
CNNs for Interference Mitigation and Denoising in Automotive Radar Using Real-World Data

Johanna Rock¹, Mate Toth^{1,2}, Paul Meissner², Franz Pernkopf¹

¹Graz University of Technology, Austria

²Infineon Technologies Austria AG

johanna.rock@tugraz.at

Abstract

Radar sensors are crucial for environment perception of driver assistance systems as well as autonomous cars. Key performance factors are a fine range resolution and the possibility to directly measure velocity. With a rising number of radar sensors and the so far unregulated automotive radar frequency band, mutual interference is inevitable and must be dealt with. Sensors must be capable of detecting, or even mitigating the harmful effects of interference, which include a decreased detection sensitivity. In this paper, we evaluate a Convolutional Neural Network (CNN)-based approach for interference mitigation on real-world radar measurements. We combine real measurements with simulated interference in order to create input-output data suitable for training the model. A finite sample size performance comparison shows the effectiveness of the model trained on either simulated or real data as well as for transfer learning. A comparative performance analysis with the state of the art emphasizes the potential of CNN-based models for interference mitigation and denoising of real-world measurements, also considering resource constraints of the hardware.

1 Introduction

Automotive radar sensors are key elements of current driver assistance systems and autonomous driving applications. In the automotive context, *frequency modulated continuous wave (FMCW)/chirp sequence (CS)* radars are prevalent. They transmit sequences of linear chirp signals in a shared and non-regulated spectrum. Ever larger *radio frequency (RF)* transmit bandwidths are required to fulfill the demand on fine range resolution. Because of these larger bandwidths and because of a rising number of deployed radar sensors, mutual interference is becoming increasingly likely. Non-coherent interference, in which radar sensors with non-identical transmit signal parameters interfere, is the most common form of mutual interference [1]. This leads to a reduced object detection sensitivity [2]. Therefore, interference mitigation is a crucial part of current and future radar sensors used in a safety context. Several conventional signal processing algorithms have been proposed in order to mitigate mutual interference. The most basic method is to zero out the interference-affected signal samples. More advanced methods use nonlinear filtering in slow-time [3], iterative reconstruction using Fourier transforms and thresholding [4], estimation and subtraction of the interference component [5], or beamforming [6]. Some machine learning techniques were discussed in the context of interference detection and classification in [7]. *Convolutional Neural Networks (CNNs)* have been successfully used for image denoising, e.g. in [8]. CNN-based interference mitigation and denoising methods presented in [9] can be applied to *range-Doppler (RD)* maps. A two-channel representation of the complex spectrogram data (i.e. real and imaginary data) is used as network input. Experimental results show a strong denoising and interference mitigation capability in comparison to state-of-the-art signal processing algorithms, though evaluated only on simulated data. The applicability of these

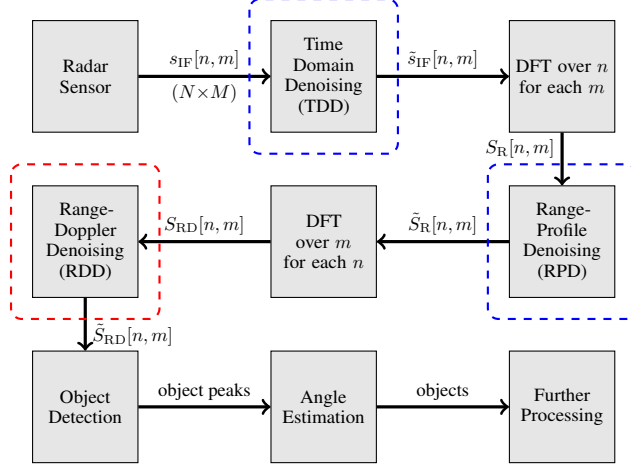


Figure 1: Block diagram of a basic FMCW/CS radar processing chain. Dashed boxes indicate the locations of optional interference mitigation steps, including the CNN-based approach (red) and classical methods (blue).

models for robust interference mitigation on real-world data has not been investigated so far. In this paper, we analyze the suitability of CNN-based models from [9] to perform interference mitigation and denoising on real-world radar measurements. Therefore an extensive measurement campaign in a typical inner-city road traffic scenario has been carried out. The interference is simulated for both, simulated object scenarios and real-world measurements. Due to the absence of labeled object positions in the target RD maps, we use the *cell averaging constant false alarm rate (CA-CFAR)* algorithm [10] to identify the most likely object positions. These positions are the basis for our performance comparison using the *signal-to-interference-plus-noise ratio (SINR)* [11].

