High-Resolution Range-Doppler Imaging from One-Bit PMCW Radar via Generative Adversarial Networks

Jingxian Wang
FAU Erlangen-Nürnberg
Erlangen, Germany
jingxian.wang@fau.de

Moritz Kahlert
HELLA GmbH & Co. KGaA
Lippstadt, Germany
moritz.kahlert@forvia.com

Tai Fei
University of Applied Sciences and Arts
Dortmund, Germany
tai.fei@fh-dortmund.de

Changxu Zhang
HELLA GmbH & Co. KGaA
Lippstadt, Germany
changxu.zhang@forvia.com

Zhaoze Wang
HELLA GmbH & Co. KGaA
Lippstadt, Germany
zhaoze.wang@forvia.com

Markus Gardill
Brandenburg University of Technology
Cottbus, Germany
markus.gardill@b-tu.de

***Abstract*—Digital modulation schemes such as phase-modulated continuous wave (PMCW) have recently attracted increasing attention as possible replacements for frequency-modulated continuous wave (FMCW) modulation in future automotive radar systems. A significant obstacle to their widespread adoption is the expensive and power-consuming analog-to-digital converters (ADCs) required at gigahertz frequencies. To mitigate these challenges, employing low-resolution ADCs, such as one-bit, has been suggested. Nonetheless, using one-bit sampling results in the loss of essential information. This study explores two range-Doppler (RD) imaging methods in PMCW radar systems utilizing neural networks (NNs). The first method merges standard RD signal processing with a generative adversarial network (GAN), whereas the second method uses an end-to-end (E2E) strategy in which traditional signal processing is substituted with an NN-based RD module. The findings indicate that these methods can substantially improve the probability of detecting targets in the range-Doppler domain.**

***Index Terms*—Analog-digital-conversion, one-bit sampling, phase-modulated continuous wave, quantization, range-Doppler processing**

I. Introduction

Digital modulation schemes, such as phase-modulated continuous wave (PMCW), have recently attracted attention due to their robustness against mutual interference and their inherent multiplexing capability, which is essential for multiple-input multiple-output (MIMO) systems [1]. However, PMCW and other digital modulation schemes require fast-sampling analog-to-digital converters (ADCs) to process the entire baseband, reaching up to 1 GHz at 77 GHz and 4 GHz at 79 GHz. These high-sampling ADCs are power-consuming and generate significantly larger data volumes than frequency-modulated continuous wave (FMCW) radar systems, which only need to sample a narrowband beat signal [2]. To address this issue, with a focus on waveform design and signal processing, stepped frequency solutions have been proposed for PMCW [3], providing a tradeoff between performance and hardware requirements. The carrier frequency changes linearly over time to reduce the bandwidth of each pulse, which in turn lowers the ADC sampling requirements. An alternative approach to reducing sampling rates involves lowering the resolution of the ADCs [4] and, for example, employing them in mixed-ADC setups, as shown in [5]. Low-resolution ADCs, e.g., one-bit ADCs, could be a promising solution to reduce the amount of data, power consumption, and costs. However, amplitude information is lost by one-bit sampling, which provides crucial information for target detection. In [6], the authors focused on designing the transmit code and receive filter in the presence of one-bit sampling.

This study investigates the effectiveness of neural network (NN)-based methods, specifically end-to-end (E2E) and hybrid approaches, to improve target detectability when applying one-bit quantization. Our contributions include the development of an E2E approach for range-Doppler (RD) map generation based on high-resolution and one-bit quantized ADC data and the introduction of a denoising network to mitigate noise introduced by quantization. Furthermore, we propose a hybrid approach combining conventional RD processing with a denoising network. Finally, we provide a comparative performance evaluation of the E2E and hybrid approaches, highlighting their advantages.

