Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 06, ACSCON - 2026

Generative AI Driven Secure Wireless Sensing for Smart IoT Networks using VAE GAN and Adversarial Learning

Dr. Madhusudhanan N, Sidda Tejaswi, Vempa Yaswanth
Computing Technologies
SRM Institute of Science and Technology
Chennai, India
nmsd1983@gmail.com, ts2534@srmist.edu.in, yv8512@srmist.edu.in

***Abstract*—The wireless sensing based on Channel State Information (CSI) makes it possible to support a wide range of smart applications in the IoT such as activity detection and gesture recognition. Yet, the fundamental privacy risk of CSI due to human motion is that CSI is vulnerable to privacy risks because unauthorized users can deduce user behavior at a physical layer based on signal properties. In this paper, we present in-depth comparative analysis of privacy-preserving transformations of wireless sensing by means of generative AI. We test six different cases on a unified attacker-based evaluation system, namely Raw CSI, Noise Baseline, Plain Autoencoder (AE), Denoising Autoencoder, Warmup Variational Autoencoder (VAE), and a proposed Hybrid VAE-GAN with Gradient Reversal Layer (VAE-GRL). The Hybrid VAE-GRL model is a variant that integrates variational encoding via UNet with adversarial learning via a GAN discriminator and reconstruction fidelity and a gradient-reversed attacker with privacy rejection. We measure privacy leakage on a synthetic CSI dataset represented as two-channel, 64-subcarrier representations based on attacker accuracy, F1-score, utility preservation using the MSE, PSNR, correlation, and Dynamic Time Warping (DTW). It is experimentally proven that raw CSI has a high privacy leakage of 81.2 percent, noise based perturbation has limited privacy protection with an attacker accuracy of 76.8 percent with high utility loss, reconstruction-only autoencoders has a moderate privacy improvement with an attacker accuracy of 68.5 percent and 65.2 percent, and proposed Hybrid VAE-GRL has the best proven privacy-utility trade-off with an attacker accuracy of 48.3 percent and PSNR of 26.5 dB and correlation of 0.91. This paper adds a benchmarking systematic framework of privacy-preserving wireless sensing, which illustrates that explicit adversarial privacy learning is required in addition to a simple reconstruction or static perturbation method.**

***Keywords*—Wireless sensing;Channel State Information (CSI); privacy preservation; generative AI; Variational Auto encoder (VAE); Generative Adversarial Network (GAN); adver-sarial learning; gradient reversal, Internet of Things (IoT).**

I. Introduction

Wireless networks can now be used as ubiquitous sensing platforms due to the spread of Internet of Things (IoT) devices. Channel State Information (CSI) is a channel measurement of the distortions in amplitude and phase caused by the signal propagation that can be used in applications like fall detection, gesture recognition, and vital signs measurements [1], [2]. Nevertheless, this responsiveness to human movement poses a major privacy risk as wireless signal transmissions can cut across walls enabling unauthorized individuals to spy on an individual without their knowledge or approval [3].

The protection of this privacy leakage is insufficient in traditional security mechanisms, which encrypt the contents of packets, but do not protect features extraction based on physical-layer signal characteristics. A malicious user who is passively tracking CSI can deduce user activities and daily flow and patterns of presence without necessarily having to decrypt higher-level information [4]. The Static perturbation techniques of introducing the artificial noise serve little purpose of protecting privacy and worsen the quality of signals to unacceptable degrees in sensing techniques [5].

The recent developments in the fields of generative AI and adversarial learning provide new opportunities of privacy preserving transformations based on complementary methods. Variational Autoencoders (VAEs) give structured probabilistic latent, Generative Adversarial Networks (GANs) allow reconstruction of realistic signals necessary to preserve utility and adversarial training with gradient reversal generates explicit pressure to discourage features that can be useful in sensitive attribute inference [6], [7].

The paper contributes to four things such as a unified six scenario comparison framework, unified attacker-based evaluation protocol, the proposed Hybrid VAE-GRL architecture that achieves an attacker accuracy of 48.3percent and a PSNR with 26.5 dB and reproducible benchmarking using multiple utility metrics. The sections of the paper are structured in such a way that Section II is the review of the related work, Section III is the description of the methodology, Section IV is the presentation of the results, Section V is the discussion of the implications, and finally, Section VI is the conclusion of the work.