Main contributions of this paper are:

- We consider real-world radar measurements combined with simulated interference for CNN-based interference mitigation and denoising of RD maps.
- We analyze the effect of finite sample sizes on model performance and robustness.
- We present numerical results using application-related performance metrics in a comparison with the state of the art, i.e. Zeroing, IMAT and Ramp filtering.
- We show that an excellent level of noise reduction and hence an improvement of detection sensitivity can be achieved on real-world measurements.

2 Signal model

The *range-Doppler (RD) processing* chain of a common FMCW/CS radar is depicted in Fig. 1. A measurement is performed by transmitting a succession of M linearly modulated RF chirp sequences. By mixing each chirp, also termed *ramp*, with the received object reflections, the *intermediate frequency (IF)* signal is obtained. It consists of sinusoidal components corresponding to objects. The objects' distances and velocities are contained in the sinusoidals' frequencies and their linear phase change over successive ramps [12, 13], respectively. We obtain N samples per ramp; from a data processing point of view the IF signal can be interpreted as a $N \times M$ data matrix $s_{IF}[n, m]$. Note that the indices n and m are commonly referred to as *fast-* and *slow-time*, respectively. Subsequently in the processing chain, discrete Fourier transforms (DFTs) are computed over both dimensions, yielding a two-dimensional spectrum on which peaks can ideally be found at positions corresponding to the objects' distances and velocities. After peak detection, further processing can include angular estimation, tracking, and classification. Besides object reflections, real radar measurements may also include receiver noise and interference. In order to model these disturbances, we define the IF signal as

$$s_{IF}[n, m] = \sum_{o=1}^{N_O} s_{O,o}[n, m] + \sum_{i=1}^{N_I} s_{I,i}[n, m] + v[n, m], \quad (1)$$

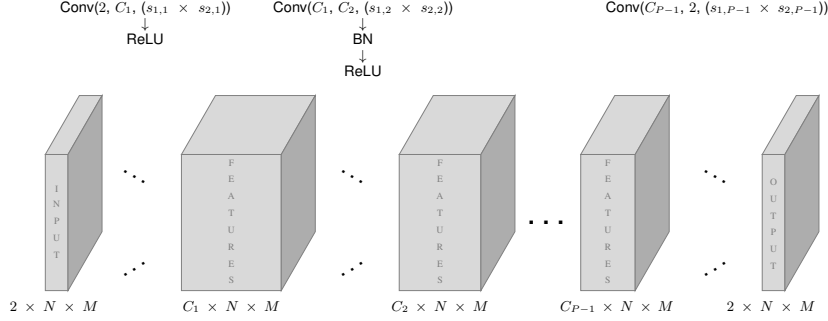


Figure 2: CNN architecture for radar signal denoising. It uses *ReLU*, *Batch Normalization (BN)* and the convolution operation $\text{Conv}(i, o, (s_1 \times s_2))$, for i input channels, o output channels and a kernel size of $s_1 \times s_2$.

where $s_{O,o}[n, m]$ is the o^{th} object reflection, $s_{I,i}[n, m]$ corresponds to interference from the i^{th} interferer assuming N_I interferers, and $v[n, m]$ models the noise. In the simulated radar signal we use *additive white Gaussian noise (AWGN)* to model the receiver noise and point objects with random distances and velocities to model the object reflections. Real measurements already contain object reflections mixed with receiver noise. In both cases, the interference is simulated and included in the time domain IF signals. Noncoherent mutual interference essentially generates time-limited broadband disturbances, see [1, 14] for details. State-of-the-art (“classical”) interference mitigation methods are mostly signal processing algorithms that are applied either on the time domain signal $s_{\text{IF}}[n, m]$ or on the frequency domain signal $S_{\text{R}}[n, m]$ after the first DFT. The CNN-based method used in this paper, also denoted *Range-Doppler Denoising (RDD)*, is applied on the RD map after the second DFT.