II. Signal Model

Let $\mathbf{x} = [x_0, \ldots, x_{N-1}] = \exp(\jmath\boldsymbol{\phi})$ with $\boldsymbol{\phi} \in \{0, \pi\}^N$ represent a pseudo-noise (PN) binary sequence of length N, where $x_n \in \{-1, 1\}$ is denoted as a chip. The transmitted and received signals in their equivalent complex baseband (ECB) representation can be expressed by x_{BB} and y_{BB}, respectively,

$$x_{\mathrm{BB}}(t) = \sum_{n=0}^{N-1} x_n \mathrm{rect}\left(\frac{t-(n+0.5)T}{T}\right), \qquad (1)$$

and

$$y_{\mathrm{BB}}(t) = \sum_{k=0}^{K-1} \gamma_k \, x_{\mathrm{BB}}\big(t - \tau_k(t)\big) \exp\big(-\jmath 2\pi f_{\mathrm{c}} \tau_k(t)\big), \quad (2)$$

where T is the chip duration, $\mathrm{rect}(\cdot)$ denotes the rectangular function, K is the number of point targets, γ_k is the scaling factor of the kth target reflection, including free-space attenuation, radar cross-section (RCS), and the reflection phase, $\tau_k(t)$ is the round-trip delay, and f_{c} is the carrier frequency. The round-trip delay can be expressed by $\tau_k(t) = 2r_{0,k}/c_0 + 2v_{\mathrm{r},k}t/c_0$, where $r_{0,k}$ is the range between the radar and the kth target at the beginning of the coherent processing interval (CPI), $v_{\mathrm{r},k}$ is the relative velocity, and c_0 is the light speed. In addition, it is assumed that M sequences with sequence duration T_{seq} are transmitted in a single CPI, and $y_{\mathrm{BB}}(t)$ is sampled with a rate of $1/T$ at time steps $t_{\mathrm{s}} = nT + mT_{\mathrm{seq}}$ with $0 \le n < N, 0 \le m < M$, resulting in the two-dimensional matrix representation of the sampled baseband signal expressed by $\mathbf{Y} = y_{\mathrm{BB}}(t_{\mathrm{s}}) \in \mathbb{C}^{N \times M}$.

Further, it is assumed that the output of the ADCs can be either high-resolution or one-bit. The output of the ADCs after one-bit quantization can be expressed by $\tilde{\mathbf{Y}} = \mathcal{Q}(\mathbf{Y}) = \mathrm{sign}\big(\Re(\mathbf{Y})\big) + \jmath\,\mathrm{sign}\big(\Im(\mathbf{Y})\big)$, where $\mathcal{Q}(\cdot)$ denotes the complex quantization operator, $\mathrm{sign}(x)$ is the sign function, and $\Re(\cdot)$ and $\Im(\cdot)$ denote the real and imaginary parts, respectively. We define that $\mathrm{sign}(x) = 1$ for $x \ge 0$, otherwise $\mathrm{sign}(x) = -1$.

To retrieve the range and Doppler information from the baseband signal, we first apply cross-correlation along the fast-time domain of $\mathbf{Y}$ as follows,

$$p_{rm} = \sum_{n=0}^{N-1} x^*_{\mathrm{mod}\,(n-r,N)} y_{nm}, \quad (3)$$

where (y_{nm}) is the (r,m)th element in $\mathbf{Y}$, $\mathbf{P} = (p_{rm}) \in \mathbb{C}^{N \times M}$ is the range profile, r denotes the range bin index, $(\cdot)^*$ is the complex conjugate, and $\mathrm{mod}\,(\cdot)$ is the modulo operator. Subsequently, the relative velocities (i.e., Doppler shifts) can be extracted by applying discrete Fourier transforms (DFTs) along the slow-time domain, resulting in the RD map $\mathbf{Q} = (q_{rv}) \in \mathbb{C}^{N \times M}$. The (r,v)th element can be expressed by

$$q_{rv} = \sum_{m=0}^{M-1} p_{rm} \exp\left(-\jmath 2\pi v \frac{m}{M}\right). \quad (4)$$

Note that by replacing $\mathbf{Y}$ with $\tilde{\mathbf{Y}} = (\tilde{y}_{nm})$ in (3), $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{S}}$ can be calculated similarly without losing generality.

III. Range-Doppler Neural Networks

This work presents two NN-based approaches for generating and denoising RD maps from 1-bit ADC data: an E2E approach and a hybrid approach. Both networks are adversarially trained within a generative adversarial network (GAN) framework. The E2E method integrates convolutional neural networks (CNNs) [7], residual neural networks (ResNets) [8], and frequency domain operations. The hybrid approach combines conventional signal processing techniques with a denoising NN to improve performance.

