A Physics-Guided Framework for Solar-Aware Load Disaggregation in PV-Integrated Households

Muhammad Affan Khan*†, Giulia Tanoni*, Enrik Xhani*, Stefano Squartini* and Emanuele Principi*
*Department of Information Engineering, Università Politecnica delle Marche, Italy
†Department of Information Engineering, Electrical Engineering, and Applied Mathematics, University of Salerno, Italy
Email: m.a.khan@pm.univpm.it

***Abstract*—The ever increasing adoption of residential photovoltaic (PV) systems fundamentally alters true household consumption and therefore challenges the applicability of conventional Non-Intrusive Load Monitoring (NILM) techniques. In this paper, we propose a physics-guided neural network for the disaggregation of household appliances, as well as behind the meter solar generation using the net load signal and local irradiance measurements. Unlike conventional methods that regress solar power directly, our method introduces a latent Correction Factor (CF) that is estimated alongside appliance power profiles. This predicted CF scales the irradiance input to reconstruct solar generation, strictly enforcing physical constraints (e.g. zero output under no irradiance) directly into the learning process. This formulation enables the CF to implicitly capture PV system efficiency variations due to environmental and operational factors without explicit supervision. Experimental evaluation on real world datasets demonstrate that the proposed approach achieves improved disaggregation accuracy and enhanced consistency compared to unconstrained solar regression baselines in effectively decoupling household consumption from distributed generation in PV-integrated environments.**

***Index Terms*—Non-Intrusive Load Monitoring, Physics-Guided Neural Networks, PV Awareness, Deep Learning.**

I. Introduction

The techniques for Non-Intrusive Load Monitoring (NILM) provide a means to infer appliance level energy consumption from aggregate household measurements, eliminating the need for intrusive and costly sub-metering infrastructure [1, 2]. This technical insight is increasingly vital for smart grid ecosystems, enabling demand response participation, data-driven energy efficiency strategies and much more effective grid management [3–5]. However, the rapid growth of residential photovoltaic (PV) adoption, driven by declining installation costs and renewable energy incentives, has diminished the visibility of appliance load profiles in smart meter data [6] and rendered conventional NILM approaches less effective in solar integrated environments.

To address the influence of PV systems, several efforts have extended the NILM framework to account for solar generation. As demonstrated by Chen et al. in [7], they used proxy solar data and iterative algorithms to infer PV output characteristics and improve disaggregation when ground truth generation data were unavailable. Another framework proposed by Farzana et al. [8], integrates physical PV system models with statistical load estimators such as hidden Markov models, iteratively estimate solar and native load from net measurements with reduced error. Some recent real-time disaggregation frameworks aim to estimate PV generation and load jointly from smart meter data without detailed system dynamics [9, 10], showing improved robustness and accuracy over conventional algorithms. One of the approaches also employ hybrid data-driven deep learning systems that learns to separate demand and generation from net load using bidirectional recurrent architectures or CNN-LSTM models [11], treating solar output as an additional regression target.

Despite the innovative progress, a limitation remains that many of the existing solar disaggregation models treat PV generation as an unconstrained regression target, thereby ignoring the well established physical relationship between irradiance and PV output. Unconstrained data-driven regression may provide physically implausible predictions, such as non-zero generation during night hours or negative power values. It may also encounter difficulties in generalizing across seasonal and environmental variations. To overcome these deficiencies, research in PV estimation and forecasting has increasingly explored hybrid machine learning approaches [12–15], that integrate physical knowledge with data-driven learning. Comprehensive reviews of hybrid methods for photovoltaic power estimation [16] indicate that embedding physical models or constraints into machine learning algorithms, enhances generalization, robustness and interpretability by limiting predictions to physically plausible regions and amalgamating the advantages of both modeling approaches. For example, physics-guided models that incorporate physical PV conversion principles alongside learned components have been shown to prevent illogical outputs and better capture environmental dependencies inherent in PV generation processes [17], especially under variable irradiance and weather conditions.

