How Much Can We Salvage in Disrupted RF Vital Signs Monitoring? A Measurement Study of Post-processing

Leo He*, Zongxing Xie†, Fan Ye*
*Department of Electrical and Computer Engineering, Stony Brook University, {leo.he, fan.ye}@stonybrook.edu
†Department of Computer Science, Kennesaw State University, zxie1@kennesaw.edu

***Abstract*—Continuous monitoring of vital signs offers valuable insights into health status through patterns of changes over time. Radio Frequency (RF) solutions have gained significant attention due to their non-invasive and privacy-preserving nature. Nevertheless, the reliability of RF-based vital signs monitoring remains an ongoing challenge, as RF signals are inherently fragile and susceptible to disruptions, especially due to random body movements. Most of methods proposed to enhance the robustness of RF vital signs sensing reply on accessing and analyzing raw RF signals. However, the diversity of RF configurations and the restricted accessibility of commercial off-the-shelf (COTS) RF solutions make direct analysis of raw RF data impractical for improving accuracy. To address this gap, we propose a novel framework that focuses solely on post-processing to recover the vital signs from disrupted RF signals. Our approach implements a suite of classical smoothing and denoising algorithms, alongside representative data-driven techniques, to rectify noisy and disrupted RF vital signs estimations through data reconstruction. We evaluate these post-processing techniques using a dataset containing 58 hours collected from 3 subjects in cluttered, free-living environments. Our results show that applying Temporal Convolutional Network (TCN) to RF heart rate (HR) estimations doubles the percentage of data below 5 bpm error against ground truth. We additionally find that RF respiration rate (RR) estimations is relatively robust and a simple moving average can increase the percentage of data below 2 bpm error by over 20%. We assess the generalizability of these methods through a leave-one-out evaluation and analyze their respective computational costs, shedding light on practical trade-offs between accuracy and resource requirements.**

***Index Terms*—Radio Frequency (RF), Non-contact vital signs monitoring, Post-processing, Ultra-Wideband (UWB)**

I. Introduction

Monitoring vital signs, such as respiration rate (RR) and heart rate (HR), is crucial for evaluating health status [1]. This practice supports early detection of health changes and enables ongoing tracking of the progression of chronic diseases.

Significant progress has been made in integrating electrocardiogram (ECG) and photoplethysmography (PPG) technology into mobile devices and wearables (e.g., Fitbit, Apple Watch) for vital signs monitoring. However, these devices require frequent charging and wearing, presenting both physical and cognitive challenges, particularly for older adults. Non-contact methods using cameras, sound sensors, or radio frequency(RF) offer a promising alternative. Of those methods, RF-based solutions stand out as they are not as affected by environmental factors (such as lighting and sound) and are less likely to cause privacy concerns (unwanted visual images or recording of conversations).

While RF signals are sensitive in detecting variations in the physical environment, they are inherently fragile [2]. Reliable measurement of subtle chest displacements, such as those caused by heartbeats, is especially challenging, as even slight body movements can easily disrupt RF signal accuracy [3]. For that reason, researchers have proposed multiple methods to strengthen the reliability of RF-based vital signs monitoring. However, many of these methods require modifications to the entire pipeline that processes raw data [4, 5]. Different RF systems vary widely in their hardware specifications, signal processing algorithms, and proprietary protocols, making it difficult to apply a standardized approach to enhancing accuracy. Additionally, COTS RF solutions often have restricted access to low-level signal data, which limits the ability to directly manipulate or optimize signal interpretation. As a result, achieving robust and accurate RF-based vital signs monitoring requires approaches other than relying on raw data analysis.

In this study, we focus exclusively on evaluating the effectiveness of various post-processing methods in recovering vital signs estimations from disrupted RF signals without raw signal access. To evaluate the performance, we employ three metrics: the percentage of the estimations below a specified error, the similarity of vital signs trajectory based on Dynamic Time Warping (DTW), and inference time. Specifically, we aim to answer research questions including: RQ1: How much disturbed RF vital signs estimations can be salvaged solely based on post-processing? RQ2: How do various post-processing methods compare regarding capturing the trend of the vital signs? These insights could lead to choosing the best method to apply after a COTS RF solution to construct a more accurate and representative vital signs estimation.

