A Unified Privacy-Preserving Framework for Non-Intrusive Load Monitoring

Chenhan Xiao, **Hamad Alduaij**, **Yang Weng**
Department of Electrical, Computer and Energy Engineering
Arizona State University, Tempe, Arizona, United States
Email: {cxiao20, halduaij, yang.weng}@asu.edu

***Abstract*—Non-Intrusive Load Monitoring (NILM) estimates appliance-level energy usage from aggregate measurements, providing critical insights for smart grid energy management and demand response. However, the aggregate load data can also expose sensitive behavioral patterns, such as occupancy, appliance usage, or socio-economic status, raising growing cybersecurity and privacy concerns. This paper presents a unified privacy-preserving framework for NILM that integrates both traditional probabilistic models and modern machine-learning architectures under a common selective protection mechanism. The framework injects calibrated differential-privacy noise into aggregated power signals before third-party transmission, while providing authorized entities with noise metadata (e.g., variance) as lightweight access keys. These metadata allow authorized NILM models to perform noise-aware inference, maintaining disaggregation accuracy, whereas unauthorized models experience intentional degradation. Evaluations on the GREEND dataset demonstrate that the proposed framework effectively balances data utility and privacy, preserving accurate disaggregation for trusted analytics while mitigating behavioral leakage to untrusted observers.**

***Index Terms*—Non-Intrusive Load Monitoring, Differential Privacy, Cybersecurity**

I. Introduction

Non-Intrusive Load Monitoring (NILM) [1] has emerged as a key technique for inferring appliance-level energy consumption from aggregate power measurements, thereby reducing the need for costly per-appliance instrumentation. It supports consumption optimization, strategic load scheduling, and demand-side management for utilities [2]. With high-resolution smart meters and large public datasets [3], [4], NILM has advanced from probabilistic models such as Factorial Hidden Markov Models (FHMMs) [5] to deep architectures like sequence-to-point Convolutional Neural Networks (CNNs) [6], greatly improving accuracy and scalability. Yet these powerful algorithms also enable adversaries to infer sensitive behaviors, such as residential occupancy, appliance usage, or socio-economic status [7], [8], making NILM both a tool for grid intelligence and a potential privacy threat.

Specifically, as utilities increasingly outsource analytics tasks (e.g., load forecasting, anomaly detection, and NILM) to third-party providers with specialized machine-learning (ML) expertise and high-performance computing infrastructure [9], [10], power consumption data become vulnerable during transmission, storage, and processing [11]. These external vendors often operate outside traditional regulatory oversight, creating privacy and accountability gaps. In fact, the unauthorized sharing of raw power usage data has already prompted residents to file lawsuits against utility companies [12]. Recent incidents, including the 2025 Nova Scotia Power breach [13] and the 2023 MOVEit data leak affecting Texas utilities [14], also highlight this risk. Once exposed, aggregated power data can be exploited by advanced NILM algorithms to reconstruct household routines and behaviors [7], [15], underscoring the urgent need for privacy-preserving NILM frameworks compatible with ML-driven grid operations.

To safeguard customer privacy in NILM applications, existing research falls into two main categories. The first targets the training phase, where federated learning frameworks such as Fed-NILM [16] and FederatedNILM [17] keep data local while sharing only model updates. DP^2-NILM [18] further integrates differential privacy into the federated setting to provide formal guarantees against information leakage. While effective in reducing training-time privacy risks, these approaches do not protect data during the inference phase, when trained models process new residential data. The second category addresses this setting by perturbing aggregate signals before transmission to third-party NILM systems. Prior studies have applied differential privacy through calibrated noise injection [8] or adversarial perturbations that mask appliance signatures while preserving aggregate-level statistics [19]. However, these methods primarily seek to degrade NILM accuracy universally, without considering how to selectively preserve disaggregation accuracy for authorized parties while blocking unauthorized inference.

