ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 979-8-3503-6874-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICASSP49660.2025.10890731

Advancing High-Resolution and Efficient Automotive Radar Imaging through Domain-Informed 1D Deep Learning

Ruxin Zheng[†], *Shunqiao Sun*[†], *Hongshan Liu*[†], *Holger Caesar*[‡], *Honglei Chen*[§], *Jian Li*[¶]

[†]The University of Alabama, Tuscaloosa, AL, USA
[‡]Delft University of Technology, Delft, The Netherlands
[§]Mathworks, Inc, Natick, MA, USA
[¶]University of Florida, Gainesville, FL, USA

ABSTRACT

Millimeter-wave (mmWave) radars are critical for autonomous vehicles' perception tasks, offering reliable performance in adverse weather conditions. However, their application is often hindered by insufficient spatial resolution for detailed semantic scene interpretation. Traditional super-resolution methods derived from optical imaging fail to accommodate the unique properties of radar signals. Addressing this, our study redefines radar imaging super-resolution as a one-dimensional (1D) signal super-resolution spectra estimation problem, leveraging domain-specific insights to innovate data normalization and introduce a domain-informed signal-to-noise ratio (SNR)-guided loss function. Our custom deep learning network, tailored for automotive radar imaging, achieves substantial improvements in parameter efficiency, and inference speed while enhancing image quality and resolution. Comprehensive tests demonstrate that our SR-SPECNet establishes a new standard for high-resolution radar range-azimuth imaging, surpassing previous methods. Source code and new radar dataset will be made publicly available at https://github.com/ruxinzh/SR_DOA.

***Index Terms*—** Automotive Radar, High-Resolution Radar Imaging, Real-World Dataset, Direction-of-Arrival Estimation, Deep Learning

I. INTRODUCTION

Radar technology, especially millimeter wave radars, is pivotal for advanced driver assistance systems (ADAS) and autonomous vehicles, excelling over traditional RGB cameras and LiDAR in adverse weather and low visibility [1]–[6]. It provides robust, cost-effective, and reliable sensing across almost all environmental conditions. Within these systems, Frequency-modulated continuous-wave (FMCW) signals in the millimeter-wave band are preferred for their affordability and high-resolution capabilities, supporting crucial functions like free space detection, 360° sensing, object detection, classification, and SLAM [7]–[9]. Despite their strengths, these radars typically have limited azimuth angular resolution, a gap bridged by employing multiple-input multiple-output (MIMO) radar technology. MIMO radars use a synthetic large virtual array aperture to significantly enhance angular resolution, efficiently balancing the number of transmit and receive antennas [3], [10]–[12].

Advanced signal processing techniques are being developed to enhance angular resolution beyond the capabilities of digital beamforming, typically implemented via fast Fourier transform (FFT). Super-resolution direction of arrival (DOA) estimation algorithms, such as compressive sensing (CS) [13]–[15] and the iterative adaptive approach (IAA) [16], [17], offer substantial improvements. However, their high computational demands pose significant challenges for real-time applications in automotive contexts. Deep learning (DL) techniques have recently shown promising results in image super-resolution within computer vision [18]–[21]. Adapting these methods to enhance azimuth resolution in range azimuth (RA) heatmaps could lead to major advancements. Nonetheless, few studies have explored the generation of super-resolution RA heatmaps directly from raw radar signals, incorporating crucial radar domain knowledge. Simplistic image-to-image or volume-to-volume approaches often neglect essential signal processing insights, resulting in data-heavy models that may not meet the demands for rapid inference and compact design critical in automotive radar applications, where on-chip implementation is key.

Research on super-resolution RA heatmap generation for automotive radar is nascent, with most studies depending on FFT-generated ground truths from expansive antenna arrays. To date, no methods have utilized RA heatmaps from super-resolution algorithms as ground truths. This paper addresses these gaps by introducing the Super-Resolution Angular Spectra Estimation Network (SR-SPECNet). SR-SPECNet, designed with radar signal processing expertise, redefines RA heatmap enhancement as a one dimensional (1D) azimuth super-resolution task. It incorporates a novel data normalization method and an SNR-guided loss function to advance angular spectra generation. Thoroughly tested across various training dataset sizes with a dedicated real-world dataset, SR-SPECNet demonstrates exceptional parameter efficiency and superior imaging quality.