By answering the above questions, our contributions are summarized as follows:

- We propose a novel framework to salvage disrupted RF vital signs estimations, focusing solely on post-processing the estimated trajectories rather than relying on raw RF signal analysis, which works around the challenges posed by heterogeneity and limited accessibility of COTS RF methods.
- We conduct comprehensive experiments to evaluate the proposed post-processing methods using data collected

This work is supported in part by NSF grants 1951880, 2119299

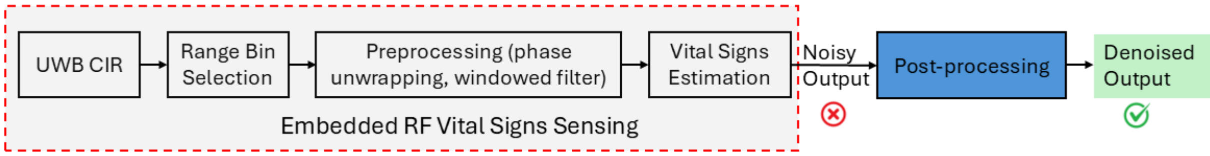

Fig. 1: A generic pipeline for RF vital signs sensing (illustrated using UWB as the RF front-end, without loss of generality).

over 58 hours from cluttered, free-living settings, demonstrating error reduction in vital signs monitoring.
- We analyze the generalizability of the post-processing methods using leave-one-out evaluation, and compare their computational costs, highlighting their practical trade-offs, implying their usefulness in respective application scenarios.

The contribution of this study lies in understanding how different post-processing methods can improve the results of RF solutions when there is no access to raw RF signals. These insights can lead to better accuracy when utilizing COTS RF-based vital signs monitoring solutions where there is limited access to raw RF signals. To the best of our knowledge, this is the first systematic study focused exclusively on post-processing approaches to recovering disrupted RF vital signs estimation without access to raw signals.

II. Related Work

Radar-based non-contact vital signs detection has garnered significant research interests over the past decade [6, 7]. Much of this work has concentrated on developing innovative embedded RF sensing systems to improve vital signs detection accuracy [2]. Although a wide range of frameworks [4, 6] and specific components [3, 8] have been explored for improved robustness, they all rely on accessing and analyzing raw RF signals. For example, the differential measurement of RF signals between a pair of radar sensors was introduced to mitigate interference from random body movements [6]. With recent advances in deep learning approaches, one representative work, MoVi-Fi ([5]) uses contrastive learning on noisy RF vital signals to detect and recover disrupted vital signs estimation.

To address the limited access of raw RF signals, a subset of studies focuses exclusively on post-processing techniques. For instance, a customized CNN was introduced for post-processing [9], demonstrating its ability to derive reasonable vital sign estimations from noisy outputs.

Despite these advancements, there remains a lack of comprehensive investigations into the effectiveness of various post-processing methods applied specifically to noisy outputs. We address this gap by systematically evaluating various post-processing methods for improved RF vital signs monitoring.

III. Preliminaries

In this section, we focus on the general way in which vital signs are extracted using RF sensors. RF-based vital signs extraction typically involves the transmission of RF signals, which are reflected back from the human body. Respiration causes periodic chest movements. Similarly, the chest cavity expands and contracts with each heartbeat. RF signals can detect subtle variations in the chest's expansion and contraction to obtain a vital signs estimate.

A. Embedded RF Vital Signs Sensing

The process of converting the raw data, or signal received from the RF device, into vital signs is grounded in the following signal modeling approach:

$$y(t) = \alpha(t) e^{-j2\pi f_c \frac{2d(t)}{c}} s\left(t - \frac{2d(t)}{c}\right), \tag{1}$$

where $\alpha(t)$ is the instant amplitude, f_c is the carrier frequency, c denotes the speed of radio waves, $d(t)$ is the distance from the chest wall to the RF sensor, and $s(t)$ is the waveform (agnostic to either IR-UWB, FMCW, or others).