For enabling selective privacy preservation in NILM, we propose a framework that perturbs aggregated power data with calibrated Gaussian noise while sharing compact statistical descriptors, such as the mean and variance of the injected noise, with authorized parties that conduct NILM algorithms. These metadata serve as lightweight access keys that allow trusted entities to incorporate the noise model into their NILM inference process and recover accurate appliance-level consumption. These keys are easier to protect than raw power data and can be periodically refreshed to prevent leakage. Unauthorized parties (e.g., attackers) that lack this information must rely on the obfuscated signal, leading to degraded disaggregation performance, as illustrated in Fig. 1. Unlike conventional obfuscation schemes that uniformly degrade accuracy, the proposed approach achieves asymmetric protection by preserving data utility for authorized users while hindering unauthorized inference, and it can be efficiently implemented at the utility side by perturbing aggregated data before transmission to third-party NILM service providers.

To demonstrate the generality of the proposed framework, we adapt both traditional and modern NILM algorithms into

979-8-3315-8138-1

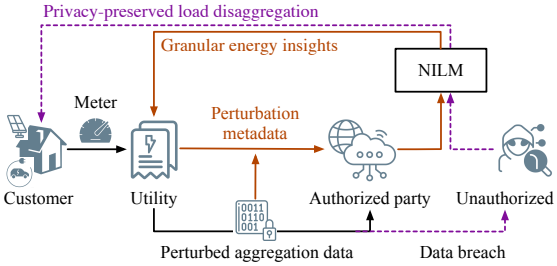

Fig. 1. Overview of the selective privacy-preserving NILM framework.

noise-aware versions that compensate for differentially private perturbations. For probabilistic models such as the FHMM, the emission distribution is modified to include the variance of the injected Gaussian noise. During inference, the emission likelihood is evaluated using the combined intrinsic and noise variance, ensuring statistical consistency between the perturbed aggregate signal and the latent appliance states. This allows authorized entities to recover accurate disaggregation results without retraining. For machine-learning-based NILM, such as CNN models, the noise metadata are incorporated as auxiliary conditioning variables within a noise-aware loss function, enabling the models to learn mappings that account for known noise statistics while naturally degrading performance when metadata are absent. Conceptually, the framework parallels asymmetric cryptographic schemes but offers greater computational efficiency and compatibility with real-time smart grid operations.

We validate the proposed framework on the publicly available GREEND dataset [4], evaluating NILM performance under both authorized and unauthorized conditions. In the authorized setting, FHMM- and ML-based models are adapted to the injected noise statistics, while in the unauthorized setting, standard (vanilla) counterparts are applied directly to the perturbed data. Results show that the framework satisfies formal differential-privacy guarantees, significantly degrading unauthorized inference while maintaining high disaggregation accuracy for authorized users.

The remainder of this paper is organized as follows. Section II formulates the NILM problem and introduces the differential privacy framework used to quantify privacy guarantees. Section III describes the proposed noise calibration procedure and the adaptation of both FHMM- and CNN-based NILM models for noise-aware inference. Section IV presents the experimental results, and Section V concludes the paper.

II. Preliminaries

A. System Modeling of NILM

NILM aims to estimate the power usage of individual appliances from the total power measured by a single sensor over time. Consider M known appliances, each consuming power $p_m[n] \in \mathbb{R}$ for $m = 1, \ldots, M$ at time step n. The aggregated household power is given by $p_{\text{agg}}[n] = \sum_{m=1}^{M} p_m[n] + e[n] = \sum_{m=1}^{M+1} p_m[n]$, where $e[n] := p_{M+1}[n]$ represents measurement noise or the contribution of unmodeled devices (also referred to as ghost power [20]). The goal of NILM is to recover the individual appliance consumptions $p_m[n]$ for $m = 1, \cdots, M$, using only the aggregate signal $p_{\text{agg}}[n]$ [1]. NILM can achieve this disaggregation because individual appliances exhibit characteristic power signatures, such as distinct activation patterns, magnitudes, and temporal dynamics, which can be learned or modeled from data [5], [6]. Although such disaggregation provides valuable insights for energy management, it also introduces privacy risks: fine-grained load patterns can reveal occupant presence, activity schedules, and lifestyle habits. This motivates the need to integrate formal privacy guarantees into NILM algorithms.