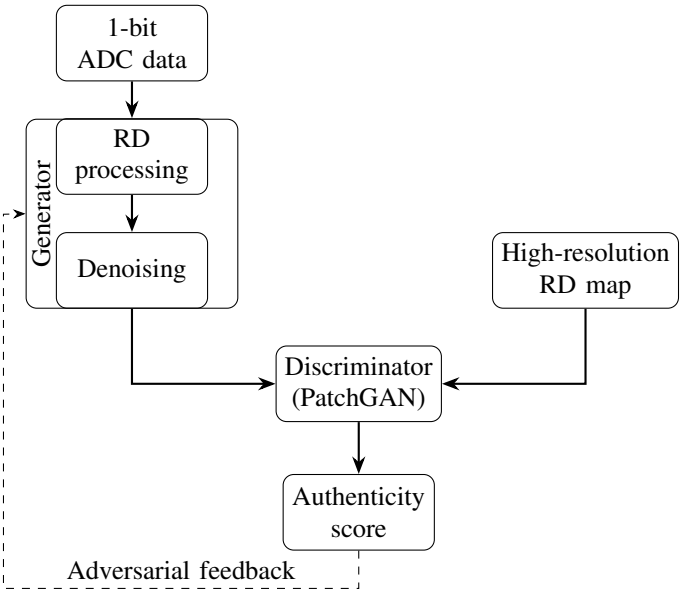

Fig. 1. Overall neural network architecture for range-Doppler map generation and denoising based on 1-bit ADC data.

A. Network Architecture

In our GAN framework (Fig. 1), the generator processes 1-bit ADC data to reconstruct high-quality RD maps, while the discriminator, implemented as a PatchGAN, assesses their authenticity by comparing them to high-resolution (HR) RD maps from a full-precision ADC. The HR RD map serves as a reference, enabling the discriminator to refine the generator iteratively by providing authenticity feedback, thereby improving the fidelity of the generated outputs.

In the E2E approach, all generator layers are trainable. In contrast, in the hybrid approach, the initial layers are replaced with classical RD processing, as described in (3) and (4), while the remaining layers, i.e., the backbone network, focus on noise reduction.

B. Generator in E2E Approach

To design the generator in the E2E approach, we adopt the idea from [9] and integrate a similar learnable signal processing layer into the backbone network, which is the Pix2Pix generator from [10], as illustrated in Fig. 2. This modified Pix2Pix generator is responsible for both RD processing and noise reduction to reconstruct high-quality RD maps.

The learnable signal processing layer in [9] was originally designed for FMCW radar, where range estimation is performed through DFT. However, in PMCW radar, the range profile is obtained by correlation, as described in (3). To accommodate this difference, we replace the DFT-based range estimation with a frequency-domain correlation operation. Specifically, we employ a fixed-length, trainable PN sequence as the correlation kernel, which is optimized during training to improve the correlation response to target echoes, thereby improving the accuracy of the range profile. Unlike the hybrid

TABLE I
GENERATOR ARCHITECTURE AND ITS COMPONENTS FUNCTIONS.

Component	Layer Type	Main Function
RD processing	Correlations + DFTs	Extracts range and Doppler information
Initial Conv	3×3 Conv + ReLU	Feature extraction
Encoder	4×4 Conv + batch normalization (BN) + ReLU + residual block (×4)	Downsampling and feature enhancement
Residual block	3×3 Conv + BN + ReLU + skip connection	Mitigates vanishing gradient and enables deeper network training
Bottleneck	3 residual blocks	Deep feature transformation
Decoder	4×4 transposed Conv + BN + ReLU + residual block (×4)	Upsampling path for reconstruction
Global residual	1×1 Conv	Directly connects input to output
Final output	3×3 Conv + Tanh	Generates final RD map

TABLE II
DISCRIMINATOR ARCHITECTURE AND ITS COMPONENTS FUNCTIONS.