The basic structure of the pipeline for RF vital signs estimation is shown in Figure 1 [10]. We will now go over each component, outlining its general function as well as what it does in the specific UWB-based pipeline we chose to use [1], without loss of generality following the signal modeling in (1).

1) UWB Channel Impulse Response (CIR): The pipeline begins with the raw RF data. In our case, since we use a UWB sensor as the RF front-end, we will analyze UWB CIR from the captured reflections (in baseband).

2) Range Bin Selection: When a RF signal is sent and reflected back from an object, the time it takes to return to the receiver is measured, allowing for distance estimation. Because of that, the raw data, the data captured from the RF device is split up into range bins, each representing a specific distance interval. To determine which range bin to use, there are multiple strategies, such as identifying the bin with the highest signal peak. To ensure reliable RF vital signs monitoring, precise detection and ranging of the target (chest wall) for processing corresponding reflections is crucial; however, this is beyond the scope of this work.

3) Preprocessing: For our pipeline, we started by extracting the phase from the raw signal. We used a 30 second interval so that we would have a collection of data over a period of time. We then applied filters to isolate specific frequencies for respiration rate and heart rate respectively.

4) Vital Signs Estimation: From the preprocessed data, the periodicity of cyclic variations in RF signals are extracted based on both temporal and spectral analysis, creating a noisy vital sign estimate.

B. Post-processing

When using data from external sources (eg. other studies) or COTS RF devices, the ability to modify or control the raw signal is often limited or entirely unavailable. The only "data" we would have available would be the "Noisy Output" from after the Embedded RF Vital Signs Sensing pipeline.

Data collection, especially in real-world environments, is often hindered by the ease with which raw signals can be disrupted. Radio signals are particularly vulnerable to interference from environmental factors, multipath effects, and movement. Traditionally, the approach to addressing this issue has been to employ signal quality analysis. By analyzing the raw signal quality, researchers can assess its reliability and choose which pieces to keep and to what extent. Without that data, we would have to search for alternative methods to improve the accuracy of the results

As shown in Figure 1, the only step that modifies noisy output without access to the Embedded RF Vital Signs Sensing pipeline is Post-processing. By investigating the improvements that various postprocessing methods can offer, we can identify methods to enhance results, even in the absence of control over earlier stages of the pipeline.

C. Modeling Post-processing Problem

Given that we do not have access to the raw signal data, the post-processing problem can be formally described as attempting to recover the true vital sign signals from the noisy post-pipeline data. This is represented by the equation:

$$p(t) = vs(t) + n(t) \tag{2}$$

where $p(t)$ is 'Noisy Output' data, $vs(t)$ is the true vital signs, and $n(t)$ is the noise, the distribution of which is, unknown. In essence, our objective is to examine the performance of various methods across different conditions for recovering vital signs trajectory, thus enhancing the accuracy and reliability of RF vital signs sensing.

IV. METHOD

A. Data Collection

To obtain real-world results, we collected data from three participants (two males and one female) over eight nights, each lasting 6–8 hours, resulting in a total dataset over 58 hours.

Data collection was conducted using four IR-UWB (X4-XeThru) sensors [11], operating at 10 frames per second, deployed in each participant's room: one beneath the bed, one beside the bed, one on the ceiling, and one at the foot of the bed. To measure the groundtruth of vital signs, we used a Masimo Pulse Oximeter [12], a FDA approved device, which provided a vital signs reading once per second.

B. Post-processing Methods

For the experiment, we used the post processing methods shown in Table I. We chose the windowed filters and the forecasting method as four simple traditional denoising approaches that require no training.

We explore Random Forest and Gradient Boosting Regressor (GBR) as baseline traditional non neural network supervised learning methods [13][14].

