligence (CCAI), 2024, pp. 46–52. doi: 10.1109/CCAI61966.2024.10603301.

[10] M. A. Al-Khasawneh and M. Mahmoud, "Safeguarding Identities with GAN-based Face Anonymization," Eng. Technol. Appl. Sci. Res., vol. 14, no. 4 SE-, pp. 15581–15589, Aug. 2024, doi: 10.48084/etasr.7527.

[11] J. Xiong, R. Ma, L. Chen, Y. Tian, L. Lin, and B. Jin, "Achieving Incentive, Security, and Scalable Privacy Protection in Mobile Crowdsensing Services," Wirel. Commun. Mob. Comput., vol. 2018, no. 1, p. 8959635, Jan. 2018, doi: https://doi.org/10.1155/2018/8959635.

[12] X. Ding et al., "Privacy-preserving task allocation for edge computing-based mobile crowdsensing," Comput. Electr. Eng., vol. 97, p. 107528, 2022, doi: https://doi.org/10.1016/j.compeleceng.2021.107528.

[13] T. Hukkelaas, R. Mester, and F. Lindseth, "DeepPrivacy: A Generative Adversarial Network for Face Anonymization," in Lecture Notes in Computer Science (LNCS), ECCV Workshops, 2020, pp. 565–580.

[14] Y. Wang, M. Yue, Z. Yao, Z. Chen, R. Hu, and B. Jin, "MAP-GAN: multi-attribute facial privacy protection model without losing identification," Cybersecurity, vol. 8, no. 1, p. 112, 2025, doi: 10.1186/s42400-025-00437-7.

[15] F. Hellmann et al., "GANonymization: A GAN-Based Face Anonymization Framework for Preserving Emotional Expressions," ACM Trans. Multimed. Comput. Commun. Appl., vol. 21, no. 1, Dec. 2024, doi: 10.1145/3641107.

[16] K. Saleh, S. Szénási, and Z. Vámossy, "Generative Adversarial Network for Overcoming Occlusion in Images: A Survey," 2023. doi: 10.3390/a16030175.

[17] H. Yang et al., "G2Face: High-Fidelity Reversible Face Anonymization via Generative and Geometric Priors," IEEE Trans. Inf. Forensics Secur., vol. 19, pp. 8773–8785, 2024, doi: 10.1109/TIFS.2024.3449104.

[18] H. Tian, T. Zhu, and W. Zhou, "Fairness and privacy preservation for facial images: GAN-based methods," Comput. Secur., vol. 122, p. 102902, 2022, doi: https://doi.org/10.1016/j.cose.2022.102902.

[19] Y. Yang, K. Mu, and R. H. Deng, "Lightweight Privacy-Preserving GAN Framework for Model Training and Image Synthesis," IEEE Trans. Inf. Forensics Secur., vol. 17, pp. 1083–1098, 2022, doi: 10.1109/TIFS.2022.3156818.

[20] I. Fontana, M. Langheinrich, and M. Gjoreski, "GANs for Privacy-Aware Mobility Modeling," IEEE Access, vol. 11, pp. 29250–29262, 2023, doi: 10.1109/ACCESS.2023.3260981.

[21] D. Tang, S. Zhou, H. Jiang, H. Chen, and Y. Liu, "Gender-Adversarial Networks for Face Privacy Preserving," IEEE Internet Things J., vol. 9, no. 18, pp. 17568–17576, 2022, doi: 10.1109/JIOT.2022.3155878.

[22] R. More, A. Maity, G. Kambli, and S. Ambadekar, "Privacy-Preserving Video Analytics Through GAN-Based Face De-Identification," in 2024 Second International Conference on Networks, Multimedia and Information Technology (NMITCON), 2024, pp. 1–6. doi: 10.1109/NMITCON62075.2024.10698920.

[23] Z. Chen, T. Zhu, P. Xiong, C. Wang, and W. Ren, "Privacy preservation for image data: A GAN-based method," Int. J. Intell. Syst., vol. 36, no. 4, pp. 1668–1685, Apr. 2021, doi: https://doi.org/10.1002/int.22356.

[24] A. G. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv Prepr. arXiv1704.04861, 2017, [Online]. Available: https://arxiv.org/abs/1704.04861

[25] X. Zhang, X. Zhou, M. Lin, and J. Sun, "ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018, pp. 6848–6856.

[26] H. Shen, G. Bai, Y. Hu, and T. Wang, "P2TA: Privacy-preserving task allocation for edge computing enhanced mobile crowdsensing," J. Syst. Archit., vol. 97, pp. 130–141, 2019, doi: https://doi.org/10.1016/j.sysarc.2019.01.005.

[27] L. Ma, X. Liu, Q. Pei, and Y. Xiang, "Privacy-Preserving Reputation Management for Edge Computing Enhanced Mobile Crowdsensing," IEEE Trans. Serv. Comput., vol. 12, no. 5, pp. 786–799, 2019, doi: 10.1109/TSC.2018.2825986.

[28] X. Wang, S. Garg, H. Lin, G. Kaddoum, J. Hu, and M. S. Hossain, "PPCS: An Intelligent Privacy-Preserving Mobile-Edge Crowdsensing Strategy for Industrial IoT," IEEE Internet Things J., vol. 8, no. 13, pp. 10288–10298, 2021, doi: 10.1109/JIOT.2020.3032797.

[29] T. Saleem and V. Sivakumar, "Mobile Deep Learning: Delving into the Future of Portable AI," in 2025 International Conference on Metaverse and Current Trends in Computing (ICMCTC), 2025, pp. 1–14. doi: 10.1109/ICMCTC62214.2025.11196529.

[30] D. Ngo, H.-C. Park, and B. Kang, "Edge Intelligence: A Review of Deep Neural Network Inference in Resource-Limited Environments," 2025. doi: 10.3390/electronics14122495.