

# MODEL-BASED DIFFUSION FOR MITIGATING AUTOMOTIVE RADAR INTERFERENCE

*J. Overdevest<sup>§\*</sup>, X. Wei<sup>§†</sup>, H. van Gorp<sup>†</sup>, R.J.G. van Sloun<sup>†</sup>*

<sup>\*</sup>Signal Processing Algorithms, NXP Semiconductors, Eindhoven, The Netherlands

<sup>†</sup> Dept. of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

## ABSTRACT

Mitigating automotive radar-to-radar interference is a challenging task, especially when the observed signal is densely corrupted with highly correlated interference signals. In this paper, we propose to remove this interference using joint-conditional posterior sampling with score-based diffusion models. These models use three individual scores: a target score, an interference score, and a joint data consistency score. Leveraging the sparsity of clean target signals in the Fourier domain, we propose a model-based score estimator for the target signals, derived from the proximal step of the  $\ell_1$ -norm. For the interference score, we use a neural network with denoising score-matching, given that it is difficult to obtain analytical statistical models of the interference signals. Lastly, the target and interference scores are connected by a data-consistency score. Experimental results show that our solution results in superior performance over state-of-the-art methods, in terms of normalized mean squared error (NMSE) and receiver operating characteristic (ROC) curves.

**Index Terms**— Automotive Radar Interference, Deep Learning, Diffusion Models

## 1. INTRODUCTION

Automotive inter-radar interference has been intensively studied due to the severe impact it can have on driving performance, especially with more radars being deployed on vehicles in the next years. Avoiding interference altogether is the first attempt to reduce its impact; either in a coordinated or uncoordinated fashion [1]. However, once the interference enters the receiver front-end, mitigation using signal processing is required. In the digital domain, many degrees of freedom can be exploited to reduce the impact, i.e. time [2], frequency [3], time-frequency [4], and/or space [5, 6]. Besides, several deep neural networks have been proposed [7, 8], of which some exploit a signal prior [9]. While Bayesian algorithms have achieved promising performance, this performance hinges on accurate signal priors. The adoption of highly expressive generative models is a promising direction to model these signal priors.

Score-based diffusion models, a subset of generative probabilistic models, have demonstrated the capability to

approximate complex data distributions [10, 11]. They are state-of-the-art for image generation as well as inverse problem solving across a variety of applications [12, 13, 14]. Stevens *et al.* recently proposed to remove structured noise from natural images [15] and ultrasound data [16] using a joint conditional diffusion process.

In this paper, we propose a joint-conditional probabilistic diffusion model that can effectively separate radar interference signals from target signals in the time domain, specifically designed to remove highly correlated interference, as well as densely corrupted Frequency-Modulated Continuous Wave (FMCW) chirp signals. To our knowledge, we are the first to apply diffusion models to radar signals. For interference mitigation, we adopt two separate scores to incorporate two independent signal priors. Firstly, a novel model-based score is introduced within a diffusion model, called a proximal score function, that is designed for sparse target signals in the Fourier domain, offering denoising in the complex domain with reduced computational complexity. Secondly, a learned score function for interference signals exploits the signal structure in the time domain. We demonstrate that our approach outperforms other one-dimensional state-of-the-art reconstruction methods [17, 18].

## 2. SIGNAL MODEL

In automotive radar, the received signals can consist of target reflections (desired), inter-vehicle interference (undesired), and thermal noise (undesired). Given that FMCW is the most common waveform in automotive radar applications, we shall focus our signal model to FMCW radars. For a single receive antenna and a single chirp, the sampled signal  $\mathbf{y} \in \mathbb{R}^N$  can be expressed as [17]:

$$\mathbf{y} = \sum_k \mathbf{x}^{(k)} + \sum_l \mathbf{i}^{(l)} + \mathbf{n} = \mathbf{x} + \mathbf{i} + \mathbf{n}, \quad (1)$$

where  $\mathbf{x}^{(k)} = \alpha_k \sin(\omega_k \mathbf{v} + \phi_k)$  with  $\mathbf{v} = [0, \dots, M-1]$  denotes the  $k$ -th target reflection with  $\alpha_k$ ,  $\omega_k$  and  $\phi_k$  representing the received signal amplitude, its frequency and phase, respectively, and  $M$  is the number of fast-time samples. Moreover,  $\mathbf{i}^{(l)}$  refers to the  $l$ -th interference signal which is generally known to be time-limited after analog down-mixing and anti-aliasing filtering. Lastly,  $\mathbf{n}$  is the thermal noise, modeled using a zero-mean Gaussian distribution.