While GBR optimizes around loss, Random Forest focuses on reducing variance through an ensemble of decision trees, each trained on different data subsets. This distinction allows GBR to be more sensitive to subtle patterns by iteratively correcting errors, while Random Forest tends to be more robust to overfitting, particularly with highly variable data.

Group	Method	Setup
1	Mean	Windowed filter with mean kernel
	Gaussian	Windowed filter with Gaussian kernel
	Median	Windowed filter with median kernel
	Kalman	**E**: 0, **Var**(**Process**: 1e-5; **Measure**: 0.25)
2	Random Forest	**Trees**: 100, **Max Depth**: None
	GBR	**Learning Rate**: 0.1, **Trees**: 100, **Max Depth**: 3
	SVR	**Degree**: 3, **Epsilon**: 0.1
3	TCN	**Filters**: 64, **Kernel Size**: 3, **Dilations**: 1,2,4,8,16
	LSTM-FCN	**Batch Size**: 128, **Kernel Size**: 8,5,3
	Multirocket	**Dilation**: 32, **Kernel Size**: 4, **Kernels**: 10,000
	MLP	**Activation**: ReLu, **hidden layers**: 100

TABLE I: Post-processing methods categorized into: 1) classical smoothing techniques; 2) supervised learning methods for denoising; 3) representative neural network models.

Support Vector Regression (SVR) approaches regression differently by finding a hyperplane that best fits the data within a specified margin of error. SVR aims to minimize errors by maximizing the margin around the hyperplane, making it especially effective for handling high-dimensional data and cases where noise needs to be carefully managed.

For time series regression, the two most commonly used neural network architectures are Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks. CNNs are effective at capturing local patterns, while LSTMs excel in retaining long-term dependencies. We used Temporal Convolutional Network (TCN)[15] as a representative CNN variation designed for capturing temporal relationships in time series. Instead of using a standalone LSTM, we opted to experiment with an LSTM/CNN hybrid model, specifically the LSTM Fully Convolutional Network (LSTM-FCN) [16], to leverage the strengths of both architectures.

Additionally, we used the MultiRocket neural network [17]. Although technically a type of CNN, MultiRocket is specifically designed for time series classification and regression tasks. It applies multiple convolutional kernels with varying dilation rates and sizes and uses the results as features for a SVM-like logistic regressor; the kernels allow for it to potentially capture both long and short term trends.

Finally, we selected a Multi-Layer Perceptron (MLP) [18] for its straightforward, fully connected architecture that can efficiently model complex relationships within the data. Unlike CNNs and LSTMs, MLPs do not assume any spatial or temporal structure, making them a flexible baseline model.

C. Post-processing Setup

1) Non data based methods: For the mean, median, and gaussian filters, we used a moving window of 30 seconds to predict one second of ground truth. For the Kalman filter, we forecast ground truth by running the filter through the pipeline data.

2) Supervised Learning Methods: Most of the implemented methods were directly sourced from established libraries sklearn.ensemble and sktime.regression [19][20]. Exceptions to this include the Temporal Convolutional Network (TCN), which was implemented using Keras [21], and MultiRocket, which utilized its transformation neural network from sktime.transformations while employing the logistic regressor as described in the original MultiRocket paper [17].

For all methods requiring supervised learning, a 30-second moving window derived from the pipeline was utilized to predict a one-second segment of the ground truth. Two distinct approaches to training and testing data splits were employed.

- **Global Evaluation:** We integrate all of the data from all of the nights and sensors into one large data set. We then use a 80/20 split (80% for training, 20% for testing) as a traditional methodology for evaluting methods' performance.
- **Cross-validation Evaluation:** We perform leave-one-out training for each sensor. For each sensor's measurements from each night, we train the method using all remaining data from that sensor across other nights. This approach allows us to evaluate method performance on test data that is reasonably distinct from the training data, assessing the methods' generalizability with a limited dataset.

3) Neural Networks (NNs): For all NN methods, we used the ADAM optimizer and measured loss using mean squared error. During training, the number of epochs was determined by an early stopping criterion, where training halted if no improvement was observed for 10 consecutive epochs.