Motivated by the limitations of unconstrained solar regression and the benefits of physics-guided learning, we propose a physics-guided neural network for simultaneous disaggregation of appliances and solar energy from net meter data. Rather than directly regressing the PV output, our model predicts a latent Correction Factor (CF) alongside appliance power signatures using net load and local irradiance as inputs. Solar generation is reconstructed as the product of irradiance and predicted CF, explicitly enforcing the physical relationship between environmental conditions and PV production. This formulation naturally ensures physically consistent behavior, such as zero output during non-irradiance periods, while

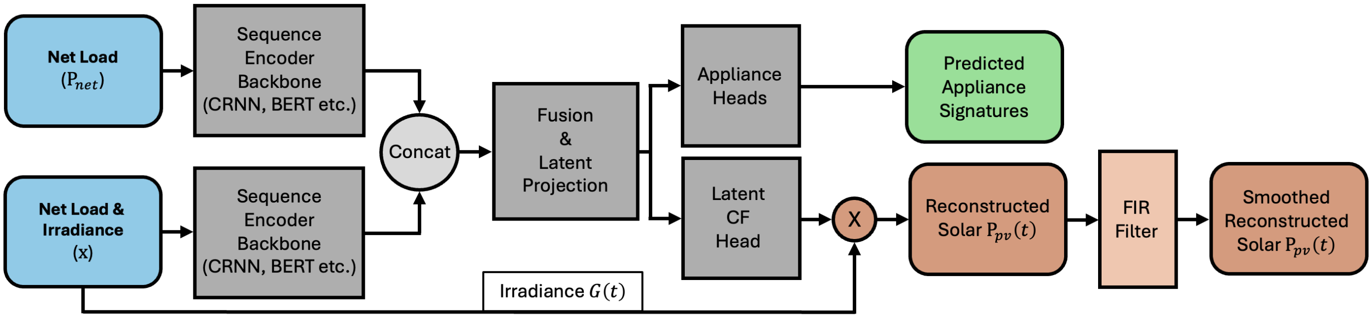

Fig. 1. Schematic view of the designed Deep Learning Architecture based on the Physics-Guided paradigm.

allowing CF to implicitly capture variations in PV system efficiency. The network is trained with a multi-task loss function defined on appliance consumption and reconstructed solar output, thereby avoiding direct supervision of the latent CF. In comparison to the unconstrained data-driven baselines, the proposed approach improves disaggregation accuracy and resilience by embedding the physical structure into the learning process. The main contributions of this paper can be summarized as follows:

- We propose a physics-guided neural network for joint disaggregation of appliances and solar from the net load data, in which solar is reconstructed using an irradiance scaled CF, rather than direct regression, ensuring physical consistency.
- A multi-task learning framework is developed with loss over appliance consumption and reconstructed solar output, enabling implicit learning of PV efficiency variations and improving disaggregation accuracy over unconstrained baselines.

The remainder of the paper is organized as follows: Section II gives details on the proposed method. Section III provides information about the experimental setup. Results are discussed in Section IV and Section V concludes the paper.

II. Proposed Methodology

We consider a residential household equipped with PV system and multiple electrical appliances, where only a single smart meter measurement is available. Objective is to disaggregate the net meter signal into individual appliance level consumption profiles and the underlying solar generation.

A. Net Load Representation

The net power measured by the smart meter, denoted as $P_{net}(t)$, can be modeled as follows:

$$P_{net}(t) = \sum_{i=1}^{N} P_i(t) - P_{pv}(t) + \epsilon(t) \tag{1}$$

where $P_i(t) \in \mathbb{R}$ is the power consumption of appliance i, N denotes the number of target appliances, $P_{pv}(t) \in \mathbb{R}$ gives PV generation (where $P_i(t)$ and $P_{pv}(t)$ are ≥ 0) and $\epsilon(t)$ captures measurement noise and/or unmodeled loads. In this work, $N = 5$, corresponding to kettle, microwave, fridge, washing machine, and dishwasher.

B. Designed Neural Network Architecture

The proposed physics-guided Neural Network (NN), builds upon the sequence modeling principles originally introduced for NILM tasks in [18, 19], and is generalized to support multiple sequence encoder backbones, including CRNN, BERT, Transformer and TransUNet, as illustrated in Figure 1. Each backbone is treated as a modular sequence encoder, enabling a unified architectural formulation across different network designs.

At each time step t, the network receives the net load P_{net} and the corresponding solar irradiance $G(t)$, forming the input vector,

$$\mathbf{x}(t) = [P_{net}(t), G(t)]. \tag{2}$$

This allows explicit incorporation of environmental information for PV modeling. The framework operates on two complementary inputs: the net load signal and the net load with measured irradiance. Each input is processed by an independent instance of the selected sequence encoder to learn consumption-driven and generation-aware temporal representations, respectively. The resulting features are fused through concatenation and latent projection to form a shared representation.