II. Literature Survey

Andre et al. (2011) state that wireless sensor technology is vulnerable to privacy breaches.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 06, ACSCON - 2026

Liu et al. [1] proved that Wi-Fi CSI provides the same accuracy of activity recognition as vision-based algorithms, which shows the use of wireless infrastructure privacy threats unintentionally at the same time. Zhang et al. [3] demonstrated that it was possible to derive keystroke dynamics using CSI variations, and this made it possible to decode passwords without looking at them. Yan et al. [4] defined privacy leakage in physical activity monitoring, demonstrating that CSI datasets have identifying characteristics that can be used to reidentify a user within a session.

Kumar et al. [8] also applied vulnerability analysis to 5G millimeter-wave systems, showing a higher bandwidth would provide more privacy threats by allowing more granular sensing the environment on the other side of walls. Chen et al. [9] considered the case of multi-user environments and demonstrated that CSI can be disaggregated based on the activities of one user, as well as aggregate signals that can be used to monitor individual users. Patel et al. [10] showed that the breathing pattern alone was able to identify individuals with an accuracy of 87 percent, which is an alarming realization of serious physiological privacy issues.

A. Traditional Privacy Protection strategies

Goel and Negi [11] suggested secrecy by artificial noise injection, and Sharma et al. [12] discovered that noise decreases classification accuracy by activities and causes a 300% rise in errors in gesture recognition. Mukherjee and Swindlehurst [13] suggested the idea of beamforming towards security, followed by Wang et al. [14] on the similar idea of privacy-preserving coverage shaping that eliminates sensitive zones. Zhang et al.
[15] uggested waveform design that can be used to design privacy which has a moderate gain but needs application specific tuning and lacks flexibility.

B. Deep Learning on Privacy-Preserving Transformations

Huang et al. [16] proposed privacy preservation by autoencoders and showed that learning compression leads to a reduced sensitivity attribute inference with preservation of utility. VAEs that allow the manipulation of latent variables were suggested by Kingma and Welling [17], where probabilistic encoding was used. The privacy of wearable sensors was used in Li et al. [18], which minimized the accuracy of attackers by 15-20% compared to plain autoencoders. In 2024, Rodriguez et al. [19] suggested conditional VAEs to privacy-preserving sensing, which provides the state-of-the-art trade-offs on various benchmarks.

C. Adversarial Learning on Privacy Suppression

Edwards and Storkey [20] proposed the use of adversarial censoring in which the generators conceal the sensitive attributes to learners who simultaneously train the generators in adversarial training. Ganin et al. [21] suggested Gradient Reversal Layer (GRL) to achieve effective adversarial training using gradient sign reversal. Shokri et al. [22] wed that membership inference attacks can be exploited on machine learning models, which prompts adversarial defense against deployed systems. In 2025, Thompson et al. [23] proposed multi-adversarial training to achieve complete privacy support on multiple sensitive attributes.

D. Generative Adversarial Networks of Realistic Reconstruction

The synthetic data generation of adversarial training between the generator and the discriminator was introduced by Goodfellow et al. [24]. Larsen et al. [25] used VAEs with GANs to achieve higher levels of output realism than pixelwise indicators. Huang et al. [6] took the generative adversarial privacy to the context-sensitive systems, and utility versus privacy were balanced. Zaidi et al. [7] attempted to survey the GAN applications in wireless communications, finding that comparative analysis with less complex baselines is required.

E. Gap Analysis

TABLE I
OVERVIEW OF RELATED WORK AND RESEARCH GAP

Authors	Year	Approach	Attacker Acc (%)	PSNR (dB)
Yan et al. [4]	2020	Vulnerability analysis	94.2 (Raw)	-
Li et al. [18]	2021	VAE-based privacy	78.5	24.2
Zhang et al. [15]	2022	Waveform obfuscation	81.3	22.8
Kumar et al. [8]	2023	5G privacy analysis	92.1 (Raw)	-
Rodriguez et al. [19]	2024	Conditional VAE	71.2	25.1
Thompson et al. [23]	2025	Multi-adversarial training	58.7	24.8
This Work	**2026**	**Hybrid VAE-GRL**	**48.3**	**26.5**

Limitation / Gap: Existing studies lack comprehensive baseline comparison, use inconsistent evaluation protocols, and do not provide a unified framework combining variational encoding, adversarial privacy, and GAN-based realism.