4) Normalization: Data was normalized using global normalization: We found the mean and standard deviation of all the training data (from both the pipeline and the ground truth). We then normalize all training data around this global mean and standard deviation before training. During testing, we normalize the testing input around the same global mean and standard deviation and then de-normalize the output after applying the post-processing method.

V. Evaluation

A. Evaluation Metrics

To measure performance, we used three metrics: [1]

- **Percentage of data within 10% error of the ground truth.** After post processing, we find the percentage of the data within 5 bpm of the ground truth for heart rate and within 2 bpm of the ground truth for respiration rate. This metric gives us an idea of what percentage of the data would be "salvageable". The choice of 2 bpm for respiration rate and 5 bpm for heart rate would be an approximately 10% error based off of average respiration/heart rate respectively.
- **DTW score.** DTW score is a relative score that determines the similarity between two time series. We will be using this to see which methods are best capable of capturing the trend of the data regardless of the actual amount of data the method can salvage.
- **Inference time.** Inference time, measured in miliseconds, is the time it takes for a single instance to go through the method. We will use the average testing time for supervised learning methods. This gives us an idea of how well these methods would perform in real-world environments where we may need instantaneous calculations.

[1]For cross-validation evaluation, DTW and time are averaged across all sessions with different sensor/camera combinations.

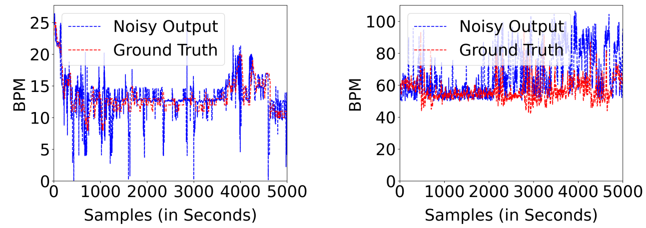

(a) RR (b) HR

Fig. 2: Shown above is a sample of Noisy Output vs Ground Truth. Before post processing, the data for Respiration Rate reasonably closely follows ground truth with a few evident outliers.

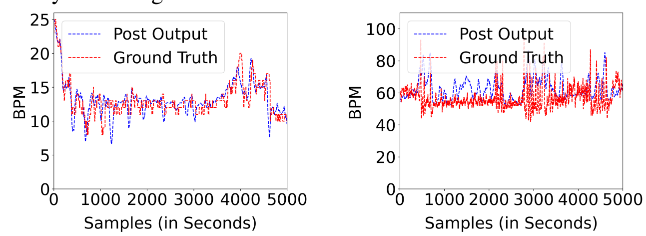

(a) RR (b) HR

Fig. 3: Shown above is a sample of Denoised Output vs Ground Truth. This is using the best performing denoising method: mean for respiration rate and TCN for heart rate.

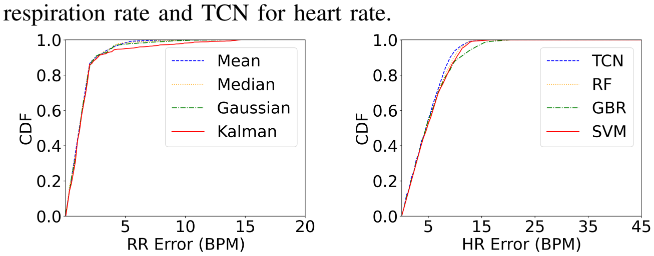

(a) RR (b) HR

Fig. 4: Shown above is the CDF of error between post processing results and ground truth for the best four performing methods, according to "percentage within 10% error", for respiration rate and heart rate respectively.

B. Respiration Rate

1) Global Evaluation: Our global evaluation allows for us a simple way to asses overall performance of post processing methods assuming sufficient similar data. Following the global data split for training and testing, analysis of the noisy output data extracted from the pipeline indicates that 65% of predictions fall within a 2 bpm margin of the ground truth. A plot of noisy output vs ground truth can be seen in Figure 2. These results suggest that even without post-processing, the pipeline achieves a substantial degree of accuracy.