II. RADAR DATASETS

Our approach demands detailed radar configuration parameters and intensive raw analog-to-digital converter (ADC) data processing to integrate super-resolution algorithms effectively and assess network performance across various antenna apertures. Hence, we created our own dataset by driving a Lexus RX450h SUV equipped with multi-modal sensors, including a TI imaging radar, Teledyne FLIR Blackfly S stereo cameras, and a Velodyne Ultra Puck VLP-32C LiDAR sensor, along urban streets, highways and campus roads. The centerpiece of our dataset is the TI cascaded imaging radar system [22], configured for MIMO operations with an array of 12 transmit (TX) and 16 receive (RX) antennas. The operational 9 TX and 16 RX antennas were arranged to form a virtual uniform linear array (ULA) of 86 elements, with half-wavelength spacing, rendering an azimuth resolution of roughly 1.2 degrees via FFT.

The work of R. Zheng, S. Sun and H. Liu was supported in part by National Science Foundation (NSF) under Grants CCF-2153386 and ECCS-2340029.

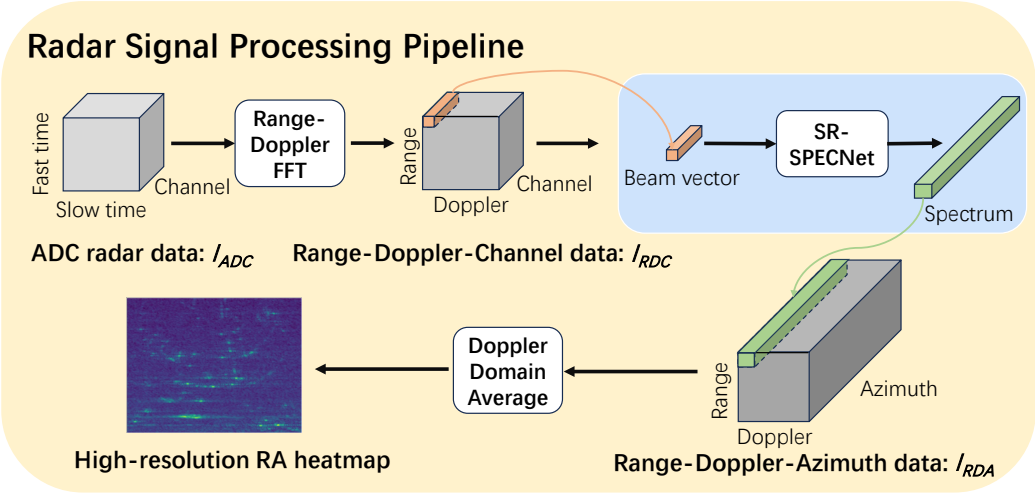

Fig. 1: Radar signal processing pipeline

III. PROPOSED METHOD

We aim to transform raw ADC data into high-resolution RA maps. Unlike recent approaches that derive high-resolution ground truth from RA maps using an expanded antenna array [20], our method relies on RA maps generated with the same number of antennas but refined using IAA algorithms as our benchmark.

Figure 1 depicts our processing workflow. The input ADC data, $I_{\mathrm{ADC}} \in \mathbb{C}^{N_{\mathrm{fast}} \times N_{\mathrm{slow}} \times N_{\mathrm{ch}}}$, encapsulates three dimensions: N_{fast} for fast time samples, N_{slow} for slow time samples (or chirps), and N_{ch} for channels (or receivers). Through a 2D FFT to I_{ADC} across both fast and slow time dimensions, we obtain range-Doppler-channel data, $I_{\mathrm{RDC}} \in \mathbb{C}^{N_{\mathrm{Range}} \times N_{\mathrm{Doppler}} \times N_{\mathrm{ch}}}$. Subsequently, beam vectors $\mathbf{y} \in \mathbb{C}^{1 \times 1 \times N_{\mathrm{ch}}}$ are extracted from each range-Doppler bin. These vectors are processed by SR-SPECNet to generate a super-resolution spectrum. This operation, performed on all beam vectors across all range-Doppler bins, yields the range-Doppler-azimuth data, $I_{\mathrm{RDA}} \in \mathbb{R}^{N_{\mathrm{Range}} \times N_{\mathrm{Doppler}} \times N_{\mathrm{Azimuth}}}$. Notably, this procedure is highly parallelizable, treating the dataset as a 2D matrix with a batch size of $N_{\mathrm{Range}} \times N_{\mathrm{Doppler}}$. The final high-resolution RA maps, $M \in \mathbb{R}^{N_{\mathrm{Range}} \times N_{\mathrm{Azimuth}}}$, are achieved by averaging I_{RDA} over Doppler dimension.