From this shared latent space, the network branches into task-specific heads for appliance power estimation and latent correction factor prediction. The predicted CF modulates the irradiance signal to reconstruct solar generation in a physically guided manner, with an FIR filtering stage applied to enforce temporal smoothness. This design enables coordinated learning of appliance disaggregation and solar modeling while embedding physical constraints ensuring stable and interpretable performance across diverse sequence encoder architectures.

C. Correction Factor Learning

The Correction factor is introduced as a latent variable representing the effective conversion of solar irradiance into PV generation. It implicitly captures system efficiency and environmental effects, such as shading, temperature variations and seasonal changes. Rather than supervising CF directly, the designed model learns it implicitly through the reconstruction of solar power in the forward pass. This approach encourages the network to produce physically plausible PV outputs, including zero generation during periods of no irradiance, without constraining the network to predefined efficiency values.

By treating CF as a latent physics-guided variable, the model decouples data-driven feature learning from physical PV behavior, supporting generalization across different households and PV systems.

D. Physics-Guided PV Generation

Rather than directly regressing PV power, we model PV generation through a physics-guided formulation. PV generation at time t can be expressed as:

$$P_{pv}(t) = CF(t) \cdot G(t), \tag{3}$$

where, as mentioned before, $G(t)$ denotes the measured solar irradiance and $CF(t)$ is the Correction Factor. This specific formulation enforces fundamental physical constraints while allowing CF to capture time varying efficiency effects arising from environmental and operational conditions.

E. Learning Objective

The goal is to jointly estimate the power consumption of N appliances and PV generation from net meter data, while enforcing physical consistency between irradiance and solar output. The network predicts appliance powers $\hat{P}_i(t)$, $i = 1, ..., N$, and the Correction Factor $\hat{CF}(t)$. PV generation is reconstructed as described earlier and given in Equation (3). The CF is not directly supervised and is learned implicitly through the PV reconstruction error.

Appliance disaggregation is trained using a regression loss:

$$\mathcal{L}_{app} = \frac{1}{NL} \sum_{i=1}^{N} \sum_{t=1}^{L} [P_i(t) - \hat{P}_i(t)]^2, \tag{4}$$

where $P_i(t)$ is the ground truth and $\hat{P}_i(t)$ denote the predicted power of appliance i and L is the dimension of input window. Solar generation is supervised via the physics-guided reconstruction loss taken as:

$$\mathcal{L}_{pv} = \frac{1}{L} \sum_{t=1}^{L} [P_{pv}(t) - \hat{CF}(t) \cdot G(t)]^2, \tag{5}$$

where $P_{pv}(t)$ gives the ground truth and $\hat{CF}(t) \cdot G(t)$ gives the reconstructed solar, respectively.

In addition to the appliance level and PV reconstruction losses, a net load consistency loss is introduced to enforce global power balance. This loss ensures that the sum of predicted appliance consumption and reconstructed PV generation remains consistent with the measured net load, as given below:

$$\mathcal{L}_{cons} = \frac{1}{T} \sum_{t=1}^{T} [(P_{app}(t) - P_{pv}(t)) - (\hat{P}_{app}(t) - \hat{CF}(t) \cdot G(t))]^2 \tag{6}$$

where $P_{app}(t)$ gives the sum of ground truth appliance power, $P_{pv}(t)$ provides the sum of ground truth PV power, $\hat{P}_{app}(t)$ is the sum of predicted appliance power, $\hat{CF}(t) \cdot G(t)$ indicate the sum of reconstructed PV power and T is the total number of samples. This term penalizes inconsistencies between the measured net power and the reconstructed net signal obtained from predicted appliance consumption and PV generation, thereby coupling appliance and solar estimates in a physically meaningful manner.

The overall training loss is then taken as:

$$\mathcal{L}_{total} = \mathcal{L}_{app} + \lambda_{pv} \cdot \mathcal{L}_{pv} + \lambda_{cons} \cdot \mathcal{L}_{cons}, \tag{7}$$

here λ_{pv} and λ_{cons} are weighting coefficients that balance appliance disaggregation accuracy, PV reconstruction fidelity and net load consistency. This joint optimization encourages accurate appliance estimation while enforcing physically consistent solar generation and overall power balance.