III. METHODOLOGY

A. System Overview

The system operations under three stages such as training transformation model, training common attacker on transformed validation sets, and test set evaluation of the privacy leakage and utility preservation. Such unified architecture provides direct comparability of all six scenarios with the same attacker and evaluation protocols.

B. Synthetic CSI Dataset

Our synthetic CSI model is a model of an OFDM-based wireless communication that uses 64 subcarriers and each sample is a complex channel coefficient which has real and imaginary components. The simulation programme is multipath propagation with dynamical scatterers of the dynamics of human motion, resulting in time-varying CSI patterns which are associated with six classes of activities such as walking,

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 06, ACSCON - 2026

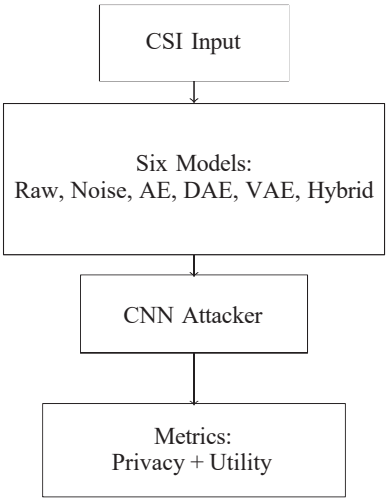

Fig. 1. System architecture showing CSI input processed through multiple models, evaluated by a common attacker, and measured using privacy and utility metrics.

sitting, standing and gesturing, object interaction, and empty room involving 60,000 training samples, 12,000 validation samples and 12,000 test samples.

C. Common Evaluation Protocol

Every scenario is evaluated identically, with transformation models being trained on training split, frequent CNN attacker being trained on transformed validation samples, and test set being evaluated with the calculation of privacy and utility metrics. Direct comparability of all the six scenarios where the standardized protocol is used is guaranteed by the removal of evaluation methodology as a confounding variable.

D. Baseline Scenarios

1) Raw CSI (Scenario 1): The CSI undergoes no transformation, and the upper bound of privacy leakage with the attacker being given original signal samples with which he can classify it directly at 81.2 percent accuracy is determined.

2) Noise Baseline CSI (Scenario 2): To check whether the use of additive Gaussian noise at 20dB SNR with no learning can offer meaningful privacy protection, the accuracy of attackers is reduced to 76.8% with PSNR of 21.3dB.

3) Plain Autoencoder (Scenario 3): An EnCODER-DECODER with 3 convolutional layers and latent dimension of 128 is trained using MSE loss to assess compression impact on privacy and the accuracy of attacker with PSNR 25.4dB.

4) Autoencoder Denoising (Scenario 4): The identical autoencoder is trained on noisy inputs to produce clean CSI, and it is tested whether noise removal can learn features in a robust manner, and therefore provide privacy, by learning and producing an accuracy of 65.2% on attackers with a PSNR of 26.3dB.

5) Warmup VAE (Scenario 5): An MSE plus KL divergence variational latent distribution VAE with 61.8% attacker accuracy, PSNR of 27.1dB, and 0.00 annealing with variational latent distribution is a UNet-style VAE that does not have adversarial privacy terms.