<sup>§</sup>Equal Contribution

## 2.1. Source Separation Approach

Our goal is to mitigate the interference  $\mathbf{i}$  while preserving the targets  $\mathbf{x}$ , which we propose to solve using a source separation approach. Given the ill-posed nature of the problem, and assuming that the target and interference signals are independent with distinct data distributions, we introduce priors and sample from the joint conditional posterior using:

$$\mathbf{x}, \mathbf{i} \sim p(\mathbf{x}, \mathbf{i} | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{x}, \mathbf{i}) \cdot p(\mathbf{x}) \cdot p(\mathbf{i}), \quad (2)$$

where  $p(\mathbf{x}, \mathbf{i} | \mathbf{y})$  is the joint conditional posterior distribution. Additionally,  $p(\mathbf{x})$  is the prior for the targets, which are known to be sparse in the Fourier domain. Besides,  $p(\mathbf{i})$  represents the prior for the interference, explicitly modeled in the time domain using a deep generative network. Lastly,  $p(\mathbf{y} | \mathbf{x}, \mathbf{i})$  denotes the joint data consistency term. We now propose a three-stage approach to get samples from the joint conditional posterior. First, we define separate priors for interference and targets. Subsequently, we incorporate the priors into a diffusion process. Finally, we derive the estimated interference and subtract it from the noisy observations. This aids in preserving the weak targets buried under the noise floor, setting the stage for their recovery in subsequent signal processing steps.

## 3. METHOD

Diffusion models are generative models that evolve data in a forward diffusion process, ultimately leading to random Gaussian noise. This process can be represented by a stochastic differential equation (SDE):

$$d\mathbf{x} = f(t)\mathbf{x} + g(t)d\mathbf{w}, \quad (3)$$

where  $f(t)$  and  $g(t)$  are the drift and diffusion coefficients, respectively, at time step  $t$ . Here, we choose to utilize Variance Preserving (VP) SDE for perturbing both the targets and interference. Given a clean sample  $\mathbf{x}_0$  from the target distribution, the perturbed data at time step  $t$  can be represented as:  $\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\mathbf{z}$  with  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$  with  $\alpha_t$  being a scalar determined from  $f(t)$  and  $g(t)$ .

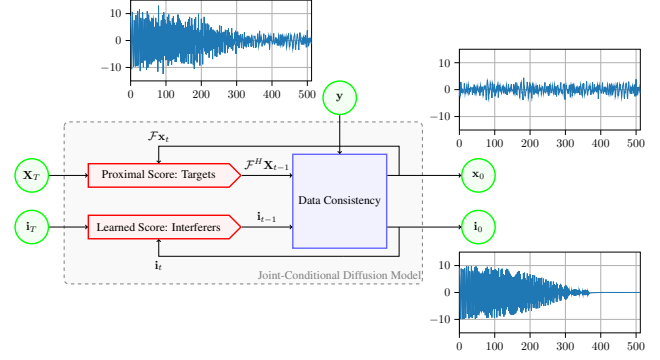
Time-reversing the process enables us to generate novel samples from the target data distribution. The generation relies on the gradient of the log-probability density function, commonly known as the *score*, and can be analytically expressed through a reverse-time SDE. For joint conditional data generation, the corresponding reverse-time SDE can be formulated as follows:

$$d(\mathbf{x}_t, \mathbf{i}_t) = \left[ f(t)(\mathbf{x}_t, \mathbf{i}_t) - g(t)^2 \nabla_{\mathbf{x}_t, \mathbf{i}_t} \log p(\mathbf{x}_t, \mathbf{i}_t | \mathbf{y}) \right] dt + g(t) d\mathbf{w}_t. \quad (4)$$

However,  $\nabla_{\mathbf{x}_t, \mathbf{i}_t} \log p(\mathbf{x}_t, \mathbf{i}_t | \mathbf{y})$  is not directly solvable. To address it, we resort to (2) and take the derivative with respect to  $\mathbf{x}_t$  or  $\mathbf{i}_t$ . This results in the following equations:

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t, \mathbf{i}_t | \mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t, \mathbf{i}_t), \quad (5)$$

$$\nabla_{\mathbf{i}_t} \log p(\mathbf{x}_t, \mathbf{i}_t | \mathbf{y}) = \nabla_{\mathbf{i}_t} \log p(\mathbf{i}_t) + \nabla_{\mathbf{i}_t} \log p(\mathbf{y} | \mathbf{x}_t, \mathbf{i}_t). \quad (6)$$



**Fig. 1:** Block diagram of the proposed joint-conditional diffusion model.

Consequently, for generating interference-free signals via the joint conditional reverse-time SDE, it is essential to obtain knowledge of the scores of the targets, interference, and joint data consistency, which we address in Sections 3.1, 3.2, and 3.3. We refer to the procedure as a joint conditional diffusion process. See Fig. 1 for an overview.