B. Differential Privacy as a Privacy Metric

To quantify privacy protection, we adopt the framework of *differential privacy* (DP) [21], which offers worst-case guarantees against information leakage. In this setting, the objective is to prevent adversaries from inferring the contribution of any single appliance to the aggregated power signal. Two aggregate signals $p_{\text{agg}}[n]$ and $p'_{\text{agg}}[n]$ at time n are defined as *neighboring* if they differ in the consumption of exactly one appliance, i.e., there exists $m_0 \in \{1, \ldots, M\}$ such that $p_{m_0}[n] \neq p'_{m_0}[n]$, while all other appliance consumptions remain identical. Let $\mathcal{M} : \mathbb{R} \to \mathbb{R}$ denote a randomized mechanism that perturbs $p_{\text{agg}}[n]$ before data sharing. $\mathcal{M}$ is said to satisfy (ε, δ)-differential privacy if, for any pair of neighboring inputs and any measurable subset $\mathcal{S} \subseteq \mathbb{R}$,

$$\mathbb{P}[\mathcal{M}(p_{\text{agg}}[n]) \in \mathcal{S}] \leq e^{\varepsilon} \mathbb{P}[\mathcal{M}(p'_{\text{agg}}[n]) \in \mathcal{S}] + \delta. \quad (1)$$

This condition ensures that the output distribution changes only slightly when any single appliance's power value is modified, thereby limiting an adversary's ability to infer its presence or magnitude. The parameter $\varepsilon > 0$ is the privacy budget: smaller ε corresponds to stronger privacy, while $\delta > 0$ captures the probability of rare violations.

III. Methodology

A. Differentially Private Data Perturbation at the Utility Level

To mitigate adversarial inference risks, we propose a noise-injection framework that perturbs the aggregate power signal before transmission to third-party NILM services (see Fig. 1). Instead of accessing the raw aggregate data $p_{\text{agg}}[n]$, external NILM algorithms operate on a perturbed version

$$\tilde{p}_{\text{agg}}[n] = \mathcal{M}(p_{\text{agg}}[n]) = p_{\text{agg}}[n] + \eta[n], \quad (2)$$

where $\eta[n] \sim \mathcal{D}$ is a zero-mean random noise. The distribution choice of $\mathcal{D}$ and its parameters determine the privacy-accuracy trade-off. We adopt Gaussian noise here due to its analytical tractability and natural compatibility with FHMM-based NILM models, where appliance emissions are often assumed Gaussian.

Proposition 1. *Let $\eta[n] \sim \mathcal{N}(0, \sigma^2)$ be independent Gaussian noise added to each time step $p_{agg}[n]$ of the aggregate signal. Then the mechanism in (2) satisfies (ε, δ)-differential privacy provided that $\sigma \geq \frac{\Delta \cdot \sqrt{2\log(1.25/\delta)}}{\varepsilon}$ [22], where Δ is an upper bound on the change in $p_{agg}[n]$ resulting from modifying one appliance's power usage [22].*

This result follows directly from the classical Gaussian mechanism in differential privacy [22]. In practice, a conservative choice for Δ is $\Delta = \max_m P_m^{\text{rated}}$, where P_m^{rated} denotes the rated power of appliance m. This ensures strong protection against appliance-level inference while allowing the utility to tune σ according to the desired privacy budget ε. The mechanism is easily implementable by system operators, since it only requires adding calibrated random noise to the aggregate signal before data sharing, without altering existing metering or communication infrastructure.