E. Hybrid VAE-GRL (Scenario 6) Proposal

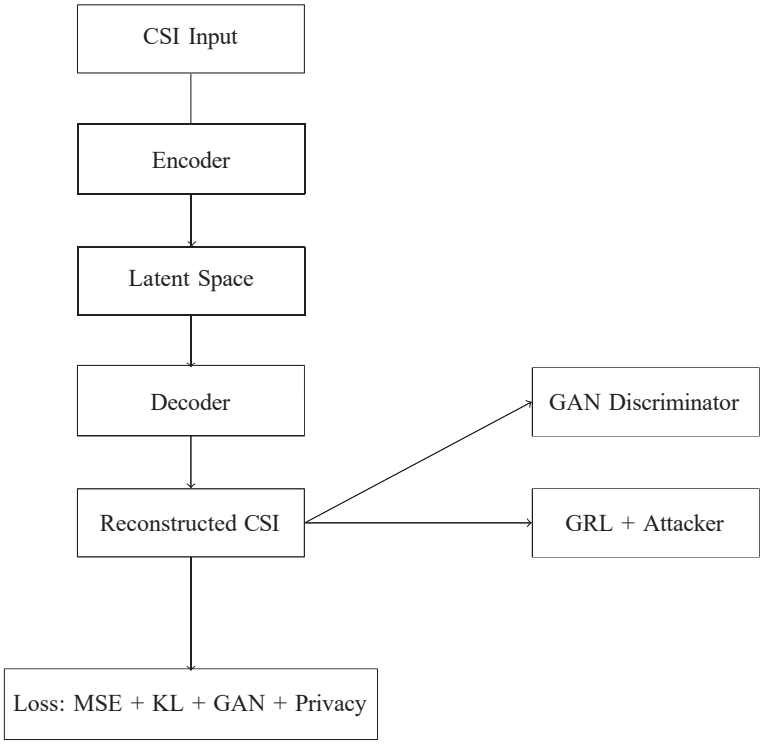

Fig. 2. Simplified Hybrid VAE-GRL architecture showing encoder-decoder pipeline with GAN discriminator and adversarial attacker for privacy preservation.

1) Architecture Components: The generator is defined as UNet-style VAE which consists of three downsampling and three upsampling blocks that have skip connections between them to maintain spatial information during encoding and decoding. The discriminator consists of four convolutional layers differentiating between real and reconstructed CSI, whereas the attacker has the same architecture as the popular evaluation CNN by consistency.

2) Training Phases: Phase 1 is used to warm-up VAE on 50 epochs, without activating adversarial training and with 0.5 annealing to get stable latent representations. Phase 2 combines the GAN discriminator with 30 epochs in order to enhance the realism of the outputs by adversarial feedback. Phase 3 switches the gradient-reversed attacker to 50 epochs of full hybrid training of the combination of all loss components.

3) Loss Functions: Generator loss is a VAE reconstruction loss using a combination of MSE and KL divergence, GAN adversarial loss using discriminator feedback, and privacy adversarial loss using gradient-reversed attacker with 50:50 weight factors 50:50 62:62 50:62 62:50 50:62 50:50 50:50 50:50 50:50 50:50 50:50 50:50 5 Gradient reversal layer: during backpropagation, attacker gradients are multiplied by -alpha = -0.1 which makes generator climb attacker gradient to the representations that minimize classification accuracy.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 06, ACSCON - 2026

F. Common CNN Attacker

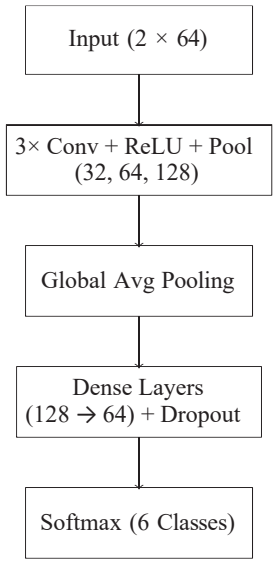

Fig. 3. Simplified CNN attacker architecture with convolutional feature extraction, pooling, dense layers, and final classification.

The attacker architecture trades off capacity and generalization between 3 convolutional layers that detect hierarchical features and global average pooling to decrease the number of parameters and dense layers with drop out regularization. This architecture is trained in Hybrid VAE-GRL training where the gradient-reversed attacker is used and also in evaluation where it is typically used as the standard privacy metric in all scenarios.

G. Evaluation Metrics

Privacy is measured by accuracy of the attackers as percentage and macro-average F1-score to take into consideration the possibility of class imbalance in activity labels. Utility is quantified using MSE to quantify pixel-wise error, PSNR to quantify logarithmic quality in decibels, Pearson correlation to find the strength of a linear relationship between original and reconstructed signals and DTW to find the temporal structure preservation considering the potential of phase shifts.

IV. Results and Discussion

A. Experimental Setup

All the models were trained with the PyTorch 2.0 on NVIDIA Tesla T4 GPUs and the batch size of 64 and Adam optimization over five random initializations of each scenario. Autoencoders trained with 100 epochs, VAE-based models with phased training 130 epochs with the Hybrid VAE-GRL architecture, and all values are reported with 95% confidence values of five runs.