Method	RR			HR		
	% in 2 bpm	DTW	Inference Time (ms)	% in 5 bpm	DTW	Inference Time (ms)
Noisy Output (Baseline)	65.33	0.121	–	22.23	2.01	–
Mean	**87.22**	**0.082**	0.0001	31.92	1.19	0.0001
Gaussian	86.23	0.093	0.0002	32.24	1.39	0.0003
Median	85.46	0.094	0.0001	33.19	1.23	0.0001
Kalman	85.21	0.101	0.0002	35.62	1.08	0.0004
Random Forest	78.21	0.107	3.22	48.66	**0.35**	2.89
GBR	77.32	0.118	1.88	47.23	0.42	1.49
SVR	76.27	0.123	100.21	45.66	0.44	100.54
TCN	79.23	0.108	2.42	**48.92**	0.41	2.82
LSTM FCN	72.97	0.181	10.92	33.21	0.63	12.93
Multi Rocket	74.88	0.113	30.31	44.98	0.70	30.25
MLP	79.23	0.113	34.66	37.83	0.73	33.29

TABLE II: method Performance in Global Split. Best-performing results in each category are bolded.

As shown in Table II, although the unprocessed data demonstrates reasonable accuracy, the application of all methods yields enhanced performance, both in percentage of data considered "salvageable" and in similarity to ground truth. The best performing methods for respiration rate were actually the non data based methods with the best method, mean, increasing salvageable data from 65% to 87.22%. However, the CDF of error between denoised output and ground truth, Figure 4, shows that there is similar performances between all non data based methods. Table II indicates that non data-driven methods achieve the highest similarity to the ground truth. However, the extent of improvement by DTW is comparatively limited with only mean standing out. Applying mean to our noisy data creates a significantly closer to ground truth output as shown in Figure 3. For data that is easy to determine from radio frequency, mean offers promising results with little to no trade off or complexity.

2) Cross-validation Evaluation: For the sake of space limitations, we did not use tables/figures to display the results from cross-validation evaluation (leave one out split). The purpose of cross-validation evaluation is to test the ability of data-driven methods to generalize.

Unsurprisingly, with less data, supervised learning methods performed worse. For cross-validation, the best-performing supervised learning method was TCN. After applying TCN, 75.33% of data fell within 2 bpm of the ground truth. Comparatively, applying "mean" had 86.88% of the data fall within 2 bpm of ground truth. When comparing by similarity to ground truth through the metric of DTW, we find that the best performing supervised learning method, TCN, has an average distance of 0.087 from the ground truth, barely better than the average DTW distance for Noisy Output data which is 0.091. Alternatively, mean, brings the average DTW down to 0.063 and the worst performing non-learning method, Kalman Filter, brings DTW down to 0.072.

C. Heart Rate

1) Global Evaluation: Compared to respiratory rate, the pre-post processing estimation for heart rate was considerably less accurate, although its general trend follows ground truth, only 18.43% of Noisy Output data is actually within 5 bpm of ground truth. As Figure 2 shows, the difference between noisy data and ground truth data for heart rate is significant.

Results in Table II show that the supervised learning methods performed significantly better than non-learning based methods by both increasing the similarity to ground truth (reducing DTW) as well as increasing salvageable information (as defined by percentage of data within 5 bpm of ground truth). Although there remains a lot of data not within 5 bpm after post processing, the supervised learning methods almost double the amount of data within 5 bpm of the ground truth, with Random Forest and TCN having almost 50% of the data within 5 bpm of ground truth. GBR, SVR, and MultiRocket show similar levels of performance while other supervised learning methods performed noticeably worse. In terms of absolute error, these top performing methods also perform similarly as shown in Figure 4 although TCN does show slightly more promise than the other methods.