We further note that since the aggregate signal $\tilde{p}_{\text{agg}}[1:N]$ is released as a sequence of N independently perturbed values, the total privacy loss accumulates across time. Under the advanced composition theorem [22], the full-sequence guarantee satisfies $(\varepsilon_{\text{tot}}, \delta_{\text{tot}})$-differential privacy with: $\varepsilon_{\text{tot}} = \sqrt{2N\ln(1/\delta')}\cdot\varepsilon + N\varepsilon(e^{\varepsilon}-1), \delta_{\text{tot}} = N\delta + \delta'$, for any $\delta' > 0$. This provides a formal quantification of privacy degradation due to repeated releases of perturbed data.

Despite the privacy protection, the injected noise inevitably degrades the accuracy of all NILM algorithms, even for authorized third parties [8]. To mitigate this effect, we propose that the utility shares perturbation metadata, such as the known noise variance σ^2, with authorized entities. This side information acts as a lightweight correction key that enables authorized NILM algorithms to perform noise-aware inference, whereas unauthorized parties who lack access to σ^2 operate on obfuscated signals with degraded performance. As illustrated in Fig. 1, this selective protection framework allows utilities to share aggregate power data in a privacy-preserving way. In the next sections, we detail how the NILM inference process can be adapted to account for the noise variance σ^2 in order to maintain accuracy.

B. Noise-Aware Adaptation of FHMM-Based NILM

The FHMM [5] provides a classical probabilistic formulation of NILM, modeling household power consumption as the superposition of multiple appliance-level processes. Each appliance $m \in \{1, \cdots, M\}$ is represented by a discrete-time hidden Markov chain with latent state $s_m[n] \in \mathcal{S}_m$ that captures operating modes such as *off*, *idle*, and *active*. The aggregate power is expressed as a Gaussian emission centered on the sum of the mean appliance powers:

$$\mathcal{P}(p_{\text{agg}}[n] \mid \mathbf{s}[n]) = \mathcal{N}(p_{\text{agg}}[n] \mid \sum_{m=1}^{M} \mu_{s_m[n]}, \tau^2), \quad (3)$$

where $\mu_{s_m[n]}$ denotes the mean power of appliance m in state $s_m[n]$, τ^2 represents intrinsic modeling variance, and $\mathbf{s}[n] = (s_1[n], \ldots, s_M[n])$ is the joint latent state.

Under the proposed privacy mechanism, the aggregate signal is perturbed according to Eq. (2), introducing independent Gaussian noise with variance σ^2. To maintain statistical consistency between the perturbed observations and the FHMM emission model, the likelihood function in Eq. (3) is modified to include the total observation variance:

$$\mathcal{P}(\tilde{p}_{\text{agg}}[n] \mid \mathbf{s}[n]) = \mathcal{N}(\tilde{p}_{\text{agg}}[n] \mid \sum_{m=1}^{M} \mu_{s_m[n]}, \tau^2 + \sigma^2). \quad (4)$$

This adjustment accounts for the independent privacy noise by inflating the emission variance while keeping the emission mean unchanged. Since the FHMM is trained on clean data, this modification is applied only during inference and requires no retraining. Authorized users, who possess the correct σ^2, evaluate likelihoods using the total variance $\tau^2 + \sigma^2$, maintaining relatively accurate disaggregation. Unauthorized users, unaware of the noise statistics, assume incorrect variances and experience likelihood mismatch and degraded performance.

C. Noise-Aware Adaptation of ML-Based NILM

To illustrate compatibility with modern machine-learning (ML)-based NILM, we extend the same principle to a data-driven model, represented by the sequence-to-point convolutional neural network (CNN) [6]. Given a sliding window of aggregate power readings, the CNN learns a nonlinear mapping to the target appliance's power:

$$\hat{p}_m[n] = f_\theta(\tilde{p}_{\text{agg}}[n-w : n+w]), \quad (5)$$

where f_θ denotes the neural network parameterized by weights θ, and $2w+1$ is the input window size.