3 CNN model architecture

The interference mitigation and denoising CNN architecture is based on the models from [9]. *Range Doppler Denoising (RDD)*, as labeled in Figure 1) is used for evaluation and comparison, because of its superior performance on past experiments with simulated data. Figure 2 illustrates the CNN-based architecture, which consists entirely of convolutional composite layers. The first layer performs convolution operations and *ReLU* [15] activation functions; subsequent layers additionally include *Batch Normalization (BN)* [16]. An exception is the last layer, which uses a linear activation function and two kernels in order to map to the real and imaginary data. The amount of kernels in a layer is chosen to be a power-of-two and decreasing for subsequent layers, e.g. $[2^6, 2^5, 2^4, 2^2]$ or $[2^6, 2^5, 2^5, 2^4, 2^2]$, as inspired by [17]. RDD is applied to radar snapshots after the second DFT, hence the input samples are complex valued patches of size $N \times M$. Square kernels are used in combination with zero-padding, such that the inputs and outputs for each layer have the same spatial dimension. We use two input channels in order to represent the real- and imaginary parts of the complex valued input. For training the network we use the *mean squared error (MSE)* loss function and the *Adam* [18] algorithm.

4 Experimental setup

4.1 Data sets

In our experiments we evaluate two data sets. The first one is purely simulated including object reflections, noise and interference. The second data set consists of real-world measurements, that are combined with simulated interferences. This way we have access to training inputs and their corresponding targets with a limited measurement expenditure. Nevertheless, realistic scenarios are the basis of training and evaluation, and thus give an insight of interference mitigation performance on real-world data. Both data sets are split into three partitions for training (2500 snapshots), validation (250 snapshots) and testing (250 snapshots) the models. Data from a single measurement, consisting

of 32 snapshots and sixteen antennas, are exclusively contained either in the training, validation, or test set.

Simulation: The basic receive IF signal is generated according to (1) and processed as described in Section 2. The signals are generated based on several parameters, that are sampled from uniform distributions $\mathcal{U}(min, max)$ in the respective domains. Among them are the number of objects $\mathcal{U}(1, 20)$ and for each object the relative distance $\mathcal{U}(0m, 100m)$ and velocity $\mathcal{U}(-20m/s, 20m/s)$. The interferer parameters are uniformly sampled within the ranges listed in Table 1, while the ego radar parameters (see Table 2) are constant for all simulations. The *signal-to-noise ratio (SNR)* and the *signal-plus-noise-to-interference ratio (SNIR)* are used to scale the noise and interference powers relative to the object-signal power and object-signal-plus-noise power respectively, when generating the interfered and noisy time domain signal $s_{IF}[n, m]$. Figure 3 shows a RD map processed from a simulated signal with five objects, where Figure 3(a) shows an interfered signal and Figure 3(b) shows the corresponding clean data with AWGN.

Table 1: Ranges of interference and noise parameters.

Parameter		Value	Value
$f_{0,1}$	Sweep start frequency	78.9GHz	79.1GHz
B_1	Sweep bandwidth	0.15GHz	0.25GHz
T_1	Sweep duration	12 μ s	24 μ s
SNR	Signal-to-noise-ratio	-15.5dB	-0.5dB
SNIR	Signal-plus-noise-to-interference-ratio	15dB	35dB

Table 2: Victim radar and signal processing parameters for simulation and measurements.