Component	Layer Type	Main Function
Input	Concatenation of RD map pairs	Distinguish real vs. generated data
Downsampling	$3 \times (4 \times 4$ Conv + Leaky ReLU)	Hierarchical feature extraction
Fully Connected 1	4×4 Conv + Leaky ReLU	Compresses feature representation
Fully Connected 2	4×4 Conv + Sigmoid	Outputs authenticity score

approach, where the PN sequence is predefined and fixed *a priori*, the E2E approach enables data-driven learning of the PN sequence. Additionally, analogous to (4), Doppler processing is implicitly learned within the network since both the Doppler window and the complex DFT kernel are trainable parameters. These modifications improve both accuracy and robustness in the RD processing.

The Pix2Pix generator, originally based on a U-Net [11] architecture, employs skip connections to transfer low-level spatial information from the encoder to the decoder. However, applying Pix2Pix directly to 1-bit ADC data leads to training instability and suboptimal reconstruction, as radar signals exhibit characteristics distinct from natural images. To mitigate these issues, we introduce three key modifications. First, the network depth is increased with a bottleneck layer to enhance feature extraction and stabilize training, improving gradient flow and preventing vanishing gradients. Second, residual learning is incorporated, where, instead of relying solely on U-Net's skip connections, multiple residual blocks [8] are integrated within both the encoder and decoder, improving gradient propagation and ensuring stable training. Third, a global residual connection is introduced using a 1×1 convolution (Conv), directly linking the input to the output to preserve low-frequency components in the final RD map. These enhancements improve stability, robustness, and accuracy, enabling effective training on 1-bit ADC data and high-fidelity RD map reconstruction. For the simulations in Sec. IV, a learning rate of $2 \cdot 10^{-4}$ is employed.

C. Discriminator

The discriminator distinguishes between reference, specifically HR, and generated RD maps, providing adversarial feedback to the generator and motivating it to produce more authentic output. The design utilizes a PatchGAN discriminator [10], focusing on classifying localized regions instead of whole images. This approach enables the discriminator to discern intricate details. The components of the discriminator are outlined as follows. Firstly, feature extraction is performed using three 4×4 Conv layers with a stride of 2 for downsampling. Each layer is succeeded by Leaky rectified linear unit (ReLU) activation to derive hierarchical features from the input data. Secondly, PatchGAN employs a patch-based discrimination approach that evaluates smaller sections rather than making an overarching decision on the full RD map, thereby improving the retention of detailed features. In the third and last discriminative stage, the concluding pair of layers utilizes fully connected Convs to generate a singular scalar output representing the probability of an authentic input RD map. For the simulations in Sec. IV, a learning rate of $1 \cdot 10^{-4}$ is employed.

Using a PatchGAN enhances the ability of the discriminator to help the generator preserve intricate structural information, thereby improving the reconstruction performance. The structure of the discriminator network is outlined in Table II.

D. Loss Functions and Training Strategy

Different loss functions guide the training of the generator and discriminator to ensure stable adversarial learning and high-quality reconstruction.

1) Generator Loss: The loss function of the generator mainly comprises three key components, which are as follows. Firstly, the *L1 loss* denoted as $\mathcal{L}_{\mathrm{L1}}$, which measures the absolute difference between the generated RD map and the HR RD map, ensuring structural consistency. Secondly, the *structural similarity index (SSIM) loss* [12], denoted as $\mathcal{L}_{\mathrm{SSIM}}$, encourages perceptual similarity by preserving structural information. Thirdly, the *adversarial loss* [13] denoted as $\mathcal{L}_{\mathrm{GAN}}$, which assigns an authenticity score to the generated RD map. This loss function ensures that the generator learns to produce RD maps indistinguishable from real data by maximizing the output of the discriminator, thus fooling the discriminator into classifying the generated samples as real. The total loss of the discriminator is formulated as $\mathcal{L}_{\mathrm{G}} = \lambda_{\mathrm{L1}}\mathcal{L}_{\mathrm{L1}} + \lambda_{\mathrm{SSIM}}\mathcal{L}_{\mathrm{SSIM}} + \mathcal{L}_{\mathrm{GAN}}$. The hyperparameters λ_{L1} and λ_{SSIM} were set to reasonable values within the range of 1 to 50, and minor variations did not significantly affect the final results.