For authorized parties, to explicitly incorporate the known perturbation, the noise level σ is provided to the model as an auxiliary conditioning variable through input concatenation: $\big[\tilde{p}_{\text{agg}}[n-w : n+w]; \sigma\big]$. This conditioning allows the CNN to adapt its internal representations to different privacy levels without modifying its architecture, enabling seamless deployment under varying noise magnitudes.

To further capture the uncertainty introduced by the perturbation, the model is optimized using a *Gaussian negative log-likelihood (NLL)* loss:

$$\mathcal{L}_{\text{NLL}} = \frac{1}{N}\sum_{n=1}^{N} \frac{(p_m[n] - \hat{p}_m[n])^2}{2(\hat{\tau}^2[n] + \sigma^2)} + \frac{1}{2}\log(\hat{\tau}^2[n] + \sigma^2), \quad (6)$$

where $\hat{\tau}^2[n]$ is an additional output predicted by a parallel branch of the network that estimates the intrinsic data variance. This formulation jointly models data noise and predictive uncertainty, allowing the CNN to produce calibrated and robust estimates even under privacy-preserving perturbations.

The two case studies demonstrate the versatility of the proposed privacy-preserving framework. Probabilistic NILM models, such as FHMM, incorporate noise statistics using variance correction during inference, while data-driven models, such as CNN, embed them during training through conditioning and uncertainty-aware optimization. This unified design ensures compatibility across NILM methods and supports practical real-world deployment.

The proposed framework also supports customer-specific noise injection to enhance robustness and deployment flexibility. Rather than using a uniform noise level, each household can adopt a personalized variance σ_i^2. This allows utilities to offer user-specific trade-offs between privacy and data fidelity according to customer preferences or contractual requirements.

IV. Numerical Results

To evaluate the proposed privacy-preserving NILM framework, experiments are conducted on the GREEND dataset [4],

a standard benchmark for residential energy disaggregation. GREEND contains active power measurements at both appliance and aggregate levels across multiple European households, originally sampled at 1 Hz. To emulate realistic smart meter reporting and mitigate high-frequency noise, the data are downsampled to 1-minute intervals. A continuous 1-month segment is used for model training, followed by a 10-day testing period during which privacy-preserving perturbations are applied. For the authorized-party scenario, both FHMM- and CNN-based NILM models [5], [6] are evaluated using the noise-aware adaptations in Sections III-B and III-C. For unauthorized scenario, the same models are tested in their standard forms, operating without access to the noise statistics.

A. Noise Calibration for Differential Privacy

Gaussian noise is injected into the aggregate data under three variance settings, $\sigma \in \{500, 1000, 1500\}$ W, based on an estimated single-appliance sensitivity of $\Delta \approx 10^3$ W over the 10-day GREEND test window. With the privacy parameter fixed at $\delta = 10^{-2}$, the corresponding privacy budgets ε are computed using Proposition 1. Table I summarizes the resulting configurations and their qualitative privacy strengths.

TABLE I
NOISE CONFIGURATION AND ACHIEVED PRIVACY GUARANTEES

Noise Std. (σ)	Sensitivity Δ	ε (with $\delta = 10^{-2}$)	Privacy Level
500 W	1000 W	6.21	Weak
1000 W	1000 W	3.11	Moderate
1500 W	1000 W	2.07	Strong

These calibrated noise magnitudes enable systematic evaluation of the privacy-accuracy trade-off for unauthorized parties. As σ^2 increases, the privacy strengthens (lower ε), while disaggregation accuracy deteriorates. Fig 2 illustrates this effect by visualizing the inferred on/off states of a television under three noise levels using a standard FHMM without noise awareness. Each dot marks an inferred "on" moment, with the ground truth shown in the bottom row. At $\sigma = 500$ W, the model still partially recovers activation patterns, whereas at 1000 W and 1500 W, the inferred transitions become increasingly distorted and misaligned with the ground truth. This confirms that privacy-preserving noise significantly hampers the ability of unauthorized parties to reconstruct appliance-level activity.

