

Automotive Radar Interference Mitigation with Unfolded Robust PCA based on Residual Overcomplete Auto-Encoder Blocks

Nicolae-Cătălin Ristea¹, Andrei Anghel¹, Radu Tudor Ionescu², Yonina C. Eldar³

¹University Politehnica of Bucharest, Romania,

²University of Bucharest, Romania,

³Weizmann Institute of Science, Israel

Abstract

In autonomous driving, radar systems play an important role in detecting targets such as other vehicles on the road. Radars mounted on different cars can interfere with each other, degrading the detection performance. Deep learning methods for automotive radar interference mitigation can successfully estimate the amplitude of targets, but fail to recover the phase of the respective targets. In this paper, we propose an efficient and effective technique based on unfolded robust Principal Component Analysis (RPCA) that is able to estimate both amplitude and phase in the presence of interference. Our contribution consists in introducing residual overcomplete auto-encoder (ROC-AE) blocks into the recurrent architecture of unfolded RPCA, which results in a deeper model that significantly outperforms unfolded RPCA as well as other deep learning models. We also show that our approach achieves a faster processing time compared to state-of-the-art fully convolutional networks, thus being a suitable candidate to be deployed on devices embedded on vehicles.

1. Introduction

Road safety is a key issue in autonomous driving systems, requiring vehicles to perceive their surroundings. One of the most common proposals made by automotive companies is to employ radar sensors in order to build systems that allow cars to scan the surrounding environment. Usually, the radar sensors used in the automotive industry are frequency modulated continuous wave (FMCW) / chirp sequence radars, which transmit sequences of linear chirp signals. The signals transmitted and received by such sensors provide the means to estimate the distance and the velocity of nearby targets (*e.g.*, vehicles, pedestrians or other obstacles).

The rapidly increasing number of radar sensors [11] unavoidably leads to a higher probability of radio frequency

interference (RFI), which generates corrupted and unusable signals. The most common RFI effect involves raising the noise floor by a large margin, to the point where potential targets are completely hidden by noise, thus reducing the sensitivity of target detection methods [3]. In order to be able to accurately detect targets from signals affected by RFI, the interference has to be mitigated. To address this problem, researchers have proposed various techniques ranging from handcrafted approaches [2, 9, 13, 12, 22, 23] to deep learning methods [5, 6, 15, 16, 17].

There are many classical RFI mitigation methods [2, 9, 13, 12, 22, 23], which are usually classified in accordance with the domain in which the interference is mitigated: time, polarization, frequency, code and space. A detailed analysis of these methods is presented in [12]. When the transmitted signal is a linear chirp waveform, one of the most common methods to mitigate interference is to detect, in various ways, the samples of the beat signal affected by interference [13] and convert them to zeroes in the time domain. This is commonly known as the zeroing technique. While this approach is fairly simple, it removes part of the useful signal and can become ineffective when the interference has a long duration.

A series of recent methods [5, 6, 16, 17] rely on deep learning models to mitigate RFI. Rock et al. [17] proposed a convolutional neural network (CNN) to address RFI, aiming to reduce the noise floor while preserving the signal components of detected targets. The authors reported promising results, but they still had concerns regarding the generalization capacity on real data. Another approach that relies on CNNs is proposed in [6]. The method is based on the U-Net architecture [18], performing interference mitigation as a denoising task directly on the range-Doppler map. Fuchs et al. [6] surpassed classical approaches, but their method fails to fully preserve the phase information. Similarly, Fan et al. [5] proposed a method based on CNNs, their contribution being that of adding residual connections, inspired by the ResNet model [8], into the architecture. In another recent work, Ristea et al. [16] proposed fully convolutional

networks (FCNs) trained on synthetic data samples. The FCNs, as well as other deep models, can estimate the absolute value of range profiles, but are not able to obtain the phase information.

Different from the related works presented above, we employ a decomposition algorithm based on unfolded robust Principal Component Analysis (RPCA) [19], introducing residual overcomplete auto-encoder blocks in order to obtain an efficient RFI approach. Our model can decompose the radio signal acquired from sensors as a sum of a low-rank matrix (interference signal) and a sparse signal (clear signal with targets). Convergence is achieved by solving a convex minimization problem to retrieve the clear signal, which leads to an iterative principal component pursuit [4]. Moreover, as we combine iterative algorithms, which provide a natural recurrent architecture, with residual overcomplete auto-encoders based on convolutional layers, we exploit the benefits of both recurrent and convolutional architectures in order to attain better results. In the experiments, we show that our approach surpasses both FCN models of Ristea et al. [16]. Unlike other deep models, which only predict the amplitude, our approach is also able to estimate the phase.