Parameter		Value
$f_{0,1}$	Sweep start frequency	79GHz
B_1	Sweep bandwidth	0.27GHz
T_1	Sweep duration	12.8 μ s
$B_{IF,V}$	IF bandwidth	10MHz
N	Number of fast-time samples	512
M	Number of slow-time samples/ ramps	128
A	Number of antennas	16
w	Window type	Hann

Real-world measurements: The measurements were recorded in typical inner-city traffic scenarios. We used a cargo bicycle with a radar apparatus mounted on the front and additional measurement equipment in the cargo container. The device was configured according to the parameters in Table 2. One measurement denotes 32 consecutive snapshots recorded with sixteen antennas, where each snapshot is associated with a wide-angle camera picture for reference. An input sample for the CNN, i.e. a RD map, is processed from one snapshot of a single antenna. The measurement signal consists of object reflections (static and moving) as well as receiver noise. The interference is simulated as described in Section 4.1. The SNIR is used for scaling the interference relative to the object signal plus receiver noise power. Figure 4 shows a RD map processed from a real-world measurement, where Figure 4(a) shows an interfered signal, Figure 4(b) shows the corresponding clean signal and Figures 4(c) shows the respective camera snapshot for reference.

Experimental analysis of simulation and measurements: The simulated signal is modeled according to reflections from point objects, which results in single, clear and well-shaped object peaks in the RD map. All distances, velocities and angles are randomly sampled, i.e. there is no observable bias towards object peak positions in the RD map. In the real-world measurements on the other hand, we observe more complex objects, which consist of object peak clusters that are often distributed along the distance as well as the velocity axis of the RD map. Furthermore, a strong bias of object velocities towards the negative velocity of the measurement vehicle is present. This bias results from static objects; they are contained in velocity bins close to the negative ego velocity. Also, there exist strong reflections within the first few meters at a relative velocity of zero, i.e. the reflections of the radar and the measurement vehicle themselves. Another bias, though less severe, can be observed




Figure 3: Exemplary RD magnitude spectra of a simulated scenario with five objects in dB.

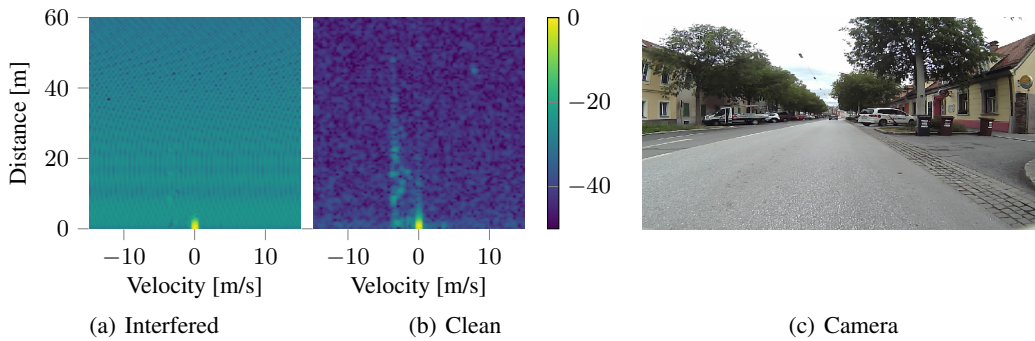


Figure 4: Exemplary RD magnitude spectra of a real-world measurement in dB and a reference camera snapshot.

regarding the physical positions of moving objects, caused by the relative position of bicycle lanes to car lanes in the measurement environment.

4.2 Performance measures

Quantitative measures: The *signal-to-interference-plus-noise ratio (SINR)* directly relates to the object detection sensitivity [11], i.e. it significantly influences the chance that an object is detected on the RD map. It is the ratio of signal power at the object peaks to the noise floor. For multi-object scenarios in the RD domain the SINR is defined as:

$$\text{SINR} = 10 \log \left(\frac{\frac{1}{N_O} \sum_{\{n,m\} \in \mathcal{O}} |\tilde{S}_{\text{RD}}[n, m]|^2}{\frac{1}{N_N} \sum_{\{n,m\} \in \mathcal{N}} |\tilde{S}_{\text{RD}}[n, m]|^2}} \right), \quad (2)$$

where n and m are row and column indices of the RD matrix, \mathcal{O} is the set of object peaks and \mathcal{N} is the set of N_N noise cells. Cells are considered as noise, when they have a minimum distance to each object peak. The CA-CFAR detector [10] is used to find the RD map positions of the most prominent object peaks in both, the simulated and the measurement data. We apply the detection algorithm with a window of 5×10 and two guard cells in each dimension. For peaks close to borders, we only consider cells lying within the RD map as reference window cells.