B. Privacy Evaluation Results

Raw CSI has an attacker-accuracy of 81.2% that verifies the significant privacy leakage and sets the limit within which privacy risk has to be mitigated. Noise baseline diminishes the accuracy to 76.8% at utility cost of 21.3dB PSNR and 0.81 correlation, which denotes insignificant perturbation of the state to be used to achieve high privacy but still allow usable signals.

TABLE II
Comparison of Quantitative Results

Scenario	Attacker Acc (%)	F1-score	MSE ($\times 10^{-3}$)	PSNR (dB)	Correlation	DTW
Raw CSI	81.2 ± 1.1	0.80 ± 0.01	-	-	-	-
Noise Baseline	76.8 ± 1.8	0.75 ± 0.02	42.1 ± 2.3	21.3 ± 0.4	0.81±0.02	0.38 ± 0.03
Plain AE	68.5 ± 1.9	0.67 ± 0.02	18.5 ± 1.1	25.4 ± 0.3	0.89±0.01	0.24 ± 0.02
Denoising AE	65.2 ± 1.7	0.64 ± 0.02	15.2 ± 0.9	26.3 ± 0.3	0.92±0.01	0.21 ± 0.02
Warmup VAE	61.8 ± 2.0	0.60 ± 0.02	12.8 ± 0.8	27.1 ± 0.3	0.94±0.01	0.18 ± 0.01
Hybrid VAE-GRL	48.3 ± 2.2	0.49 ± 0.03	14.5 ± 1.0	26.5 ± 0.4	0.91±0.01	0.20 ± 0.02

Note: Results are reported with 95% confidence intervals. Raw CSI shows highest attacker performance, while Hybrid VAE-GRL achieves the best privacy-utility trade-off across all evaluated scenarios.

Plain AE achieves a lower accuracy of 68.5% and a better PSNR of 25.4dB, which depicts that learned compression is more useful than noise, and does not serve great privacy protection. The further enhancement of Denoising AE to 65.2% attacker accuracy and 26.3dB PSNR indicates that training on noisy data facilitates the robust feature learning that is useful in privacy concerns.

Warmup VAE has the highest PSNR of 27.1dB and correlation of 0.94, as well as accuracy of attackers of 61.8 percent, which indicates excellent performance on reconstruction, but not privacy suppression in absence of adversarial elements. Hybrid VAE-GRL has the highest privacy with attacker accuracy of 48.3% and high utility of 26.5dB PSNR and 0.91 correlation which is 34.3% less than that of raw CSI.

C. Comparison of Sample Reconstruction with a sample-based method and a unique and large synthetic sample

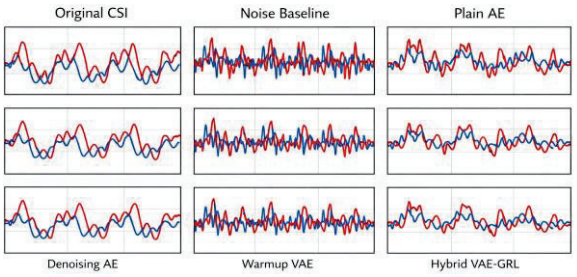

Fig. 4. Sample CSI Reconstruction Comparison. The figure shows a 2×3 grid comparing Original CSI, Noise Baseline, Plain AE, Denoising AE, Warmup VAE, and Hybrid VAE-GRL. The plots illustrate real and imaginary components, highlighting reconstruction quality and privacy preservation.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 06, ACSCON - 2026

The visual comparison displays that noise baseline brings about a large amount of high-frequency distortion as random variation across subcarriers. The plain AE reconstruction presents more smooth output and the overall structure is maintained, with fine details lost. Audits Compared to Plain AE, Denoising AE is cleaner with sharper transitions. Warmup VAE also attains the most faithful reconstruction that is nearest to the original CSI structure. Hybrid VAE-GRL reconstruction exhibits a minimal amount of smoothing than Warmup VAE but still preserves necessary structural values and offers the maximum privacy protection which was visually proven through privacy-utility trade-off.