F. Learnable FIR-based Solar Smoothing

While the proposed physics-guided formulation constraints solar generation through irradiance scaling, short-term fluctuations may still appear in the reconstructed PV signal due to measurement noise and regression uncertainty. To enhance temporal smoothness and improve overall disaggregation stability, a learnable Finite Impulse Response (FIR) filtering layer is applied to the reconstructed solar output.

Following the formulation in [20], let $\hat{P}_{pv}(t)$ denote the solar power reconstructed via the predicted CF and irradiance. The smoothed solar estimate $\tilde{P}_{pv}(t)$ is then obtained as;

$$\tilde{P}_{pv}(t) = \sum_{k=-K}^{K} h(k) \hat{P}_{pv}(t-k), \tag{8}$$

where $h(k)$ are the FIR filter coefficients and $2K + 1$ denote the filter window length. The filter is implemented as a one dimensional convolutional layer, initialized with a moving average kernel to provide stable smoothing at the start of training. Its coefficients are subsequently optimized jointly with the network parameters.

The FIR layer operates exclusively on the solar output channel and does not directly alter appliance predictions. However, by suppressing high frequency artifacts in the solar estimate, it reduces residual compensation between the solar and appliance components during net load reconstruction. Since gradients propagate through the FIR layer during training, this smoothing acts as an implicit regularizer on shared network representations, leading to improved solar stability and indirectly enhancing appliance level disaggregation.

III. Experimental Setting

This section details the experimental setup adopted to ensure full reproducibility of the results.

A. Datasets and Processing

The proposed method is evaluated on two widely used public datasets for residential energy disaggregation, that are UK-DALE [21] and REFIT [22]. These publicly available datasets include measurements from 20 UK households. Aggregate power consumption is sampled at 1 Hz, while appliance level power data are sampled at 1/6 Hz and 1/8 Hz, respectively. Both datasets provide high resolution aggregate smart meter measurements together with sub-metered appliance level power data and are commonly used as benchmarks in NILM research. For UK-DALE dataset, houses 1 and 5 are used for training while house 2 is reserved for testing. For the REFIT dataset, houses 2 and 9 are held out for testing, whereas all others are used for training. In addition to net load measurements, solar irradiance data is simulated based on the location of the household using a clear-sky model. This simulated irradiance is temporally aligned with the power signals and can be used during both training and inference, eliminating the need for additional sensors. The net load is constructed following the formulation in Section II by combining appliance consumption and PV generation. All signals are resampled to a common temporal resolution and normalized to ensure stable training.

B. Experimental Conditions

All experiments are conducted under identical conditions to ensure a fair and unbiased comparison among different evaluations. Specifically, the data preprocessing pipeline, input formatting and normalization procedures are kept consistent across all experimental runs. Likewise, all training related hyperparameters are fixed and shared across experiments to eliminate performance variations caused by parameter tuning rather than model behavior.

Early stopping is employed to mitigate overfitting, with a patience of 10 epochs and a maximum training duration of 1000 epochs. Model optimization is performed using the Adam optimizer and the batch size is fixed at 64. Overall, this controlled experimental setup ensures that any observed performance differences across evaluations can be attributed solely to the proposed modeling or learning mechanisms, rather than variations in training conditions or implementation details.

C. Evaluation Metric

The performance of the proposed framework is evaluated using a metric that quantifies the accuracy of disaggregated appliance signals. Specifically we use the Mean Absolute Error (MAE), defined as,

$$MAE = \frac{1}{T}\sum_{t=1}^{T}\left|P(t) - \hat{P}(t)\right| \tag{9}$$

where $P(t)$ is the ground truth power consumption at time step t, $\hat{P}(t)$ is the corresponding estimated power by the framework and T is the total number of samples in the evaluation window.

The MAE measures the average absolute deviation between the true and estimated power signals over time. Unlike squared error metrics, MAE is less sensitive to occasional large deviations, providing a robust and interpretable indication of typical disaggregation performance. In the context of NILM, MAE is widely used because it directly reflects the expected error in watts, allowing practitioners to assess the practical reliability of the disaggregated outputs for energy management and appliance monitoring applications.