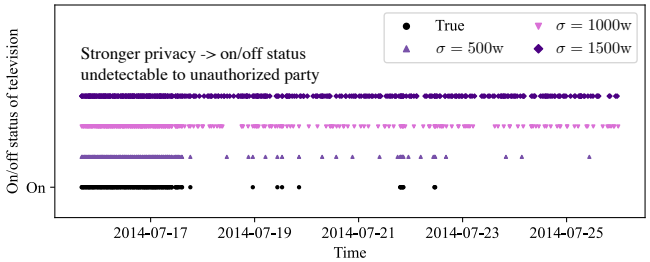

Fig. 2. Television on/off state inferred by a standard FHMM under Gaussian noise with $\sigma \in \{500, 1000, 1500\}$ W.

B. Binary On/Off State Recovery of Appliances

We first evaluate binary disaggregation accuracy at the appliance on/off state level. Fig. 3 shows the inferred on/off state of nine appliances over a 10-day window. Black dots denote ground-truth on events, while red and pink indicate predictions from the noise-adapted (authorized) and standard (unauthorized) FHMM models, respectively. The authorized model accurately recovers appliance usage patterns despite injected noise, performing well on both continuous loads (e.g., fridge freezer) and sporadic devices (e.g., coffee maker, vacuum cleaner). Only minor degradation appears for low-power appliances such as the kettle. In contrast, the unauthorized model produces frequent false positives and misses short or subtle on/off events. These results confirm that access to noise statistics allows authorized parties to maintain NILM accuracy, fulfilling the objective of selective privacy preservation.

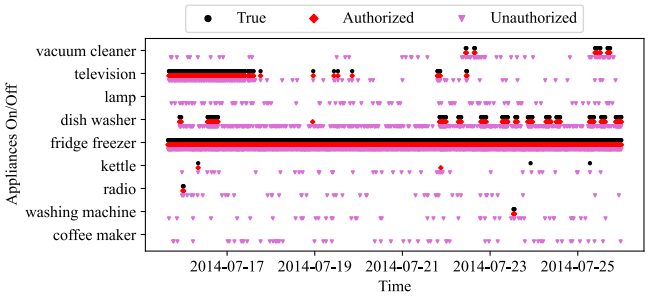

Fig. 3. Inference of on/off state of appliances between noise-adapted (authorized) and standard (unauthorized) FHMM models. $\sigma = 1000$w.

C. Continuous Power Disaggregation of Appliances

Beyond binary state detection, we further assess time-series fidelity of appliance-level power reconstruction. Fig. 4 compares disaggregated power traces for (left) the ground truth, (middle) the authorized noise-adapted CNN, and (right) the standard unauthorized CNN. Each trace represents one of nine appliances over the 10-day evaluation window. The authorized model closely follows true consumption patterns, particularly for dominant appliances such as the fridge freezer, dishwasher, and television. Despite the added perturbation, key temporal patterns and activation cycles remain well preserved. By contrast, the unauthorized model exhibits systematic distortion: short-duration or low-power activities are lost, and baseline consumption is consistently overestimated. It indicates the unauthorized model's inability to compensate for the injected noise, reinforcing the framework's effectiveness in achieving selective privacy protection.