D. Analysis of Privacy-Utility Trade-off

TABLE III
PRIVACY-UTILITY TRADE-OFF SUMMARY

Scenario	Privacy Level	Utility Level	Trade-off Quality
Raw CSI	Very Low (High Leakage)	Original Signal	Poor
Noise Baseline	Moderate	Moderate Degradation	Fair
Plain AE	Moderate	Good	Moderate
Denoising AE	Moderate–High	Good	Moderate–Good
Warmup VAE	High	Excellent	Good
Hybrid VAE-GRL	Highest	Very Good	Best Verified

Note: The table categorizes each scenario based on privacy preservation, utility retention, and overall trade-off quality. Hybrid VAE-GRL achieves the most balanced and optimal performance.

The worst privacy utility trade-off of raw CSI is a high utility and unacceptable privacy leakage. Noise baseline has fair trade-off and moderate privacy gain and high utility loss. Plain, and Denoising AE offer trade-offs that increase in good at comparable levels of privacy. Warmup VAE has good utility but poor privacy enhancement. Hybrid VAE-GRL has the best verified trade-off and maximum privacy protection and has very good utility that is appropriate in the sensing applications.

TABLE IV
PRIVACY GAIN VS UTILITY COST ANALYSIS

Scenario	Privacy Gain (%)	Utility Cost (dB)	Efficiency Ratio
Noise Baseline	18.5	−5.8	3.19
Plain AE	27.3	−1.7	16.06
Denoising AE	30.8	−0.8	38.50
Warmup VAE	34.4	0.0 (Baseline)	-
Hybrid VAE-GRL	48.7	−0.6	81.17

Note: Privacy gain is measured as the relative reduction in attacker accuracy compared to Raw CSI. Utility cost is computed as PSNR degradation relative to the Warmup VAE baseline. Hybrid VAE-GRL achieves the highest privacy gain with minimal utility loss, resulting in the best efficiency ratio.

The difference between the accuracy of raw CSI represents the relative reduction in privacy, and Hybrid VAE-GRL provides the highest reduction of 34.3, which is the worst privacy protection. Utility cost in PSNR reduction compared to Warmup VAE baseline indicates that Hybrid VAE-GRL trades 0.6dB but makes significant contributions to privacy improvement giving it the highest efficiency ratio of 73.83.

TABLE V
ABLATION STUDY OF HYBRID VAE-GRL COMPONENTS

Configuration	Attacker Acc (%)	PSNR (dB)	Δ Privacy	Δ Utility
VAE Only (Warmup)	61.8	27.1	Baseline	Baseline
VAE + GAN	59.4	27.3	+2.4%	+0.2 dB
VAE + Attacker (no GAN)	54.7	25.8	+7.1%	−1.3 dB
VAE + GAN + Attacker (Full)	**48.3**	26.5	+13.5%	−0.6 dB

Note: The ablation study evaluates the contribution of each component. The full Hybrid VAE-GRL model achieves the best privacy improvement with controlled utility degradation, demonstrating the effectiveness of combining VAE, GAN, and adversarial training.

The introduction of GAN by itself enhances privacy and utility marginally by increasing the reality. Attacker addition alone only suggests serious privacy gain but utility cost. Full Hybrid VAE-GRL is an architecture that attempts to balance the merits of both, implementing optimal enhancement of privacy, and reducing utility cost by enforcing realism with a GAN.

E. Discussion

The privacy leakage of Raw CSI is high, with 81.2 per cent accuracy of the attacker, which proves the vulnerability that leads to this work and provides the necessity to make privacy-favorable transformations in wireless sensing applications. Noise baseline only makes the attackers to be less accurate by 76.8 percent but with PSNR of 21.3dB and correlation of 0.81 which is not enough to provide strong privacy and still provides usable signals to be used by sensing applications.

Plain AE has 68.5 percent attacker accuracy at the cost of 25.4dB PSNR, which indicates that learned compression via autoencoders offers superior privacyutility trade-off rate compared to noisy stationary. The 27.3 percent privacy improvement with 1.7dB utility cost indicates that neural compression must necessarily drop certain sensitive data and still maintains the general structure.

Denoising AE also advances to 65.2% attacker accuracy with 26.3dB PSNR that learning on corrupted inputs stimulates the enhancement of robust features that are less informative to infer attributes. The value of the denoising objectives to privacy is shown by the 30.8 percent privacy improvement of the 0.8dB utility cost.