TABLE I
COMPARISON OF MAE (W) VALUES FOR SOLAR AND APPLIANCE DISAGGREGATION ACROSS DIFFERENT NETWORK ARCHITECTURES ON UK-DALE AND REFIT DATASETS.(<u>UNDERLINED</u> VALUES SHOW THE BEST SOLAR MAE, WHEREAS **BOLD** VALUES REPRESENT OVERALL BEST MAE).

Model	Appliance	UK-DALE		REFIT	
		Baseline	Proposed	Baseline	Proposed
CRNN	Solar	88.40±6.58	<u>61.54</u>±1.21	61.30±7.77	<u>40.50</u>±3.06
	Kettle	16.40±2.47	14.30±3.23	21.17±1.86	22.12±0.84
	Microwave	21.90±4.34	21.93±2.12	12.46±5.22	13.57±2.86
	Fridge	35.81±0.42	33.17±1.21	39.06±0.76	39.34±1.42
	Washing Machine	20.92±2.46	18.58±2.50	25.02±4.43	25.25±3.62
	Dishwasher	43.06±4.64	36.17±3.12	38.42±5.88	44.64±3.94
	Avg. Overall	37.75±1.81	**30.94**±1.06	32.91±2.35	**30.90**±1.27
BERT	Solar	95.56±17.89	<u>58.92</u>±4.48	77.09±10.59	<u>48.55</u>±5.69
	Kettle	23.63±9.03	25.97±3.17	22.80±2.45	24.59±2.24
	Microwave	31.73±2.57	25.21±4.27	9.42±1.30	8.43±0.90
	Fridge	37.77±0.23	35.16±0.96	40.48±0.33	40.50±0.87
	Washing Machine	28.58±5.70	28.58±3.89	29.66±2.58	31.21±3.27
	Dishwasher	42.92±3.06	47.00±3.48	47.97±6.28	46.31±3.10
	Avg. Overall	43.37±5.59	**36.80**±3.05	37.90±3.42	**33.26**±0.85
Transformer	Solar	88.16±12.56	<u>58.77</u>±5.0	98.71±12.51	<u>58.98</u>±9.0
	Kettle	23.78±7.81	29.15±3.58	30.65±7.16	28.46±1.36
	Microwave	19.83±4.93	23.76±4.37	9.16±4.38	13.97±4.46
	Fridge	37.46±0.43	36.50±0.55	41.76±1.35	40.87±2.63
	Washing Machine	24.70±2.85	25.70±1.70	24.67±7.30	29.37±3.59
	Dishwasher	48.11±3.85	49.20±5.16	69.61±7.06	71.25±8.71
	Avg. Overall	40.37±4.99	**37.18**±2.51	45.76±3.86	**40.48**±2.69
TransUNet	Solar	<u>66.60</u>±5.38	68.80±6.17	82.59±20.0	<u>35.82</u>±4.49
	Kettle	15.29±2.88	30.09±2.57	37.47±12.08	20.38±1.47
	Microwave	29.31±4.84	26.09±4.59	12.84±9.56	8.74±4.57
	Fridge	33.52±3.13	37.39±1.23	40.82±2.31	39.52±1.20
	Washing Machine	30.26±1.27	23.01±4.60	42.19±7.30	22.81±5.75
	Dishwasher	43.83±4.69	49.34±3.33	76.79±16.64	36.83±2.05
	Avg. Overall	**36.47**±1.33	39.12±1.89	48.78±12.01	**27.35**±1.84

IV. Results

A. Evaluation of the Core Physics-Guided Framework

In this section, we evaluate the proposed physics-guided solar-aware NILM framework without the inclusion of consistency loss and FIR filtering components. This configuration allows us to isolate and assess the core contribution of the CF based solar modeling strategy across multiple network architectures and datasets. The consistency loss and FIR filtering are introduced later as targeted refinements and are evaluated separately to avoid conflating their effects with the core framework.

Table I presents an MAE based comparison between the baseline and the proposed physics-guided methods, across four network architectures: CRNN, BERT, Transformer and TransUNet on both the UK-DALE and REFIT datasets. Under the baseline setting, all architectures exhibit high solar disaggregation errors on both datasets, indicating the inherent difficulty of directly regressing solar power from net meter readings. While TransUNet achieves the lowest baseline solar MAE (66.60 W), appliance level performance remains inconsistent across architectures, indicating residual solar interference in appliance learning.