Qualitative measures: During visual inspection of the RD map, we consider criteria such as object peak and noise floor magnitude, object peak location, resolution and distortion as well as artifact appearances.

4.3 Mitigation methods selected for comparison

A small representative subset of state-of-the-art signal processing algorithms have been chosen for a comparative analysis. This also allows for a discussion of the properties of the CNN-based approaches in a broader context of interference mitigation. A short summary of these methods is presented below.

Zeroing: Simple method in which samples of the IF signal affected by interference are set to a value of zero, see e.g. [19].

Iterative method with adaptive thresholding (IMAT): After an initial application of zeroing, the zeroed samples may be interpolated. IMAT [20] is a promising interpolation method based on an iterative compressed sensing algorithm.

Ramp filtering: Ramp filtering processes the signal after the first DFT, using a non-linear filter in slow-time to strongly suppress interference as well as noise. Several choices for the filter can be considered [3].

Note that both, zeroing and IMAT, require the detection of interfered IF signal samples. In this paper, it will be assumed that this operation works perfectly. However, in general, errors in interference detection may have a strong impact on the performance of mitigation algorithms [11]. Ramp filtering, as well as the proposed CNN-based approaches, are not directly based on an interference detection step.

5 Experimental results

All models are based on the architecture described in Section 3. In order to find suitable hyper parameters for real-world data, we run simulations using a different number of layers (2, 3, ..., 10) and maximal number of kernels (2^n , $n = 3, 4, \dots, 8$). A kernel size of 3×3 is used for all simulations, because of its clear superiority on past experiments. An architecture with 4 layers and [512, 32, 16, 2] kernels in these layers reaches the highest SINR performance and is selected for further evaluations. In order to analyze the effect of training set sample sizes, we perform a finite sample size performance comparison. For these evaluations we use the same test set, namely with real measurements, and vary the training set consisting either of the simulated data, real measurements, or both in the context of transfer learning. Finally, we provide a performance comparison with classical interference mitigation methods.

5.1 Finite sample size performance comparison on real-world data

We investigate the interference mitigation and denoising capabilities on real measurements dependent on the amount of training samples. Therefore the same test set, namely with measurement data, is used for all evaluations. For training we consider three different data sets consisting of: simulated data (Sim), measurement data (Real), and both, simulated data to pre-train the model and measurement data for fine-tuning, in the context of transfer learning (Transfer). In the case of simulated training data, we run simulations using also simulated data for validation (Sim), and using measurement data for validation (Sim+VReal). For each data set we train the model using 50, 100, ..., 600 samples. In the case of transfer learning, we pre-train the model with 500 simulated samples and fine-tune with the respective amount of measurements. Figure 5 shows the relation of SINR performance to the number of samples in the training set. The top and bottom figures show the mean and variance of the SINR, based on twenty simulations per configuration using a randomly selected training subset. The interfered and clean signal SINRs are indicated by a dashed horizontal line for reference. According to our evaluations, also the training with simulated data results in a very high interference mitigation and denoising performance. This indicates, that our model indeed learns to remove the interference and noise instead of learning the object scenarios. Naturally, the outstanding performance is possible, because we use simulated interference also for testing. Nonetheless, the evaluation shows, that interference mitigation can be generalized to unseen object scenarios, such as real-world measurements, as long as realistic interference is used during training. Training on simulated data seems to be very stable according to the small variance over all training set sizes. This may result from the simpler nature of simulated data, which is beneficial for the learning process. For training with measurement data and transfer learning, the SINR is reduced for fewer training samples;

it increases with a rising number of samples until it reaches the clean SINR with around 400 samples and surpasses the performance of models trained on simulated data with 450 (Real) or 500 (Transfer) samples. During this rise, the variance increases as the model’s training progress highly depends on the significance of the randomly sampled training subset. With more than 300 samples, the variance drops again and stabilizes at a low level with more than 500 samples.