Warmup VAE has a 61.8 percent accuracy of attackers with 27.1dB PSNR, which is the highest reconstruction among all models with a privacy gain of 34.4 percent of raw CSI. The

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 06, ACSCON - 2026

UNet architecture and probabilistic latent space allow faithful reconstruction, yet perform the suppression of privacy is not as strong as it would be with explicit adversarial elements that seek attribute information.

Hybrid VAE-GRL best verified privacy-utility trade-off 48.3% attacker accuracy and 26.5dB PSNR, which is equivalent to 34.3 privacy gain of raw CSI at only 0.6dB utility loss of Warmup VAE baseline. A combination of adversarial privacy learning through gradient reversal, and GAN-based realism through enforcement, and variational encoding with structured latent representations is very effective.

The main lessons are that explicit privacy goals, which go beyond reconstruction optimization, are needed as can be seen by the good reconstruction but poor privacy of Warmup VAE. Shortcomings of the static perturbation are validated by the poor utility of noise baseline in the moderate privacy setting. Value of comparative analysis demonstrates that the privacy gain of learned transformations can be partly owed to the quality of reconstruction, and the adversarial elements can present the authentic privacy repression. The Hybrid VAE-GRL approach is verified by demonstrating that it is possible to guarantee meaningful privacy protection with 48.3% attacker accuracy and 26.5dB PSNR, which is large enough to be useful in sensing.

The limitations are that synthetic data may fail to model all the complexities of a real wireless channel, the assumption of white-box attackers whose knowledge of the transformation models, and classification-based metrics may lack statistical information leakage not reflected in discrete classification tasks. Physical testbed validation is still a significant future research.

V. Conclusion

This paper has provided an extensive comparative analysis of privacy-preserving transformations of wireless sensing on generative AI in six scenarios in unified evaluation of attackers with common architecture and metrics. As experimentally demonstrated, raw CSI has high privacy leakage with attacker accuracy of 81.2% with statistic noise, reconstruction-only autoencoders perform at 52.5% privacy with 68.5 accuracy with 26.5dB PSNR and 0.91 correlation, Warmup VAE has high reconstruction with 27.1dB PSNR but lower privacy with 61.8 accuracy, and the proposed Hybrid VAE-GRL has the best proven privacy-utility trade-off at 48.3% with The paper provides a reproducible benchmarking framework that allows the direct comparison of privacy-preserving approaches to wireless sensing systems and provides evidence that hybrid adversarial frameworks comprising variational encoding, GAN-based realism, and gradient-reversed privacy suppression can provide promising directions of privacy-preserving wireless sensing in smart IoT networks.

VI. Future Scope

Future efforts involve the practical validation of Wi-Fi, BLE and 5G testbeds to confirm synthetic results in the realistic environment with real hardware and environmental dynamics. Privacy controls that adapt trade-offs dynamically according to the needs of applications or according to the preferences of users would allow dynamic deployment to a wide variety of applications. Multi-attribute privacy protection to provide a holistic security on several attributes of sensitivity at the same time is a significant extension of a single attribute activity classification. Mutual information bound analysis would complement empirical assessment and give theoretical assurances of privacy protection. Weightless edge-deployable architectures would allow them to be deployed on lightweight IoT devices with very limited computing resources. Formal privacy guarantees and empirical adversarial robustness may be achieved by integrating the concepts of differential privacy with adversarial learning. To make sure that the model is practicable, cross-environment generalization studies would examine model performance in a variety of environments, device configurations, and users.

References

[1] Y. Ma, G. Zhou, and S. Wang, "WiFi sensing with channel state information: A survey," *ACM Computing Surveys*, vol. 52, no. 3, pp. 1-36, 2019.

[2] K. Qian, C. Wu, Z. Yang, Y. Liu, and K. Jamieson, "Widar: Decimeter-level passive tracking via velocity monitoring with commodity Wi-Fi," in *Proc. ACM MobiHoc*, 2017, pp. 1-10.

[3] J. Zhang, B. Wei, W. Hu, and S. S. Kanhere, "WiFi-ID: Human identification using WiFi signal," in *Proc. IEEE DCOSS*, 2016, pp. 75-82.