2) Discriminator Loss: As proposed in [13], the discriminator is trained using the Wasserstein loss with gradient penalty, mainly comprising two components. Firstly, the *Wasserstein distance*, denoted as $\mathcal{L}_{\mathrm{W}}$, measures the discrepancy between

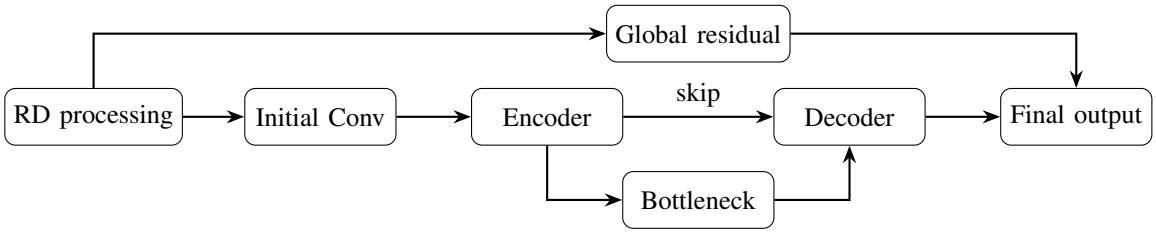

Fig. 2. Structure of the proposed generator network. The functionalities of the blocks are outlined in Table I.

the distributions of real and generated samples. Secondly, a *gradient penalty (GP)* loss, denoted as $\mathcal{L}_{\mathrm{GP}}$, enforces the Lipschitz constraint and regularizes the discriminator gradient norm to prevent vanishing or exploding gradients. The total loss of the discriminator is formulated as $\mathcal{L}_{\mathrm{D}} = -\mathcal{L}_{\mathrm{W}} + \lambda_{\mathrm{GP}}\mathcal{L}_{\mathrm{GP}}$. The gradient penalty weight is chosen as $\lambda_{\mathrm{GP}} = 10$, according to the configuration described in [13].

IV. Results

A. Data Set

To train and evaluate the models, a synthetic data set is generated using the PMCW signal model described in Section II. The data set contains 6000 two-dimensional matrices, including 1-bit ADC data and HR RD maps utilized as ground truth (GT). The HR RD maps are derived from unquantized ADC data. Moreover, the dataset is divided such that $80\,\%$ serves as training data, while $20\,\%$ is used for validation. The signal processing takes 1-bit ADC data as input, consisting of 3000 matrices divided into 1500 matrices with a signal-to-noise ratio (SNR) of $10\,\mathrm{dB}$ and 1500 matrices with an SNR of $20\,\mathrm{dB}$. For each input matrix, we generate 3000 HR RD maps from full-resolution ADC data with an SNR of $50\,\mathrm{dB}$. These HR RD maps serve as a reference to supervise model training. The *Adam* optimizer is utilized for the training, configured with the parameters $\beta_1 = 0.5$ and $\beta_2 = 0.999$ and the global batch size is set to 16.

The data set is generated using a simulated single-input single-output (SISO) $79\,\mathrm{GHz}$ automotive radar employing a maximum length sequence (MLS) of length 127. An additional chip is added to each sequence to improve the usability of NNs, resulting in a total length of $N = 128$. The chip duration and bandwidth are $10\,\mathrm{ns}$ and $100\,\mathrm{MHz}$, respectively, with a total of 10240 pulses transmitted. To improve the SNR, we apply the accumulation approach described in [14]. By accumulating each set of 20 pulses to generate a range profile, the total number of slow-time samples is reduced to $M = 512$.