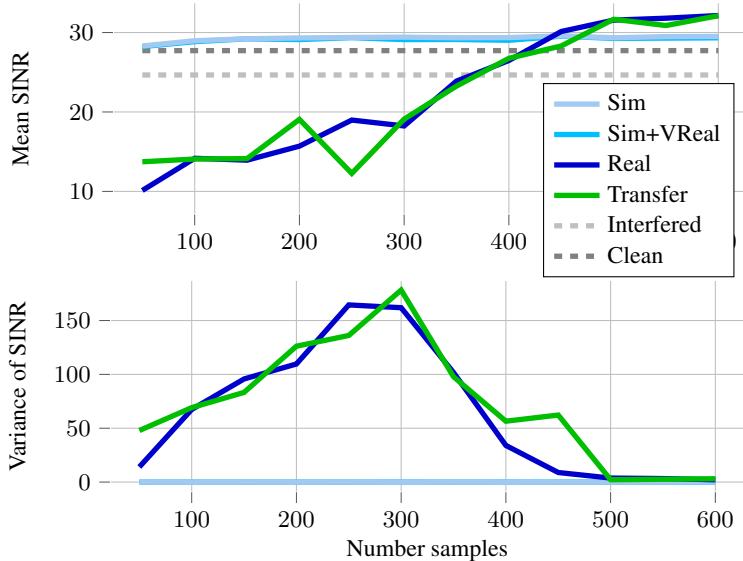


Figure 5: Relation of SINR to the number of samples in the training set.

5.2 Performance comparison with classical interference mitigation methods

Three classical interference mitigation methods, as described in Section 4.3, were implemented and evaluated using SINR. The results are statistically compared to the CNN-based model, that was trained and validated on 2500 and 250 snapshots from real measurements, respectively. For the evaluation we used a Monte-Carlo-Simulation with 250 measurement snapshots. Figure 6 shows the empirical *cumulative density function (CDF)* of their evaluated SINR values. The ‘clean’ measurement and interfered signals are included as reference. One of the snapshot RD maps with interference, without interference and with mitigation using the CNN-based model is displayed in Figure 7. The three classical methods, i.e. zeroing, IMAT and Ramp filtering, all improve the SINR slightly in snapshots with strong interference. For moderate to weak interference on the other hand, they are not capable of removing interference effects and even decrease the SINR of the interfered signal, i.e. perform worse than without mitigation. The CNN-based model outperforms the classical methods for all tested interference levels and even surpasses the SINR of the ‘clean’ signal without interference, thus it performs additional denoising of the noisy real-world measurements. The shape of the CDF indicates, that this outstanding interference mitigation and denoising performance is also robust over all test set samples.

6 Conclusion

In this paper, we use CNN-based mutual interference mitigation and denoising models on real-world data. The training and evaluation framework was extended to handle real measurement data and simulated interference was used to obtain input- and output-pairs for training and evaluating the CNN models. Our experiments show, that the CNN-based interference mitigation approach is also applicable to real-world measurements and results in an outstanding and robust performance. We illustrate the impact of training set coverage to the performance and robustness of the models. The use of simulated data for training can reduce the amount of required real measurements. In a performance comparison with classical interference mitigation methods, the CNN-based model outperforms the state of the art and shows robust behavior in our Monte-Carlo-Simulation. The most important task in

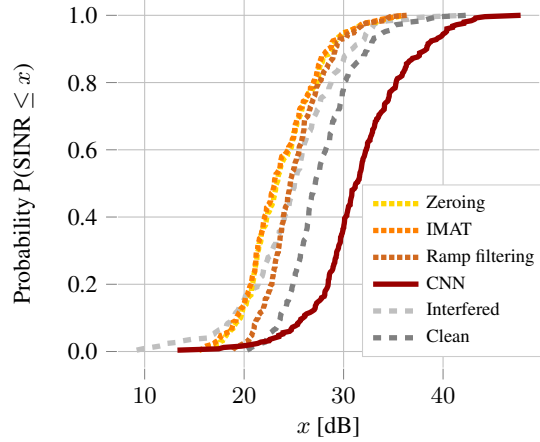


Figure 6: CDF comparison of SINR with classical methods.