2. Method

2.1. Problem Formulation

For a continuous-time signal $x(t)$, let $x[n]$ denote the discrete-time signal computed by sampling $x(t)$, $x[n] = x(n \cdot T_S)$, where T_S is the sampling period.

In FMCW radar, the antenna transmits a signal $s_{TX}(t)$, which is a chirp sequence, whose frequency usually follows a saw-tooth pattern. The receive antenna collects $s_{RX}(t)$, which, in the presence of interference from another vehicle, is a mixture of two signals: the signal reflected by targets (having the chirp modulation rate identically with the transmitted signal) and the interference signal (having the chirp modulation rate different from the transmitted signal). Consequently, the received signal is defined as:

$$s_{RX}(t) = \sum_{p=0}^{N_t} A_p \cdot s_{TX}(t - t_{d,p}) + s_{chirp,RFI}(t), \quad (1)$$

where A_p and $t_{d,p}$ are the complex amplitude and the propagation delay of target p , respectively, N_t is the number of targets, and $s_{chirp,RFI}(t)$ is the interfering signal.

The received signal, $s_{RX}(t)$, is mixed with the transmitted signal, $s_{TX}(t)$, low-pass filtered and sampled, resulting in the beat signal $s_b[n]$. Therefore, $s_b[n]$ consists in a sum of complex sinusoids (representing the targets) and an uncorrelated interference $s_{b,RFI}[n]$, which is a chirp signal that is obtained by mixing two chirp signals with different modulation rates. Hence, the beat signal in the presence of

uncorrelated interference is written as:

$$s_b[n] = \sum_{p=0}^{N_t} A_p \cdot \exp(2\pi \cdot j \cdot f_p \cdot n) + s_{b,RFI}[n], \quad (2)$$

where $j^2 = -1$, $f_p = (\alpha \cdot t_{d,p}) \cdot T_S$ is the beat frequency of target p and α denotes the slope of the transmitted chirp. The uncorrelated interference appears as a highly non-stationary component on the beat signal's spectrogram, being spread across multiple frequency bins, as opposed to the signal received from targets, which is present only at some frequency values f_p [1]. Hence, we can consider the signals received from targets as narrow band components and the interference as a wide band signal.

We propose to mitigate the interference in the Fourier domain. Therefore, we first apply the Fast Fourier Transform (FFT) to the signal defined in Eq. (2), obtaining the beat signal spectrum. We consider the FFT of the signal received from targets (the sum in Eq. (2)) as a sparse signal, because there is a limited number of targets, which translates to a few amplitudes on the range profile. The interference signal, $s_{b,RFI}[n]$, is considered a type of noise, because its spectrum contains multiple frequency bins with higher amplitudes. In order to obtain a matrix representation, which is commonly used in deep learning methods, the FFT of each discrete signal is represented as a matrix of shape $N_{FFT} \times 2$, where N_{FFT} is the number of FFT points, and 2 comes from the real and imaginary parts of FFT, which are viewed as independent vectors. Consequently, we propose a data model composed of a matrix L (the FFT of the interference signal) plus a sparse matrix S (the FFT of the signal received from targets). Our data model is described in a matrix formulation as:

$$D = L + S, \quad (3)$$

where D is the received data, L is the interference data and S is the clean data with targets. The matrices D , L and S have the same dimension, namely $N_{FFT} \times 2$.

2.2. Unfolded Robust PCA

Unfolding [14], or unrolling an iterative algorithm, was first suggested by Gregor et al. [7] to accelerate convergence. They showed that by considering each iteration of an iterative algorithm as a layer in a deep network and by concatenating a few such layers, it is possible to train unfolded networks to achieve a dramatic improvement in convergence, significantly reducing the number of training iterations. In the context of RPCA, a principled way to construct learnable pursuit architectures for structured sparse and robust low-rank models was introduced in [20]. The proposed networks, derived from the iteration of proximal descent algorithms, were shown to faithfully approximate