The proposed physics-guided framework consistently improves performance across all architectures and datasets. On UK-DALE, solar MAE is reduced for all models with CRNN, BERT and Transformer achieving values of 61.54 W, 58.92 W and 58.77 W respectively. These improvements demonstrate that decoupling solar estimation from direct regression yields more stable solar predictions. Enhanced solar modeling translates into cleaner residual load signals, resulting in systematic reductions in average appliance MAE. Notably, CRNN

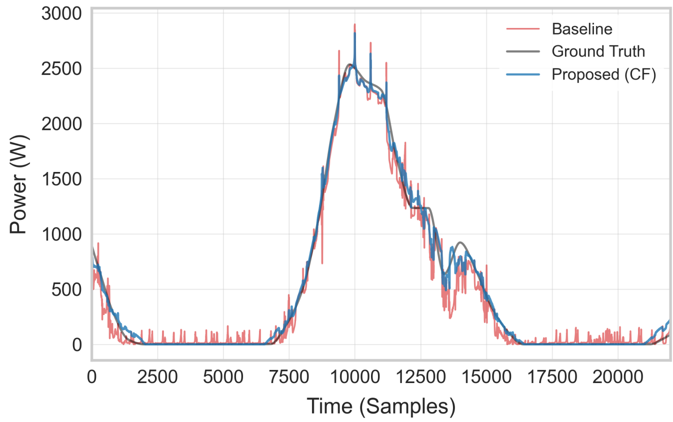

Fig. 2. Performance comparison for Solar disaggregation. The proposed Physics-Guided method (blue) strictly adheres to physical constraints (zero generation at night), whereas the unconstrained Baseline (red) exhibits significant noise during non-generation periods. The ground truth is shown in gray.

achieves the lowest average appliance MAE and the lowest overall MAE (30.94 W) on UK-DALE, indicating a strong balance between solar and appliance disaggregation.

A similar but more pronounced trend is observed on the REFIT dataset, which is characterized by higher PV variability and household diversity. Under the baseline configuration, solar MAE is substantially higher than appliance MAE for all architectures, particularly for Transformer and TransUNet, leading to elevated overall MAE values (48.78 W for TransUNet). The proposed physics-guided method yields significant and consistent reductions in solar MAE across all models, with the most notable improvement observed for TransUNet, where solar MAE drops from 82.59 W to 35.82 W. These gains are accompanied by reductions in average appliance MAE, confirming that explicit solar decoupling produces cleaner residual loads. Under the proposed framework, TransUNet achieves the best overall MAE (27.35 W) on REFIT, highlighting the advantage of architectures with strong spatio-temporal modeling capacity in more complex PV-integrated environments.

Overall the results in Table I demonstrate that the proposed physics-guided CF framework consistently outperforms the baseline across architectures and datasets, confirming its robustness and generalizability. Figures 2 and 3 further illustrate representative disaggregation examples on UK-DALE dataset with CRNN model for solar and three appliances (kettle, dishwasher and fridge). While CRNN yields the best overall performance on UK-DALE, TransUNet emerges as the most effective model on REFIT, underscoring the influence of dataset complexity on architectural suitability.

B. Evaluation of Consistency and FIR Enhanced Framework

Due to the space limitations, a full-scale evaluation of all framework variants across multiple architectures and datasets is not feasible. Consequently, the assessment of consistency loss and FIR-based filtering is restricted to the CRNN architecture, which demonstrated superior performance based on the averaged results across the UK-DALE and REFIT datasets during core evaluation. The targeted analysis enables a controlled assessment of the incremental benefits introduced by these refinement components while avoiding redundancy across architectures that already exhibit consistent trends under the CF-aware formulation. CRNN is selected as a representative model due to its strong overall performance, stable appliance level disaggregation and widespread adoption in NILM literature.

Table II reports an MAE based comparison of three CRNN variants evaluated on the UK-DALE dataset, all built upon the proposed physics-guided solar modeling framework: (i) the CF only model, (ii) the CF model with consistency loss, and (iii) the full framework incorporating both consistency loss and FIR-based filtering. The CF only configuration serves as a stable baseline within the proposed framework, achieving a solar MAE of 61.54 W with moderate appliance level performance. While the correction factor effectively decouples solar generation from the net signal, the absence of explicit temporal regularization limits further improvements in appliance disaggregation.