III-A. SR-SPECNet for Azimuth Super-Resolution

IAA is a data-dependent, nonparametric algorithm [16]. By discretizing the DOA space into an L point grid, the array manifold is defined as $\mathbf{A}(\theta) = [\mathbf{a}(\theta_1), \cdots, \mathbf{a}(\theta_L)]$ with $\mathbf{a}(\theta)$ being the array steering vector. The fictitious covariance matrix of $\mathbf{y}$ is represented as $\mathbf{R}_f = \mathbf{A}(\theta)\mathbf{P}\mathbf{A}^H(\theta)$, where P is a $L \times L$ diagonal matrix with the l-th diagonal element being $P_l = |\hat{s}_l|^2$, and $\hat{s}_l$ is the source reflection coefficient corresponding to direction θ_l. IAA iteratively estimates the reflection coefficient $\hat{s}$ and updates the fictitious covariance matrix by minimizing a weighted least-square (WLS) cost function $\|\mathbf{y} - s_l\mathbf{a}(\theta_l)\|^2_{\mathbf{Q}^{-1}(\theta_l)}$, where $\|\mathbf{X}\|^2_{\mathbf{Q}^{-1}(\theta_l)} \triangleq \mathbf{X}^H\mathbf{Q}^{-1}(\theta_l)\mathbf{X}$ and $\mathbf{Q}(\theta_l) = \mathbf{R}_f - P_l\mathbf{a}(\theta_l)\mathbf{a}^H(\theta_l)$.

SR-SPECNet's primary objective is to transform input beam vectors into super-resolution IAA spectra. The IAA, functioning as an advanced beamforming algorithm, iteratively refines a reconstructed covariance matrix to estimate the spectrum as $\hat{\mathbf{s}}_{\mathrm{IAA}} = \mathbf{W}^H\mathbf{y}$, where $\mathbf{W} \in \mathbb{C}^{N_{\mathrm{ch}} \times L}$ are the beamforming weights. At each beamforming angle θ_L, the array's response, $\hat{s}_l = \mathbf{W}^H(\theta_l)\mathbf{y}$, parallels the output process of a multi-layer perceptron (MLP), underscoring the suitability of using an MLP for this application [23].

Designed as a four-layer MLP, SR-SPECNet mirrors the mathematical operations in the IAA algorithm. It processes the input signal $\mathbf{y} \in \mathbb{C}^{N_{\mathrm{ch}}}$, by separating its real and imaginary parts and concatenating them into a real-valued input $\bar{\mathbf{y}} \in \mathbb{R}^{2*N_{\mathrm{ch}}}$. This approach ensures the preservation of crucial phase information, as complex value multiplication inherently involves both the real and imaginary parts.

III-B. Data Preprocessing

Proper data normalization is crucial for training neural networks, especially for regression tasks. Different from simulated signals with controlled factors like SNR, target reflection, and target number, real-world signals add unpredictability in SNR and reflections, challenging normalization. SNR varies significantly within a radar frame and from frame to frame. Maintaining a comparable intensity among beam vectors is crucial for constructing accurate RA heatmaps.

We introduce a frequency domain normalization method designed to produce consistent and interpretable inputs for neural network training. This approach entails determining a normalization factor, α, for each beam vector, calculated as the maximum absolute value of the frequency spectra, obtained by multiplying $\mathbf{A}^H$ to the beam vector, equivalent to an FFT operation, and then divided by the total number of elements, N_{ch}:

$$\alpha = \max\left(\left|\frac{\mathbf{A}^H\mathbf{y}}{N_{\mathrm{ch}}}\right|\right). \tag{1}$$