The analysis of the method results using DTW indicate that Random Forest, GBR, SVR, and TCN provide results most similar to the ground truth. Noticeably, after those four results, the next lowest DTW score belongs to LSTM FCN which is almost 50% farther away suggesting that the four earlier indicated methods are significantly better at capturing trend for heart rate. From this analysis, it is evident that, given sufficient data, TCNs are the best choice for processing complex signals such as heart rate. As illustrated in Figure 3, the methods are able to yield results that are significantly more consistent with the ground truth, despite their inability to fully capture the entire ground truth. TCNs demonstrate the ability to effectively preserve a significant amount of information while accurately capturing the overall trends of the ground truth.

2) Cross-validation Evaluation: In contrast to a global split, a cross-validation evaluation approach could provide clearer insights into method performance when trained on limited data and data originating from distinct sources. The change in data available to train on significantly decreases the quality of supervised learning methods, which may affect how the top performing methods for global split heart rate perform.

Compared to almost 50% of data within 5 bpm, the best performing supervised learning method, Random Forest, was only able to obtain 36.1% in 5 bpm of ground truth. The best performing overall method was Kalman filter which obtained 36.3% in 5 bpm of ground truth. When examining at DWT, we find Kalman Filter exhibits the highest DTW distance, measured at 3.05. This contrasts significantly with the average DTW distance between the ground truth and the pipeline data for each sensor, which stands at 0.64. For DTW, the best performing methods were GBR, TCN, and Random Forest which had DTW distances of 0.332, 0.312, and 0.319 respectively, almost half the DTW distance of the base Noisy Output data. This suggests that despite being unable to salvage a lot of information, these methods were able to learn to capture the trend of the ground truth.

This analysis indicates that while TCNs and other supervised learning techniques are capable of capturing the complex trends inherent in heart rate data, they may not be effective at obtaining the actual heart rate unless they are properly trained on enough data with similar characteristics.

D. Computing Resources

As presented in Table II, our four simple non-learning denoising methods exhibited extremely low execution times, completing post-processing for a single instance on the order of microseconds. SVR takes exceptionally longer for inference as it does not scale well for large datasets which would make it non optimal for real world calculations. The experiments are processed on servers equipped with two NVIDIA RTX 4090 GPUs and an AMD Ryzen Threadripper PRO 5995WX.

VI. Discussion

In this study, we examine the effectiveness of various post-processing methods on salvaging vital signs estimation from an ultrawideband RF device. Respiration rate, due to its significant effect on chest displacement, can be reliably detected and estimated from RF signals. Although the baseline RF pipeline can generate a reasonable estimate of the respiration rate, simple denoising methods, especially moving average, further led to notable performance improvements.

For heart rate estimation, we observed that supervised learning methods Random Forest, Gradient Boosting Regressor, and Temporal Convolutional Network, effectively captured trends and recovered a substantial portion of the disrupted data, with Temporal Convolutional Network demonstrating slightly superior performance compared to the others.

However, with limited training data, these methods struggled to recover the majority of data when tested on a different set, highlighting the dependency of these methods on a large and diverse dataset.

In the future, we will continue to evaluate these methods as we collect increasingly diverse datasets given the potential performance gains achieved with more data. Moreover, we will experiment with various feature extraction techniques prior to post-processing to further improve the performance. Additionally, we will investigate the trade-offs between computational efficiency and inference latency to optimize the method for practical deployment, where real-time performance is critical.

VII. Conclusion

In this paper, we examine the impact of various post-processing methods on post-pipeline RF data. Performance evaluation is conducted using multiple training and testing methods using data collected from a real world environment. Our analysis reveals that despite the significantly more complex nature of heart rate signals, Random Forest and TCN models, which achieve comparable performance, were able to recover an additional 20% more of the total data compared to the noisy output prior to post-processing. These methods recovered 10% more than all non-supervised learning methods and even some of the other representative Neural Network models. For respiration rate estimation, the moving average method demonstrates superior performance, surpassing more complex post-processing methods, recovering slightly over 20% more of the total data than the noisy output. When utilizing data from external sources or COTS RF devices, improving result accuracy remains a critical consideration. This paper offers practical recommendations for enhancing the performance of out-of-box RF vital signs monitoring in scenarios where access to raw data is unavailable.