D. Sensitivity Analysis to Perturbation Noise Variance

To quantify the effect of privacy-preserving noise on NILM performance, we conduct a sensitivity analysis by varying the noise variance σ^2 and measuring binary on/off inference accuracy using two standard metrics: F1 score (higher is better) and Hamming loss (lower is better). The reported results represent the average performance of the FHMM and CNN models. As shown in Table II, authorized NILM models maintain high F1 scores (above 0.77) and low Hamming loss, indicating that the noise-aware adaption effectively compensates for injected perturbations. In contrast, unauthorized users experience a marked performance decline as σ increases: the F1 score drops from 0.76 at 200 W to 0.41 at 1500 W, accompanied by a steady rise in Hamming loss. This widening performance

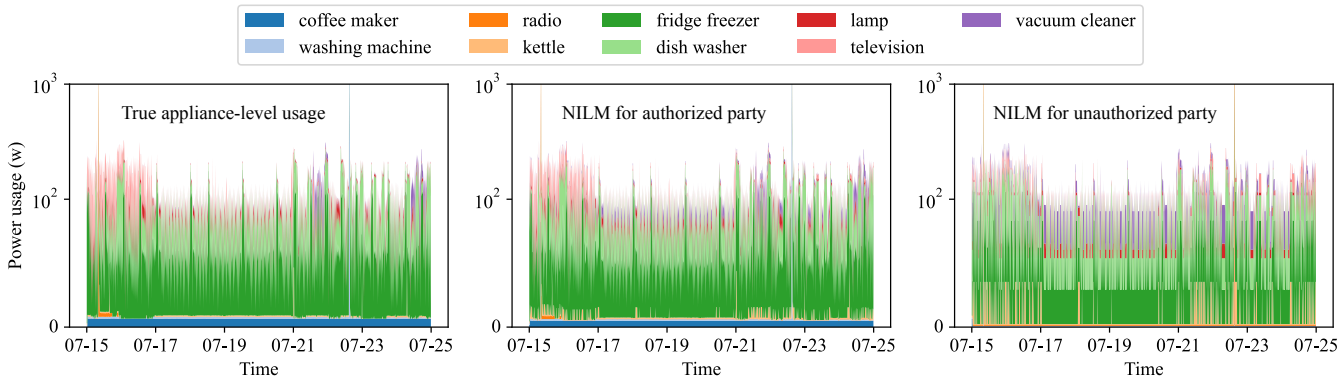

Fig. 4. Disaggregation of appliance-level power for authorized-party and unauthorized party using FHMM-based NILM algorithms.

gap confirms the ability to preserve disaggregation accuracy for trusted entities while progressively degrading unauthorized inference under stronger privacy settings.

TABLE II
DISAGGREGATION ACCURACY UNDER VARIOUS NOISE LEVEL.

σ (W)	Authorized Party		Unauthorized Party	
	F1 Score	Hamming Loss	F1 Score	Hamming Loss
200	0.796	0.092	0.759	0.112
500	0.798	0.093	0.643	0.174
1000	0.783	0.097	0.523	0.241
1500	0.776	0.102	0.407	0.331

V. CONCLUSION

This paper presents a principled privacy-preserving framework for Non-Intrusive Load Monitoring (NILM) that enables selective disaggregation under formal privacy guarantees. A calibrated Gaussian noise injection mechanism is introduced to obfuscate aggregate power data, deterring unauthorized NILM inference while allowing authorized entities to recover appliance-level usage through noise-aware inference using shared noise statistics. The approach extends naturally to both probabilistic and machine-learning-based NILM models, achieving differential privacy and lightweight, real-time compatibility. Experiments on the GREEND dataset show that the proposed method effectively balances privacy protection with NILM accuracy. It maintains high disaggregation accuracy for authorized users while substantially degrading performance for unauthorized ones.

Acknowledgment. We wish to acknowledge the support of Kuwait Foundation for the Advancement of Sciences (KFAS).

REFERENCES

[1] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.

[2] P. A. Schirmer and I. Mporas, "Non-intrusive load monitoring: A review," *IEEE Transactions on Smart Grid*, vol. 14, no. 1, 2022.

[3] J. Kelly and W. Knottenbelt, "The uk-dale dataset, domestic appliance-level electricity demand and whole-house demand from five uk homes," *Scientific data*, vol. 2, no. 1, pp. 1–14, 2015.