Subsequently, the raw signal $\mathbf{y}$ is normalized using α to yield $\mathbf{y}_{\mathrm{norm}} = \mathbf{y}/\alpha$, ensuring that the signal levels are stable across varying SNR conditions. Similarly, the label, represented by the IAA spectra $\hat{\mathbf{s}}_{\mathrm{IAA}}$, is normalized to $\mathbf{s}_{\mathrm{norm}} = \hat{\mathbf{s}}_{\mathrm{IAA}}/\alpha$. This normalization strategy effectively scales the signal and the IAA spectra so that their values fall within a comparable range, thereby facilitating more effective network training. Moreover, α preserves the relative intensity across all beam vectors in a radar frame, maintaining

Models	small			medium			large		
	NMSE↓	SSIM↑	PSNR↑	NMSE↓	SSIM↑	PSNR↑	NMSE↓	SSIM↑	PSNR↑
2D U-Net	0.321	0.782	27.780	0.233	0.820	29.182	0.104	0.877	32.736
3D U-Net	0.763	0.841	26.655	0.205	0.894	31.318	0.132	0.917	33.958
SR-SPECNet	0.168	0.909	30.683	0.092	0.946	33.446	0.077	0.955	34.326
SR-SPECNet+	**0.080**	**0.950**	**34.006**	**0.063**	**0.962**	**35.350**	**0.056**	**0.965**	**35.884**

Table I: Performance metrics of deep learning models for super-resolution RA heatmap generation.

the spatial relationships essential for accurate synthesis of RA heatmaps. This normalization process is exclusively needed during training and is not required for generating super-resolution RA heatmaps with test data. This is attributed to the linear relationship between the beam vector and its corresponding spectra.

III-C. SNR-Guided Loss Function

The normalization factor α, which represents the maximum value in the signal's frequency domain, is directly proportional to the signal's SNR. A higher α suggests a higher SNR, positively influencing the quality of the final RA heatmap. We introduce an SNR-guided loss function similar to a weighted mean squared error (MSE), designed to prioritize higher SNR signals during training. The loss function is defined as:

$$\mathcal{L}_{\text{SNR}} = \alpha \cdot \frac{1}{L} \sum_{i=1}^{L} (s_i - \hat{s}_i)^2,$$

where L is the number of angle grid points of the spectra, s_i and $\hat{s}_i$ are the actual and predicted values at θ_i. This approach ensures that our model is finely tuned to emphasize higher quality signals.

IV. EXPERIMENT

We train and evaluate our SR-SPECNet model using our own dataset, which comprises 17,000 frames of raw ADC radar data. To promote data diversity and minimize the redundancy of consecutive frames, we strategically selected every tenth frame from the dataset, yielding 1,700 frames, with the initial 1,400 frames dedicated to training the model, and the subsequent 300 frames reserved for testing. We intentionally structured the training frames into three subsets to simulate real-world data collection scenarios of limited time periods: a 'small' dataset with the initial 200 frames (akin to a 200-second data collection period), a 'medium' dataset comprising the first 700 frames, and a 'large' dataset that includes all 1,400 frames. This segmentation aims to test our model's performance and adaptability under varying lengths of data availability.

IV-A. Benchmarks

To evaluate SR-SPECNet's effectiveness, we compare it with models designed to enhance spatial resolution. This comparison includes a 2D U-Net [24], which transforms low-resolution RA heatmaps into high-resolution equivalents, and the RAD-UNet [20], referred to as a 3D U-Net, that upgrades low-resolution range-azimuth-Doppler (RAD) data to high-resolution RAD imagery. We exclude pixel-based super-resolution networks like SRGAN, which increase resolution by adding pixels. These models do not meet the specific requirements of radar imaging, where resolution is not directly related to pixel count [20].

IV-B. Evaluation Metrics

The RA map is a grayscale image, normalized between 0 to 1, for both generated and ground truth images. To comprehensively evaluate the quality of high-resolution RA heatmaps, we use established image evaluation metrics. PSNR, measured in dB, and SSIM, ranging from 0 to 1, assess image quality where higher values indicate better quality. NMSE also ranges from 0 to 1 and quantifies prediction accuracy by comparing the mean squared error to the variance of actual values, with lower values indicating more accurate predictions. Together, PSNR, SSIM, and NMSE provide a robust framework for assessing image fidelity, error magnitude, and compositional changes affecting perceived quality.