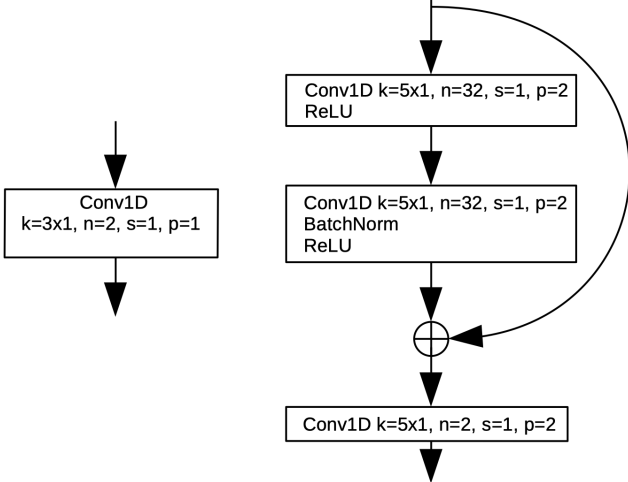


Figure 1. Convolutional layer (left) used in the unfolded RPCA of [19] versus our residual overcomplete auto-encoder block (right). The parameters are defined as follows: k is the kernel size, n is the number of filters, s is the stride and p is the padding.

the solution of RPCA, but the approach was based on a non-convex formulation in which the rank of L was assumed to be known a-priori. This poses a network design limitation, as the rank can vary between different applications or even different realizations of the same application, *i.e.* the number of targets from two signals acquired from the same radio sensor may be different. In contrast, we employ the approach proposed in [19], which does not require a-priori knowledge of the rank.

Unfolding an algorithm can be envisioned as a recurrent neural network, in which the k^{th} iteration is regarded as the k^{th} layer in a feed-forward network. Following [19], the L and S matrices for a certain step k are computed as follows:

$$\begin{aligned} L^{k+1} &= SVT_{\lambda_1^k} \{g_5^k(L^k) + g_3^k(S^k) + g_1^k(D)\}, \\ S^{k+1} &= \tau_{\lambda_2^k} \{g_6^k(L^k) + g_4^k(S^k) + g_2^k(D)\}, \end{aligned} \quad (4)$$

where the operator SVT refers to singular value thresholding and the operator τ and the regularization parameters λ_1, λ_2 are described in [19]. Each function $g_i, \forall i \in \{1, 2, \dots, 6\}$ is a transformation, which, in [19], takes the form of a convolution with a learnable kernel, and, in our approach, takes the form of a more complex function based on residual auto-encoder blocks. The parameters of each g_i are learned independently for each layer k . Although, in theory, L is supposed to be a low-rank matrix, we empirically observed that for most data samples, its rank is maximum, *i.e.* equal to 2. Nevertheless, our empirical results show that unfolded RPCA works well, even if the theoretical constraint regarding the rank of L is not met. For more details about unfolded RPCA, the reader is referred to [19].

2.3. Residual Overcomplete Auto-Encoder Blocks

In the deep formulation of unfolded RPCA, the recurrent network is based on convolutional layers. Following recent works [8, 21] showing that deeper models provide better results, we propose to replace the convolutional layers in the deep unfolded RPCA with residual overcomplete auto-encoder (ROC-AE) blocks. Our block is formed of two convolutional layers with 32 filters each, and a third convolutional layer with two filters. In a set of preliminary experiments, we observed that some weights converged to NaN values, a problem that is caused by vanishing or exploding gradients. In order to avoid this issue, we insert a skip connection from the input to the third convolutional layer. Our novel block is illustrated in Fig. 1, in comparison with the approach proposed in [19], which is based on a single layer of convolution. Both architectures use tensors of $1 \times 2048 \times 2$ components as input and output, respectively. We note that the first two convolutional layers in our block have 32 filters, generating activation maps of $1 \times 2048 \times 32$ components. Hence, our residual blocks are designed like overcomplete auto-encoders (the latent space is higher than the input space). In the experiments, we show that our residual overcomplete auto-encoder blocks significantly outperform residual undercomplete auto-encoder (RUC-AE) blocks with equivalent depth.

3. Experiments

3.1. Data Set

The automotive radar inference mitigation (ARIM) data set [16] is a large scale database consisting of 48,000 radio signal samples, synthetically generated while trying to replicate realistic automotive scenarios with one source of interference. The data set is split into a training set of 40,000 samples and a test set of 8,000 samples. We split the training data into two disjunctive sets for training (32,000 samples) and validation (8,000 samples), according to Ristea et al. [16].

Each sample is generated using randomly selected values for the following parameters: signal-to-noise ratio (SNR), signal-to-interference ratio (SIR), relative interference slope, number of targets, amplitude, phase and distance of each target. One of the biggest advantages that are provided by the ARIM data set is that we have access to clean and perturbed signal pairs. Therefore, we are able to use the model described in Section 2, as it requires access to both interference and clean target signals during training.