[4] A. Monacchi, D. Egarter, W. Elmenreich, S. D'Alessandro, and A. M. Tonello, "Greend: An energy consumption dataset of households in italy and austria," in *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. IEEE, 2014, pp. 511–516.

[5] F. Yang, B. Liu, W. Luan, B. Zhao, Z. Liu, X. Xiao, and R. Zhang, "Fhmm based industrial load disaggregation," in *2021 6th Asia conference on power and electrical engineering (ACPEE)*. IEEE, 2021, pp. 330–334.

[6] C. Zhang, M. Zhong, Z. Wang, N. Goddard, and C. Sutton, "Sequence-to-point learning with neural networks for non-intrusive load monitoring," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.

[7] P. A. Schirmer and I. Mporas, "On the non-intrusive extraction of residents' privacy-and security-sensitive information from energy smart meters," *Neural Computing and Applications*, pp. 1–14, 2023.

[8] H. Wang, J. Zhang, C. Lu, and C. Wu, "Privacy preserving in non-intrusive load monitoring: A differential privacy perspective," *IEEE Transactions on Smart Grid*, vol. 12, no. 3, pp. 2529–2543, 2020.

[9] Bidgely Inc., "Energy disaggregation services," https://www.bidgely.com, accessed: Nov, 2025.

[10] Verdigris Technologies, "Ai-powered energy monitoring," https://verdigris.co, accessed: Nov, 2025.

[11] S. Banik, M. Rogers, S. M. Mahajan, C. M. Emeghara, T. Banik, and R. Craven, "Survey on vulnerability testing in the smart grid," *IEEE Access*, 2024.

[12] Electronic Frontier Foundation, "Sacramento county resident joins eff lawsuit after illegal sharing of his electricity usage data," 2022, accessed: 2025-07-08. [Online]. Available: https://www.eff.org/deeplinks/2022/11/sacramento-county-resident-joins-eff-lawsuit-after-illegal-sharing-his-electricity

[13] Tech Monitor. (2025) Nova scotia power breach exposes customer data. Accessed: 2025-07-02. [Online]. Available: https://www.techmonitor.ai/technology/cybersecurity/nova-scotia-power-data-breach-exposes-customer-information

[14] The Record. (2023) Texas utility firm investigates moveit breach. Accessed: 2025-07-02. [Online]. Available: https://therecord.media/texas-utility-firm-investigating-potential-data-leak-moveit-breach

[15] M. Jin, R. Jia, Z. Kang, I. C. Konstantakopoulos, and C. J. Spanos, "Presencesense: Zero-training algorithm for individual presence detection based on power monitoring," in *Proceedings of the 1st ACM conference on embedded systems for energy-efficient buildings*, 2014, pp. 1–10.

[16] H. Wang, C. Si, G. Liu, J. Zhao, F. Wen, and Y. Xue, "Fed-nilm: A federated learning-based non-intrusive load monitoring method for privacy-protection," *Energy Conversion and Economics*, vol. 3, no. 2, pp. 51–60, 2022.

[17] S. Dai, F. Meng, Q. Wang, and X. Chen, "Federatednilm: A distributed and privacy-preserving framework for non-intrusive load monitoring based on federated deep learning," in *2023 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2023, pp. 01–08.

[18] ——, "Dp2-nilm: A distributed and privacy-preserving framework for non-intrusive load monitoring," *Renewable and Sustainable Energy Reviews*, vol. 191, p. 114091, 2024.

[19] J. He, J. Wang, N. Wang, S. Guo, L. Zhu, D. Niyato, and T. Xiang, "Preventing non-intrusive load monitoring privacy invasion: A precise adversarial attack scheme for networked smart meters," *arXiv preprint arXiv:2412.16893*, 2024.

[20] P. A. Schirmer, I. Mporas, and M. Paraskevas, "Energy disaggregation using elastic matching algorithms," *Entropy*, vol. 22, no. 1, p. 71, 2020.

[21] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.

[22] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.