IV-C. Implementation Details

Our radar system configuration employs fast-time samples $N_{\text{fast}} = 256$ and slow-time samples $N_{\text{slow}} = 64$, undergoing post-MIMO processing that results in 86-element beam vectors. The radar data cube, I_{RDC}, is truncated along the range axis to retain the first 100 elements, producing a truncated dataset $I_{\text{RDC}}^{\text{trunc}} \in \mathbb{C}^{100\times 64\times N_{\text{ch}}}$, focusing on significant target information within the first 50 meters. To simplify, we use a 10-element antenna array ($N_{\text{ch}} = 10$) due to the extensive computational demands of applying the IAA on larger arrays.

The angular grid size is set to $L = 256$ to ensure frequency domain uniformity. The labels derive from the IAA spectra of our 10-element arrays. SR-SPECNet is designed with four fully connected layers, the first three followed by ReLU activations with output sizes of 2048, 1024, and 512, respectively. The final layer's output size equals L. Complex input signals are handled by concatenating their real and imaginary parts into a real-valued vector for the network input. SR-SPECNet and benchmark models were implemented in PyTorch, trained using the Adam optimizer at a learning rate of 0.0001 over 500 epochs, and accelerated on four Nvidia RTX A6000 GPUs for enhanced training efficiency.

IV-D. High-Resolution RA Heatmap

We investigated the efficacy of deep neural networks in producing high-resolution RA heatmaps by comparing SR-SPECNet and SR-SPECNet+, trained using MSE loss and our novel SNR-guided loss, respectively, against established benchmarks. As shown in Table I, both models were evaluated across varying dataset sizes.

Table I demonstrates the superior performance of SR-SPECNet+, which surpasses both 2D and 3D U-Net benchmarks as well as SR-SPECNet in all dataset scenarios. This performance advantage is largely due to the SNR-guided loss function, which integrates domain knowledge to prioritize data with higher signal integrity during training. This specialized loss function enables SR-SPECNet+ to attain the lowest NMSE and the highest SSIM and

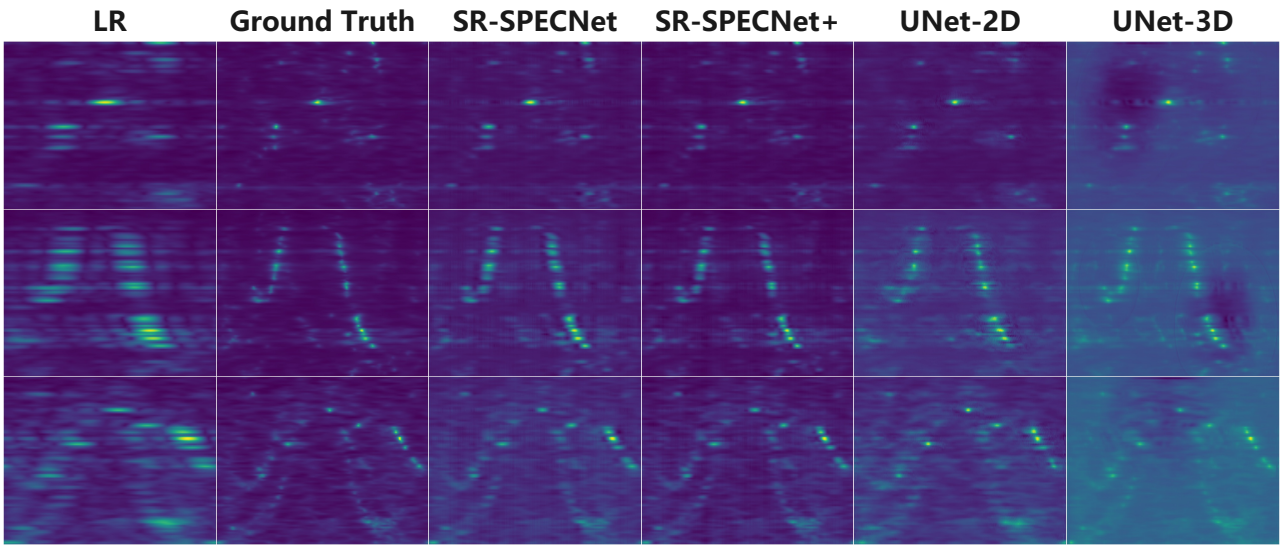

Fig. 2: RA heatmap quality comparison. Heatmaps from the same radar frame are reconstructed by SR-SPECNet, SR-SPECNet+, and baseline 2D U-Net and 3D U-Net models, alongside the ground truth. Each row corresponds to heatmaps generated from the same radar frame data.