Figure 7: Exemplary RD magnitude spectra from the measurements test set.

the future is to collect real interference measurements and to evaluate the generalization capabilities of the CNN-based models on these data.

Acknowledgments

This work was supported by the Austrian Research Promotion Agency (FFG) under the project SAHaRA (17774193). Thanks also to NVIDIA for providing GPUs.

References

- [1] M. Toth, P. Meissner, A. Melzer, and K. Witrissal, “Analytical Investigation of Non-Coherent Mutual FMCW Radar Interference,” in *2018 European Radar Conference (EURAD)*, pp. 71–74, 2018.
- [2] G. M. Brooker, “Mutual Interference of Millimeter-Wave Radar Systems,” *IEEE Transactions on Electromagnetic Compatibility*, vol. 49, no. 1, pp. 170–181, 2007.
- [3] M. Wagner, F. Sulejmani, A. Melzer, P. Meissner, and M. Huemer, “Threshold-Free Interference Cancellation Method for Automotive FMCW Radar Systems,” in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018.
- [4] F. Marvasti, M. Azghani, P. Imani, P. Pakrouh, S. Heydari, A. Golmohammadi, A. Kazerouni, and M. Khalili, “Sparse signal processing using iterative method with adaptive thresholding (IMAT),” in *2012 19th International Conference on Telecommunications (ICT)*, 2012.
- [5] J. Bechter, K. D. Biswas, and C. Waldschmidt, “Estimation and cancellation of interferences in automotive radar signals,” in *2017 18th International Radar Symposium (IRS)*, pp. 1–10, 2017.
- [6] J. Bechter, K. Eid, F. Roos, and C. Waldschmidt, “Digital beamforming to mitigate automotive radar interference,” *2016 IEEE MTT-S Int. Conf. Microwaves Intell. Mobility, ICMIM 2016*, pp. 2–5, 2016.

- [7] R. Zhang and S. Cao, "Support Vector Machines for Classification of Automotive Radar Interference," *2018 IEEE Radar Conf.*, pp. 366–371, 2018.
- [8] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [9] J. Rock, M. Toth, E. Messner, P. Meissner, and F. Pernkopf, "Complex signal denoising and interference mitigation for automotive radar using convolutional neural networks," in *2019 22nd International Conference on Information Fusion (FUSION) (FUSION 2019)*, 2019.
- [10] L. Scharf and C. Demeure, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Addison-Wesley series in electrical and computer engineering, Addison-Wesley Publishing Company, 1991.
- [11] M. Toth, P. Meissner, A. Melzer, and K. Witrissal, "Performance comparison of mutual automotive radar interference mitigation algorithms," in *IEEE Radar Conference*, 2019.
- [12] A. G. Stove, "Linear FMCW radar techniques," *IEE Proceedings F - Radar and Signal Processing*, vol. 139, no. 5, pp. 343–350, 1992.
- [13] V. Winkler, "Range Doppler detection for automotive FMCW radars," in *2007 European Microwave Conference*, pp. 1445–1448, Oct. 2007.
- [14] G. Kim, J. Mun, and J. Lee, "A Peer-to-Peer Interference Analysis for Automotive Chirp Sequence Radars," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 9, pp. 8110–8117, 2018.
- [15] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks.," in *AISTATS* (G. J. Gordon, D. B. Dunson, and M. Dudík, eds.), vol. 15 of *JMLR Proceedings*, pp. 315–323, JMLR.org, 2011.
- [16] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *CoRR*, vol. abs/1502.03167, 2015.
- [17] Y. Jiang, H. Li, and M. Rangaswamy, "Deep learning denoising based line spectral estimation," *IEEE Signal Processing Letters*, vol. 26, no. 11, pp. 1573–1577, 2019.
- [18] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [19] C. Fischer, *Untersuchungen zum Interferenzverhalten automobiler Radarsensorik*. PhD thesis, Ulm University, 2016.
- [20] J. Bechter, F. Roos, M. Rahman, and C. Waldschmidt, "Automotive Radar Interference Mitigation Using a Sparse Sampling Approach," in *2017 European Radar Conference (EURAD)*, pp. 90–93, 2017.