B. Metrics

Performance is assessed using the following metrics. First, the mean squared error (MSE) evaluates the differences between each cell of the 1-bit ADC data and the original HR RD maps, which are denoted by $\mathbf{Q}^{\mathrm{1b}} = \left(q_{rv}^{\mathrm{1b}}\right)$ and $\mathbf{Q}^{\mathrm{HR}} = \left(q_{rv}^{\mathrm{HR}}\right)$, respectively, as expressed by

$$\mathrm{MSE}\left(\mathbf{Q}^{\mathrm{1b}}, \mathbf{Q}^{\mathrm{HR}}\right) = \frac{1}{NM}\sum_{v=0}^{M-1}\sum_{r=0}^{N-1}\left(\left|q_{rv}^{\mathrm{1b}}\right| - \left|q_{rv}^{\mathrm{HR}}\right|\right)^2, \quad (5)$$

TABLE III
Comparison of evaluation metrics in validation scenario.

Method	**MSE**	**PSL** (dB)	**ISL** (dB)
HR ADC + RD	0	-16.1	-5.3
1-bit ADC + RD	$4.74e^{-4}$	-18.2	-4.6
1-bit ADC + RD + Denoise-NN	$8.00e^{-6}$	-13.5	-5.6
1-bit ADC + (RD+Denoise)-NN	$1.00e^{-5}$	-15.5	-4.7

where a lower MSE indicates better reconstruction quality. Second, the *peak-sidelobe level (PSL)* measures the maximum sidelobe level expressed by

$$\mathrm{PSL}_v(\mathbf{Q}) = 20\log_{10}\left(\max_{r\neq\hat{r}}|q_{rv}|\right), \quad (6)$$

where v is the target Doppler index in RD, and $\hat{r}$ is the range index of the main lobe peak. Lower PSL indicates better sidelobe suppression and, hence, better target detectability. Third, the *integrated sidelobe level (ISL)* measures the total sidelobes expressed by

$$\mathrm{ISL}_v(\mathbf{Q}) = 20\log_{10}\left(\sum_{r=0,r\neq\hat{r}}^{N-1}|q_{rv}|^2\right), \quad (7)$$

where lower ISL values indicate better sidelobe suppression and, similar to PSL, improved target detectability.

C. Evaluation Based on Validation Scenario

As illustrated in Table I, the generator employs an E2E approach to transform 1-bit ADC data into HR RD maps, implicitly fulfilling two tasks jointly: RD processing and denoising. This approach is hereinafter referred to as *1-bit ADC + (RD+Denoise)-NN*. In contrast, using the hybrid approach, the generator consists only of the denoising module, further referred to as *1-bit ADC + RD + Denoise-NN*. The conversion of ADC data to RD maps, as detailed in Section II, occurs before the application of the denoising NN.

As shown in Table III, each method exhibits distinct strengths in terms of MSE, PSL, and ISL. As detailed in Section IV-B, the MSE is determined considering all cells within the final RD maps, whereas PSL and ISL are explicitly computed for the velocity bin depicted in Fig. 3. Thus, the MSE serves as a more comprehensive metric, while PSL and ISL are localized metrics tailored to individual target velocity bins. The conventional *1-bit ADC + RD* approach achieves the best PSL in the given scenario; however, its total

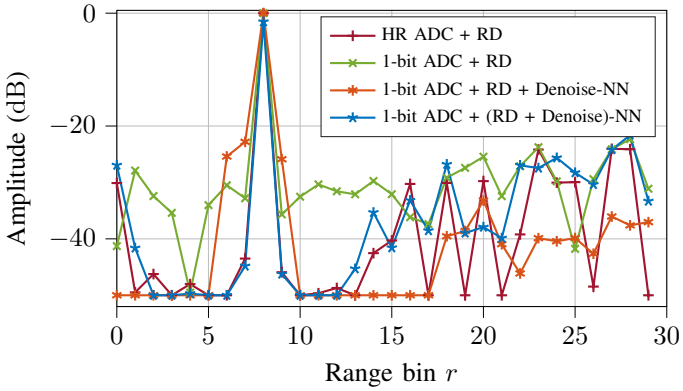

Fig. 3. Slices of range-Doppler map for a validation scenario with a single target. The curves for the different approaches have been normalized to peak power.