References

[1] Z. Xie, B. Zhou, X. Cheng, E. Schoenfeld, and F. Ye, "Vitalhub: Robust, non-touch multi-user vital signs monitoring using depth camera-aided uwb," in *2021 IEEE 9th ICHI*. IEEE, 2021, pp. 320–329.

[2] C. Li, V. M. Lubecke, O. Boric-Lubecke, and J. Lin, "A review on recent advances in doppler radar sensors for noncontact healthcare monitoring," *IEEE Transactions on microwave theory and techniques*, vol. 61, no. 5, pp. 2046–2060, 2013.

[3] Z. Xie, B. Zhou, and F. Ye, "Signal quality detection towards practical non-touch vital sign monitoring," in *Proceedings of the 12th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, 2021, pp. 1–9.

[4] Z. Xie, H. Wang, S. Han, E. Schoenfeld, and F. Ye, "Deepvs: A deep learning approach for rf-based vital signs sensing," in *Proceedings of the 13th ACM international conference on bioinformatics, computational biology and health informatics*, 2022, pp. 1–5.

[5] Z. Chen, T. Zheng, C. Cai, and J. Luo, "Movi-fi: motion-robust vital signs waveform recovery via deep interpreted rf sensing," in *MobiCom '21*. ACM, 2021.

[6] C. Li and J. Lin, "Random body movement cancellation in doppler radar vital sign detection," *IEEE Transactions on Microwave Theory and Techniques*, vol. 56, no. 12, pp. 3143–3152, 2008.

[7] A. D. Droitcour, O. Boric-Lubecke, V. M. Lubecke, J. Lin, and G. T. Kovacs, "Range correlation and i/q performance benefits in single-chip silicon doppler radars for noncontact cardiopulmonary monitoring," *IEEE Transactions on Microwave Theory and Techniques*, vol. 52, no. 3, pp. 838–848, 2004.

[8] Z. Xie, A. Nederlander, I. Park, and F. Ye, "Short: Rf-q: Unsupervised signal quality assessment for robust rf-based respiration monitoring," in *Proceedings of the 8th ACM/IEEE International Conference on Connected Health: Applications, Systems and Engineering Technologies*, 2023, pp. 158–162.

[9] S. Choi and H. Yoon, "Convolutional neural networks for the real-time monitoring of vital signs based on impulse radio ultrawide-band radar during sleep," 2023.

[10] B. Berlo, A. Elkelaney, T. Ozcelebi, and N. Meratina, "Millimeter wave sensing: A review of application pipelines and building blocks," in *IEEE Sensors Journal, Volume 21, Issue 9*. IEEE, 2021.

[11] Novelda. [Online]. Available: https://novelda.com/technology/datasheets/

[12] Masimo. [Online]. Available: https://www.masimo.com/mightysat

[13] J. Friedman, "Greedy function approximation: A gradient boosting machine," in *The Annals of Statistics, Vol. 29, No. 5, 2001.*, 2001.

[14] L. Breiman, "Random forests," in *Machine Learning, Volume 45, pages 5–32, (2001)*. Springer, 2001.

[15] C. Lea, R. Vidal, A. Reiter, and G. Hager, "Temporal convolutional networks: A unified approach to action segmentation." arxiv, 2016.

[16] F. Karim, S. Majumdar, H. Darabi, and S. Chen, "Lstm fully convolutional networks for time series classification." arxiv, 2017.

[17] C. Tan, A. Dempster, C. Bermier, and G. Web, "Multirocket: Multiple pooling operators and transformations for fast and effective time series classification." arxiv, 2021.

[18] K. He, X. Zhang, X. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification." arxiv, 2015.

[19] scikit. [Online]. Available: https://scikit-learn.org/1.5/api/sklearn.ensemble.html

[20] sktime. [Online]. Available: https://www.sktime.net/en/stable/api_reference/regression.html

[21] pypi. [Online]. Available: https://pypi.org/project/keras-tcn/2.9.3/