PSNR scores, showcasing its robustness even with smaller datasets. SR-SPECNet also exceeds the performance of traditional U-Net models, highlighting the effectiveness of approaching RA heatmap generation through a 1D super-resolution perspective. The results clearly indicate that domain-specific enhancements significantly augment SR-SPECNet+'s capacity to produce superior heatmaps.

IV-E. Complexity

The 2D and 3D U-Nets, tailored for processing low-resolution RA and RAD heatmaps, possess significant model sizes with approximately 31.0M and 51.8M parameters, respectively. Their inference times are 6.14 ms for the 2D U-Net and 6.98 ms for the 3D U-Net, reflecting the increased computational demands of handling higher-dimensional data. In contrast, our SR-SPECNet demonstrates remarkable efficiency, operating with only 1M parameters and achieving a rapid inference time of 3.12 ms. This substantial reduction in both time and model size significantly boosts the speed and resource efficiency of our approach, making it highly suitable for real-time automotive radar applications on embedded CPUs. Additionally, the IAA records inference times of 12.3 ms for 10-element vectors, highlighting its computational intensity.

IV-F. Generalizability

Our training and test datasets encompass a diverse array of signals, with each radar frame containing thousands of signals to ensure variety. To assess model generalizability, we tested our pre-trained SR-SPECNet on 10,500 radar frames from Radartron [25], which vary in N_{Range} and N_{Doppler}. Unlike the 2D and 3D U-Nets, which struggled with these variations, our 1D approach was effortlessly adapted, demonstrating superior scalability. The performance metrics for SR-SPECNet included a NMSE of 0.025, SSIM of 0.983, and PSNR of 37.858, underscoring its effective performance.

IV-G. Visualization of RA maps

To assess the quality of high-resolution RA heatmaps produced by SR-SPECNet and SR-SPECNet+, we visually compared the outputs to those from baseline models, namely 2D U-Net and 3D U-Net, as shown in Fig. 2. Each model was trained using the 'large' dataset to optimize performance. The ground truth heatmaps, generated using the IAA, exhibit significantly better resolution than the low-resolution (LR) images. Additionally, IAA's ability to suppress sidelobes results in much clearer heatmaps, highlighting the advantages of employing IAA-derived heatmaps as ground truth for training instead of relying on FFT-generated heatmaps, which require larger antenna apertures and may not be feasible due to the high costs and complexities associated with MIMO technology.

As depicted in Fig. 2, the RA heatmaps produced by SR-SPECNet and SR-SPECNet+ using an SNR-guided loss function show marked improvements over those generated by the 2D and 3D U-Nets. These enhancements align with the quantitative metrics presented in Table I, further validating the superiority of our proposed 1D approaches over traditional 2D and 3D methods.

V. CONCLUSIONS

In this study, we have advanced automotive radar imaging by introducing SR-SPECNet, a novel 1D network leveraging IAA-generated RA heatmaps as ground truth and integrating a unique SNR-guided loss function for super-resolution RA heatmap generation. Our approach, emphasizing a 1D signal processing perspective, has demonstrated superior performance on real radar measurements in terms of automotive radar imaging quality, and efficiency across varying dataset sizes. These contributions enhance the fidelity of radar imaging in autonomous vehicles. They also opens avenues for future research, especially with our commitment to sharing our own radar dataset and source code resources with the research community. This work underscores the potential of deep learning-enhanced radar processing in improving navigational safety and robust perception in autonomous vehicles.

VI. REFERENCES

[1] S. Patole, M. Torlak, D. Wang, and M. Ali, "Automotive radars: A review of signal processing techniques," *IEEE Signal Processing Magazine*, vol. 34, no. 2, pp. 22–35, 2017.

[2] F. Engels, P. Heidenreich, A. M. Zoubir, F. K. Jondral, and M. Wintermantel, "Advances in automotive radar: A framework on computationally efficient high-resolution frequency estimation," *IEEE Signal Processing Magazine*, vol. 34, no. 2, pp. 36–46, 2017.