To our knowledge, ARIM [16] is the only large scale data set that is publicly available for the radar interference mitigation task. Therefore, we evaluate our deep learning method against other competing approaches only on ARIM.

Table 1. Validation and test results on the ARIM data set attained by various versions of unfolded RPCA versus an oracle based on true labels, zeroing and state-of-the-art FCN models [16], respectively. The best results (excluding the oracle) are highlighted in bold. The symbol \uparrow means higher values are better and \downarrow means lower values are better.

Method	Validation set				Test set				Inference time		
	Δ SNR \uparrow	AUC \uparrow	MAE \downarrow (dB)	MAE \downarrow (degrees)	Δ SNR \uparrow	AUC \uparrow	MAE \downarrow (dB)	MAE \downarrow (degrees)	CPU (ms)	GPU (ms)	NX (ms)
Oracle	12.92	0.978	0	0	13.08	0.978	0	0	-	-	-
Zeroing	5.27	0.951	1.26	6.80	5.44	0.951	1.27	6.79	<1	<1	-
Shallow FCN [16]	10.34	0.965	2.20	-	10.49	0.965	2.21	-	471	62	-
Deep FCN [16]	12.90	0.972	1.21	-	13.06	0.972	1.22	-	8400	66	-
Unfolded RPCA [19]	12.14	0.968	1.47	5.58	12.33	0.970	1.47	5.56	55	45	273
Unfolded RPCA RUC-AE	9.58	0.967	2.83	5.04	9.83	0.966	2.76	5.71	76	35	242
Unfolded RPCA ROC-AE (ours)	10.15	0.975	0.53	2.45	10.46	0.976	0.55	2.55	122	40	299

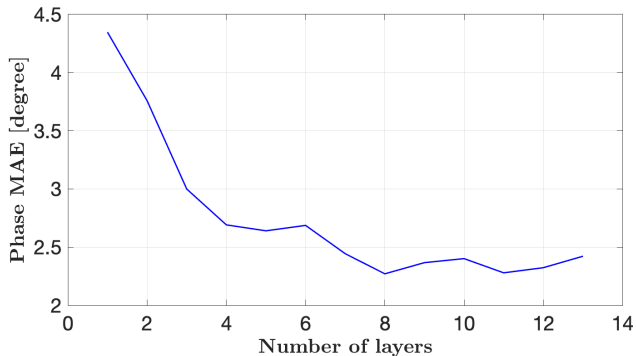


Figure 2. The influence of the number of neural network layers (from 1 to 13) on the phase MAE measured in degrees.

3.2. Evaluation Metrics

Typically, the goal in radar signal processing is to maximize the target detection performance. Therefore, an intuitive metric is the area under the receiver operating characteristics curve (AUC), which describes the ability to disentangle targets from noise at various thresholds. Another performance indicator is the mean absolute error (MAE) in decibels (dB) between the range profile amplitude of targets computed from label signals and the amplitude of targets from predicted signals. In radar signal processing, recovering the phase of targets is equally important, because it is essential in estimating other mandatory parameters, *e.g.* target velocity. Thus, we also report the MAE in degrees between the range profile phase of targets computed from label signals and the phase of targets from predicted signals. In our evaluation, we additionally report the mean SNR improvement (Δ SNR), which is computed for the target with the highest amplitude in a signal, as the difference between SNR before and after interference mitigation.

3.3. Hyperparameter Tuning

In order to minimize the chance of overfitting in hyperparameter space, we tune our hyperparameters on the validation set. The number k of unfolded network layers was validated on the evaluation set, considering values from 1 to 13. As shown in Fig. 2, the optimal value is $k = 8$. Regarding the training process, we trained our network for 100 epochs with mini-batches of 20 samples using the Adam optimizer [10] with a learning rate of $5 \cdot 10^{-4}$ and a weight decay of 10^{-6} . We also added a learning rate scheduler with a step of 30 epochs and a decay factor of 0.5.