MSE exceeds that of the NN-based methods, indicating that while it lowers peak sidelobes, the noise level is elevated due to quantization effects. Moreover, its ISL is greater than for the other methods. When integrating 1-bit ADC and RD processing with a denoising-NN in a hybrid approach (*1-bit ADC + RD + Denoise-NN*), there is a reduction in the MSE and an improvement in the ISL. However, this combination results in a smoothing effect during suppression of the peak sidelobes. While background noise is reduced, the relative amplitudes of certain sidelobes increase, causing an increase in the PSL. The loss function mainly focuses on minimizing pixel errors and lacks explicit conditions regarding peak sidelobes so that certain sidelobe levels may rise. Future research might consider incorporating the PSL and ISL metrics within the loss function. Integrating RD processing and denoising into an E2E NN (*1-bit ADC + (RD+Denoise)-NN*) leads to a more balanced performance in all metrics. The MSEs are calculated relative to the *HR ADC + RD* method, where the unquantized ADC data undergo conventional RD processing. This shows that the E2E generator module can efficiently extract range and Doppler information from the ADC data, illustrating the adaptability of the developed GAN. It can be inferred that a denoising NN on quantized ADC data can mitigate quantization impacts and improve target detectability.

D. Computational Complexity and Resource Consumption

The hybrid approach slightly reduces the floating-point operations per second (FLOPS), and compared with the E2E approach, its inference speed is greatly improved, decreasing from 3.1 s to 1.7 s. Experimental results based on *Nvidia GeForce GTX 1080 Ti* show that the hybrid approach requires an average of 2800 s per epoch during training, while the E2E approach takes about 6000 s per epoch. These results demonstrate that the hybrid approach accelerates inference and significantly shortens training time, thereby enhancing its appeal for real-time applications.

V. Conclusion and Outlook

Although the E2E generator slightly improves reconstruction quality, the experimental results indicate that conventional RD processing can achieve comparable performance when combined with a subsequent denoising NN. In addition, conventional processing benefits from reduced training time and lower computational complexity. Given the constraints on real-time processing and resource limitations, the conventional approach is more attractive for practical applications.

To prove effectiveness under various conditions, networks should be trained using various scenarios, including more targets, target RCSs, and SNRs. Furthermore, the performance of the model should be assessed by applying it to actual radar systems.

References

[1] C. Waldschmidt, J. Hasch, and W. Menzel, "Automotive Radar — From First Efforts to Future Systems," *IEEE Journal of Microwaves*, vol. 1, no. 1, pp. 135–148, 2021.

[2] V. Winkler, "Range Doppler detection for automotive FMCW radars," in *2007 European Radar Conference*, 2007, pp. 166–169.

[3] M. Kahlert, T. Fei, C. Tebruegge, and M. Gardill, "Stepped-Frequency PMCW Waveforms for Automotive Radar Applications," *IEEE Transactions on Radar Systems*, vol. 3, pp. 233–245, 2025.

[4] C.-Y. Wu, T. Zhang, J. Li, and T. F. Wong, "Parameter Estimation in PMCW MIMO Radar Systems With Few-Bit Quantized Observations," *IEEE Transactions on Signal Processing*, vol. 70, pp. 810–821, 2022.

[5] M. Kahlert, L. Xu, T. Fei, M. Gardill, and S. Sun, "High-Resolution DOA Estimation Using Single-Snapshot Music for Automotive Radar with Mixed-ADC Allocations," in *2024 IEEE 13th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, 2024, pp. 1–5.

[6] F. Foroozmehr, M. Modarres-Hashemi, and M. M. Naghsh, "One-Bit PMCW Radar: Designing Binary Transmit Code and Receive Filter via a Worst-Case Approach," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 12, pp. 19 774–19 779, 2024.

[7] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.

[8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[9] B. Yang, I. Khatri, M. Happold, and C. Chen, "ADCNet: Learning from Raw Radar Data via Distillation," 2023. [Online]. Available: https://arxiv.org/abs/2303.11420

[10] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5967–5976.

[11] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.

[12] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[13] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of wasserstein GANs," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 5769–5779.

[14] A. Bourdoux, U. Ahmad, D. Guermandi, S. Brebels, A. Dewilde, and W. Van Thillo, "PMCW waveform and MIMO technique for a 79 GHz CMOS automotive radar," in *2016 IEEE Radar Conference (RadarConf)*, 2016, pp. 1–5.