Introducing the consistency loss reduces solar MAE (61.54 W to 59.28 W), indicating improved temporal coherence in solar estimation. However, appliance level gains remain marginal and slightly inconsistent, suggesting that residual high frequency fluctuations continue to affect appliance predictions.

The full framework, yields the most substantial improvements. Solar MAE is further reduced to 54.66 W, while appliance level disaggregation improves consistently across all appliances, resulting in a significant reduction in average appliance MAE and the lowest average overall MAE of 27.47 W. These results confirm that the consistency loss and FIR filtering act as complementary refinements, enhancing temporal stability and residual signal smoothness within the physics-guided formulation.

V. Conclusions

The variability of behind the meter solar generation presents significant challenges for standard NILM algorithms, often leading to physically inconsistent predictions when applied to net load data. We proposed a physics-guided neural network approach that utilizes a latent Correction Factor to enforce

TABLE II
MAE (W) PERFORMANCE COMPARISON OF DIFFERENT CRNN VARIATIONS ON THE UK-DALE DATASET. (UNDERLINED IS BEST SOLAR AND **BOLD** IS BEST OVERALL).

Appliance	CRNN Variations		
	Core (*CF only*)	Improved (*CF* + $\mathcal{L}_{cons}$)	Enhanced (*CF* + $\mathcal{L}_{cons}$ + *FIR*)
Solar	61.54±1.21	59.28±3.14	<u>54.66</u>±5.44
Kettle	14.30±3.73	12.69±2.85	10.63±3.17
Microwave	21.93±4.12	23.40±3.62	19.43±3.49
Fridge	33.17±1.21	34.10±1.88	31.33±2.23
Washing Machine	18.58±5.50	20.48±4.12	16.86±3.56
Dishwasher	36.17±3.12	34.80±2.45	31.93±1.42
Avg. Overall	30.94±1.06	30.79±1.64	**27.47**±2.03

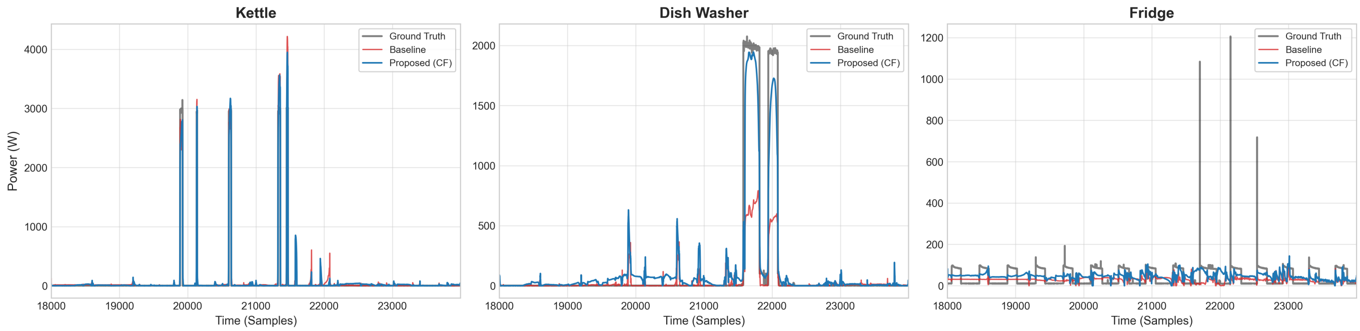

Fig. 3. Disaggregation results for individual appliances (Kettle, Dishwasher, Fridge). The plots compare the Ground Truth (gray), the unconstrained Baseline (red), and our proposed Physics-Guided CF Method (blue). The proposed method tracks the ground truth more accurately, particularly in preserving the correct activation peaks compared to the noisy baseline.

physical constraints during the simultaneous disaggregation of appliances and solar generation. Experimental results on the UK-DALE and REFIT datasets demonstrated that this approach consistently outperforms unconstrained baselines across multiple architectures. By dynamically scaling irradiance in the forward pass, the method achieved substantial reductions in solar estimation error, without compromising the accuracy of appliance disaggregation. Future works will extend this framework for even more complex environments, such as industrial or commercial locations with varying appliances and much larger solar installations.

References

[1] F. Luo. Personalized residential energy usage recommendation system based on load monitoring and collaborative filtering. *IEEE Transactions on Industrial Informatics*, 17(2):1253–1262, 2021.