3.4. Quantitative Results

We compare the unfolded RPCA models based on ROC-AE or RUC-AE blocks with the most common classical approach, called zeroing, two FCNs described in [16], the unfolded RPCA approach proposed in [19] and an oracle computed from the ground-truth labels. We present the corresponding results in Table 1. Our unfolded RPCA model based on ROC-AE blocks outperforms the zeroing method by a large margin, in terms of all performance measures. When comparing the model based on standard (undecomplete) AE blocks with the one based on overcomplete AE blocks, we observe that the latter model attains superior results, regardless of the metric. The unfolded RPCA based on undercomplete AE blocks attains lower scores even compared with the unfolded RPCA [19]. This clearly shows the necessity to use overcomplete AE blocks in order to obtain performance improvements. In terms of AUC and amplitude MAE, the unfolded RPCA [19] is below the FCN models [16]. The introduction of the ROC-AE blocks brings significant performance gains to the unfolded RPCA model, surpassing all models in terms of AUC, amplitude MAE and phase MAE. Even if our approach attains inferior results in terms of Δ SNR compared with both FCN models [16], the latter models cannot estimate the phase of targets. This is

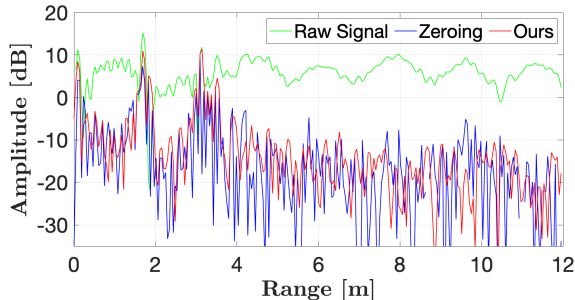


Figure 3. The beat signal spectrum for RFI mitigation results with our unfolded RPCA based on ROC-AE blocks versus the zeroing method on a real beat signal spectrum. For reference, we also added the input signal captured by the NXP TEF810X 77 GHz radar transceiver.

a major drawback of the FCN models. We emphasize that neither version of unfolded RPCA suffers from this problem.

Additionally, we observe that our model obtains lower performance in terms of the mean SNR improvement, because of the data samples having a small signal-to-interference ratio (SIR). A small SIR implies that the FFT of the signal does not meet the sparsity condition, because the targets are close to the noise floor. Having targets near or in the noise floor at training time may interfere with the proposed data modeling.

3.5. Running Time

In the radar signal domain, a key element is the capability of an algorithm to process data in real-time on low-power processing units, *e.g.* embedded devices. Therefore, we analyze the inference time of each method, *i.e.* the average time required to process a time domain signal and output the corresponding beat signal spectrum.

The CPU and GPU times reported in Table 1 were measured on a machine with Intel Core i7 CPU and NVIDIA RTX 2080Ti GPU. As expected, the zeroing method has the best inference time, but its accuracy levels are much lower compared to our deep learning approach. Moreover, we observe that our unfolded RPCA based on ROC-AE blocks is quicker than both FCNs, especially on CPU, while also offering better results. The algorithm proposed in [19] is slightly faster than our approach because of its shallower architecture.

In addition, we measured the running times of the three unfolded RPCA models on an Nvidia Jetson Xavier NX embedded system. The corresponding time measurements are reported in the last column of Table 1. We observe that our unfolded RPCA based on ROC-AE blocks requires about 22 extra milliseconds on the lower-end GPU with respect to the baseline unfolded RPCA. We conclude that our unfolded RPCA based on ROC-AE blocks provides the op-

timal trade-off between accuracy and speed. With the reported times, our approach is a viable candidate to be deployed on embedded devices placed on board road vehicles.

3.6. Qualitative Results

In addition to the results on ARIM, we assessed the generalization capacity of our approach on real data, by testing it on samples provided by the NXP company, which were captured with the NXP TEF810X 77 GHz radar transceiver, in real-world scenarios. In Fig. 3, we show an example of interference mitigation performed by our unfolded RPCA based on ROC-AE blocks on a real radar signal in comparison with the zeroing algorithm. We underline that the shown example contains three close-range targets, at roughly 0, 2 and 3 meters. Both models seem capable of reducing the noise floor, but our approach is better at estimating the targets. More precisely, the amplitude around the targets is higher for our approach compared to zeroing.

4. Conclusion

In this paper, we introduced an unfolded robust PCA model based on residual overcomplete auto-encoder blocks for automotive radar interference mitigation, which is capable of estimating both the magnitude and the phase of automotive radar signals. We compared our model with several baseline approaches in a comprehensive experiment, showing that our model provides superior results in terms of accuracy and time. We also showed the real-time processing capability of our approach, as well as its generalization capacity on real data. In future work, we aim to analyze the scenario with multiple interference sources.