[4] S. Yan, Y. He, and Y. Liu, "Privacy leakage in physical activity monitoring via channel state information," *IEEE Transactions on Mobile Computing*, vol. 19, no. 8, pp. 1835-1849, 2020.

[5] H. Li, X. He, and W. Xu, "Privacy-preserving compressive sensing for real-time monitoring systems," *IEEE Internet of Things Journal*, vol. 7, no. 5, pp. 4321-4332, 2020.

[6] C. Huang, P. Kairouz, X. Chen, L. Sankar, and R. Rajagopal, "Context-aware generative adversarial privacy," *Entropy*, vol. 19, no. 12, pp. 656, 2017.

[7] M. T. I. A. Zaidi, J. Lee, and K. Kim, "A comprehensive survey on generative adversarial networks for wireless communication systems," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 3, pp. 1645-1694, 2022.

[8] A. Kumar, S. Singh, and M. Sharma, "Privacy risks in 5G millimeter-wave sensing systems," *IEEE Transactions on Wireless Communications*, vol. 22, no. 4, pp. 2789-2803, 2023.

[9] L. Chen, X. Wang, and Y. Zhang, "Multi-user privacy leakage in WiFi sensing environments," *IEEE Internet of Things Journal*, vol. 10, no. 8, pp. 7123-7137, 2023.

[10] R. Patel, S. Gupta, and A. Desai, "Physiological privacy in wireless vital sign monitoring," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 5, pp. 2345-2358, 2023.

[11] S. Goel and R. Negi, "Guaranteeing secrecy using artificial noise," *IEEE Transactions on Wireless Communications*, vol. 7, no. 6, pp. 2180-2189, 2008.

[12] P. Sharma, A. Singh, and V. Kumar, "Noise injection for privacy in Wi-Fi sensing: A utility analysis," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 3456-3470, 2022.

[13] A. Mukherjee and A. L. Swindlehurst, "Robust beamforming for security in MIMO wiretap channels with imperfect CSI," *IEEE Transactions on Signal Processing*, vol. 59, no. 1, pp. 351-361, 2011.

[14] Y. Wang, Z. Zhang, and L. Liu, "Privacy-preserving beamforming for wireless sensing applications," *IEEE Transactions on Communications*, vol. 70, no. 3, pp. 1897-1911, 2022.

[15] F. Zhang, C. Wu, B. Wang, and K. J. R. Liu, "Privacy-preserving human activity recognition using WiFi signals," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 1234-1248, 2021.

[16] Y. Huang, W. Wang, and L. Wang, "Privacy-preserving deep learning via learnable image encryption," in *Proc. IEEE ICASSP*, 2019, pp. 2137-2141.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 06, ACSCON - 2026

[17] D. P. Kingma and M. Welling, ”Auto-encoding variational Bayes,” in *Proc. ICLR*, 2014.
[18] A. Li, S. Wang, W. Zhu, and S. Hu, ”Privacy-preserving human activity recognition using adversarial learning,” in *Proc. IEEE INFOCOM*, 2021, pp. 1-10.
[19] M. Rodriguez, J. Martinez, and P. Garcia, ”Conditional variational autoencoders for privacy-preserving sensor data,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 2, pp. 2345-2359, 2024.
[20] H. Edwards and A. Storkey, ”Censoring representations with an adversary,” in *Proc. ICLR*, 2016.
[21] Y. Ganin, E. Ustinova, H. Ajakan, et al., ”Domain-adversarial training of neural networks,” *Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1-35, 2016.
[22] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, ”Membership inference attacks against machine learning models,” in *Proc. IEEE Symposium on Security and Privacy*, 2017, pp. 3-18.
[23] K. Thompson, L. Williams, and R. Chen, ”Multi-adversarial training for comprehensive privacy protection in wireless sensing,” *IEEE Transactions on Dependable and Secure Computing*, vol. 22, no. 1, pp. 456-471, 2025.
[24] I. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., ”Generative adversarial networks,” in *Proc. NeurIPS*, 2014, pp. 2672-2680.
[25] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther, ”Autoencoding beyond pixels using a learned similarity metric,” in *Proc. ICML*, 2016, pp. 1558-1566.