[2] G. Tanoni, E. Principi, L. Mandolini, and S. Squartini. Appliance-incremental learning for non-intrusive load monitoring. In *IEEE SmartGridComm*, pages 1–6, 2023.

[3] A. Vavouris, L. Stankovic, and V. Stankovic. A non-intrusive load monitoring-enabled framework for load scheduling in the dairy industry. *Applied Energy*, 398:126456, 2025.

[4] G. Liu, G. Liang, H. Zhao, J. Zhao, J. Liu, and Z. Wu. A temporal convolutional neural network with attention mechanism for industrial non-intrusive load monitoring. In *2021 IEEE 5th (EI2)*, pages 3279–3284, 2021.

[5] G. Tanoni, E. Principi, and S. Squartini. Non-intrusive load monitoring in industrial settings: A systematic review. *Renewable and Sustainable Energy Reviews*, 202:114703, 2024.

[6] M. Salani, M. Derboni, V. Medici, L. Nespoli, and A. E. Rizzoli. Non intrusive load monitoring for demand side management. *Energy Informatics*, 3, 2020.

[7] X. Chen, M. Moghimi, and O. Ardakanian. A data-efficient approach to behind-the-meter solar generation disaggregation. *ArXiv*, abs/2105.08122, 2021.

[8] F. Kabir, N. Yu, W. Yao, R. Yang, and Y. Zhang. Estimation of behind-the-meter solar generation by integrating physical with statistical models. In *2019 IEEE SmartGridComm*, 2019.

[9] H. M. Usman, R. ElShatshat, and A. H. El-Hag. A novel non-intrusive framework for real-time disaggregation of behind-the-meter solar generation from smart meter data. *Electric Power Systems Research*, 225:109831, 2023.

[10] M. H. Saleem, M. Taha, and M. A. A. Rehmani. A comprehensive review of machine learning and deep learning models for non-intrusive load monitoring: performance, analyses, practical insights, and emerging trends. *Applied Intelligence*, 55:1020, 2025.

[11] X. Y. Zhang, S. Kuenzel, P. Guo, et al. A hybrid data-driven online solar energy disaggregation system from the grid supply point. *Complex & Intelligent Systems*, 9:3695–3716, 2023.

[12] L. O. Santos, T. AlSkaif, G. C. Barroso, and P. M. Carvalho. Photovoltaic power estimation and forecast models integrating physics and machine learning: A review on hybrid techniques. *Solar Energy*, 284:113044, 2024.

[13] Y. Li, L. Song, S. Zhang, L. Kraus, T. Adcox, R. Willardson, A. Komandur, and N. Lu. A tcn-based hybrid forecasting framework for hours-ahead utility-scale pv forecasting. *IEEE Transactions on Smart Grid*, 14(5):4073–4085, 2023.

[14] Martin J. M. Benefits of physical and machine learning hybridization for photovoltaic power forecasting. *Renewable and Sustainable Energy Reviews*, 168:112772, 2022.

[15] G. Tanoni, R. L. Taloma, E. Principi, D. Comminiello, and S. Squartini. A deep cascade framework for non-intrusive power disaggregation in solar-powered households. In *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2025.

[16] R. Ahmed, V. Sreeram, Y. Mishra, and M.D. Arif. A review and evaluation of the state-of-the-art in pv solar power forecasting: Techniques and optimization. *Renewable and Sustainable Energy Reviews*, 124:109792, 2020.

[17] J. B. Jahangir and M. A. Alam. Physics-guided machine learning predicts the planet-scale performance of solar farms with sparse, heterogeneous, public data. *Applied Energy*, 396: 126192, 2025.

[18] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 EMNLP*, pages 1724–1734, October 2014.

[19] G. Tanoni, E. Principi, and S. Squartini. Multilabel appliance classification with weakly labeled data for non-intrusive load monitoring. *IEEE Transactions on Smart Grid*, 14(1):440–452, 2023.

[20] B. Wang and J. Saniie. Learning fir filter coefficients from data for speech-music separation. In *2020 IEEE International Conference on Electro Information Technology (EIT)*, pages 245–248, 2020.

[21] J. Kelly and W. Knottenbelt. The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes. *Scientific Data*, 2(1):1–14, 2015.

[22] D. Murray, L. Stankovic, and V. Stankovic. An electrical load measurements dataset of United Kingdom households from a two-year longitudinal study. *Scientific Data*, 4(1):1–12, 2017.