Acknowledgments

The authors thank reviewers for their useful feedback leading to improvements of this work. The authors would also like to thank Adrian Şandru from SecurifAI for helping to measure the inference times of the unfolded RPCA models on Jetson NX. The authors acknowledge national funding authorities and the ECSEL Joint Undertaking, which funded the PRYSTINE project under grant 783190. This work was co-funded through the Competitiveness Operational Program 2014-2020, Axis 1, contract no. 3/1.1.3H/24.04.2019, MySMIS ID: 121784. This article has also benefited from the support of the Romanian Young Academy, which is funded by Stiftung Mercator and the Alexander von Humboldt Foundation for the period 2020-2022.

References

- [1] S. Alland, W. Stark, M. Ali, and M. Hegde. Interference in Automotive Radar Systems: Characteristics, Mitigation

- Techniques, and Current and Future Research. *IEEE Signal Processing Magazine*, 36(5):45–59, 2019.
- [2] J. Bechter, M. Rameez, and C. Waldschmidt. Analytical and experimental investigations on mitigation of interference in a dbf mimo radar. *IEEE Transactions on Microwave Theory and Techniques*, 65(5):1727–1734, 2017.
- [3] G.M. Brooker. Mutual interference of millimeter-wave radar systems. *IEEE Transactions on Electromagnetic Compatibility*, 49(1):170–181, 2007.
- [4] E.J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):1–37, 2011.
- [5] W. Fan, F. Zhou, M. Tao, X. Bai, P. Rong, S. Yang, and T. Tian. Interference Mitigation for Synthetic Aperture Radar Based on Deep Residual Network. *Remote Sensing*, 11(14):1654, 2019.
- [6] J. Fuchs, A. Dubey, M. Lübke, R. Weigel, and F. Lurz. Automotive Radar Interference Mitigation using a Convolutional Autoencoder. In *Proceedings of RADAR*, 04 2020.
- [7] K. Gregor and Y. LeCun. Learning fast approximations of sparse coding. In *Proceedings of ICML*, pages 399–406, 2010.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of CVPR*, pages 770–778, 2016.
- [9] G. Kim, J. Mun, and J. Lee. A Peer-to-Peer Interference Analysis for Automotive Chirp Sequence Radars. *IEEE Transactions on Vehicular Technology*, 67(9):8110–8117, 2018.
- [10] D.P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proceedings of ICLR*, 2015.
- [11] M. Kunert. The EU project MOSARIM: A general overview of project objectives and conducted work. In *Proceedings of EuRAD*, pages 1–5, 2012.
- [12] M. Kunert, F. Bodereau, M. Goppelt, C. Fischer, A. John, T. Wixforth, A. Ossowska, T. Schipper, and R. Pietsch. D1.5 - Study on the state-of-the-art interference mitigation technique, MOre Safety for All by Radar Interference Mitigation (MOSARIM) project. Technical report, Robert Bosch GmbH, 2010.
- [13] F. Laghezza, F. Jansen, and J. Overvest. Enhanced Interference Detection Method in Automotive FMCW Radar Systems. In *Proceedings of IRS*, pages 1–7, 2019.
- [14] V. Monga, Y. Li, and Y.C. Eldar. Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing. *IEEE Signal Processing Magazine*, 2020.
- [15] J. Mun, H. Kim, and J. Lee. A Deep Learning Approach for Automotive Radar Interference Mitigation. In *Proceedings of VTC-Fall*, pages 1–5, 2018.
- [16] N.C. Ristea, A. Anghel, and R.T. Ionescu. Fully Convolutional Neural Networks for Automotive Radar Interference Mitigation. In *Proceedings of VTC-Fall*, 2020.
- [17] 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 *Proceedings of FUSION*, 2019.
- [18] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proceedings of MICCAI*, pages 234–241. Springer, 2015.
- [19] O. Solomon, R. Cohen, Y. Zhang, Y. Yang, Q. He, J. Luo, R.J.G. van Sloun, and Y.C. Eldar. Deep unfolded robust pca with application to clutter suppression in ultrasound. *IEEE Transactions on Medical Imaging*, 39(4):1051–1063, 2019.
- [20] P. Sprechmann, A.M. Bronstein, and G. Sapiro. Learning efficient sparse and low rank models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1821–1833, 2015.
- [21] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going Deeper With Convolutions. In *Proceedings of CVPR*, pages 1–9, June 2015.
- [22] F. Uysal. Synchronous and Asynchronous Radar Interference Mitigation. *IEEE Access*, 7:5846–5852, 2019.
- [23] Z. Xu and Q. Shi. Interference Mitigation for Automotive Radar Using Orthogonal Noise Waveforms. *IEEE Geoscience and Remote Sensing Letters*, 15(1):137–141, 2018.