[3] S. Sun, A. P. Petropulu, and H. V. Poor, "MIMO radar for advanced driver-assistance systems and autonomous driving: Advantages and challenges," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 98–117, 2020.

[4] 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.

[5] Z. Peng, C. Li, and F. Uysal, *Modern Radar for Automotive Applications*. London, UK: IET, 2022.

[6] R. Zheng, S. Sun, H. Liu, and T. Wu, "Deep-neural-network-enabled vehicle detection using high-resolution automotive radar imaging," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 59, no. 5, pp. 4815–4830, 2023.

[7] S. Sun and Y. D. Zhang, "4D automotive radar sensing for autonomous vehicles: A sparsity-oriented approach," *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 4, pp. 879–891, 2021.

[8] G. Duggal, S. Vishwakarma, K. V. Mishra, and S. S. Ram, "Doppler-resilient 802.11ad-based ultrashort range automotive joint radar-communications system," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 5, pp. 4035–4048, 2020.

[9] F. Engels, P. Heidenreich, M. Wintermantel, L. Stäcker, M. Al Kadi, and A. M. Zoubir, "Automotive radar signal processing: Research directions and practical challenges," *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 4, pp. 865–878, 2021.

[10] J. Li and P. Stoica, "MIMO radar with colocated antennas," *IEEE Signal Processing Magazine*, vol. 24, no. 5, pp. 106–114, 2007.

[11] ——, *MIMO Radar Signal Processing*. Hoboken, NJ, Wiley, 2009.

[12] J. Bergin and J. R. Guerci, *MIMO Radar: Theory and Application*. Boston, MA, Artech House, 2018.

[13] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[14] E. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse problems*, vol. 23, no. 3, p. 969, 2007.

[15] E. Candes and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Communications on Pure and Applied Mathematics*, vol. 67, no. 6, pp. 906–956, 2014.

[16] T. Yardibi, J. Li, P. Stoica, M. Xue, and A. Baggeroer, "Source localization and sensing: A nonparametric iterative adaptive approach based on weighted least squares," vol. 46, no. 1, pp. 425–443, 2010.

[17] W. Roberts, P. Stoica, J. Li, T. Yardibi, and F. Sadjadi, "Iterative adaptive approaches to MIMO radar imaging," vol. 4, no. 1, pp. 5–20, 2010.

[18] K. Armanious, S. Abdulatif, F. Aziz, U. Schneider, and B. Yang, "An adversarial super-resolution remedy for radar design trade-offs," in *2019 27th European Signal Processing Conference (EUSIPCO)*, 2019, pp. 1–5.

[19] A. Geiss and J. C. Hardin, "Radar super resolution using a deep convolutional neural network," *Journal of Atmospheric and Oceanic Technology*, vol. 37, no. 12, pp. 2197–2207, 2020.

[20] Y.-J. Li, S. Hunt, J. Park, M. O'Toole, and K. Kitani, "Azimuth super-resolution for FMCW radar in autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 17 504–17 513.

[21] R. Zheng, S. Sun, H. Liu, H. Chen, M. Soltanalian, and J. Li, "Antenna failure resilience: Deep learning-enabled robust DOA estimation with single snapshot sparse arrays," in *58th Asilomar Conference on Signals, Systems, and Computers*, 2024. [Online]. Available: https://arxiv.org/abs/2405.02788

[22] Texas Instruments, *Design Guide: TIDEP-01012 Imaging Radar Using Cascaded MmWave Sensor Reference Design*, 2020, rev. A. [Online]. Available: https://www.ti.com/lit/ug/tiduen5a/tiduen5a.pdf

[23] M. Naumovski and R. Carrasco, "Neural network beamformer for narrow-band HF transmission," in *IEE Colloquium on HF Antennas and Propagation*, 1995, pp. 5/1–5/8.

[24] 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: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.

[25] S. Madani, J. Guan, W. Ahmed, S. Gupta, and H. Hassanieh, "Radatron: Accurate detection using multi-resolution cascaded MIMO radar," in *European Conference on Computer Vision (ECCV)*, 2022, pp. 160–178.