### 3.1. Proximal Score Function: Targets

The clean target signal  $\mathbf{x}$  is a sparse superposition of sinusoidal signals, as previously noted in (1). Therefore, we propose to model the prior belief of  $\mathbf{X} = \mathcal{F}\mathbf{x}$  as a Laplace distribution. Using this prior knowledge, we can calculate the score function of the target signal  $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$  from (5) more accurately in the Fourier domain. Moreover, we can define the score using Tweedie's approximation [19]:

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{X}_t) = \frac{\sqrt{\bar{\alpha}_t} \mathbf{X}_{0|t} - \mathbf{X}_t}{1 - \bar{\alpha}_t}, \quad (7)$$

where  $\mathbf{X}_{0|t}$  denotes the posterior mean  $\mathbb{E}[\mathbf{X}_0 | \mathbf{X}_t]$ , i.e., a denoised estimate from  $\mathbf{X}_t$  at time step  $t$ . Assuming independent Laplacian priors, the denoising problem is solved using the Maximum A Posterior (MAP) estimator, hence the denoised estimate  $\mathbf{X}_{0|t}$  is found from

$$\mathbf{X}_{0|t} = \arg \min_{\mathbf{X}} \frac{1}{1 - \bar{\alpha}_t} \|\mathbf{X}_t - \mathbf{X}\|_2^2 + \lambda \|\mathbf{X}\|_1, \quad (8)$$

where  $1 - \bar{\alpha}_t$  refers to noise variance at time step  $t$  and  $\lambda$  is a scalar that balances the reconstruction error against sparsity. The solution to (8) is known to be the proximal step, defined by the soft threshold operator

$$\mathbf{X}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}} \text{prox}_{\lambda \|\cdot\|_1}(\mathbf{X}_t) = \frac{1}{\sqrt{\bar{\alpha}_t}} \frac{\mathbf{X}_t}{\|\mathbf{X}_t\|_1} (|\mathbf{X}_t| - \lambda)_+, \quad (9)$$

where the threshold  $\lambda$  is determined from the noise variance and the Laplace scale parameter  $\bar{\sigma}_L$  at time step  $t$ , hence  $\lambda = \frac{\sqrt{2(1-\bar{\alpha}_t)}}{\sqrt{\bar{\alpha}_t} \bar{\sigma}_L}$ . Unlike current real-valued diffusion models, our proximal score enables denoising complex input data, i.e. Fourier transformed data. Then,  $\mathbf{X}_{0|t}$  is used to retrieve the perturbed estimate of the previous time step  $\tilde{\mathbf{x}}_{t-1}$ .

### 3.2. Learned Score Function: Interference

Given the challenging nature of modeling the interference signal, we employ a neural network to estimate the interference's score function  $\mathbf{s}_\theta(\mathbf{i}_t, t) \approx \nabla_{\mathbf{i}_t} \log p(\mathbf{i}_t)$ , which can be injected into (6). Here, we adopt a conditional U-Net architecture, following [20], while with modifications that incorporate self-attention for better temporal correlation estimation. During training, we perturb the clean interference  $\mathbf{i}_0 \sim p_0(\mathbf{i})$  to obtain the noisy samples  $\mathbf{i}_t \sim q(\mathbf{i}_t|\mathbf{i}_0)$ , which are then fed into the network. The score function is estimated by employing denoising score matching [21], as formulated below:

$$\theta^* = \arg \min_{\theta} \mathbb{E}_t \left[ \mathbf{w}_t \mathbb{E}_{\mathbf{i}_0} \mathbb{E}_{\mathbf{i}_t|\mathbf{i}_0} \left\| \mathbf{s}_\theta(\mathbf{i}_t, t) - \nabla_{\mathbf{i}_t} \log q(\mathbf{i}_t|\mathbf{i}_0) \right\|_2^2 \right], \quad (10)$$

where  $\theta$  denotes the learnable parameters within the neural network,  $\mathbf{w}_t$  is a positive weighting function,  $t$  is uniformly sampled over  $[10^{-5}, 1]$ . The objective of the training is to minimize the weighted squared error between the estimated score by the network and the score of the perturbed interference, sampled at various  $t$ . Guided by this objective, the network functions as a denoiser, steering towards the direction of clean samples by connecting the manifolds at different time steps. After adequately training the network, similar to (7), the posterior mean of the interference signal  $\mathbf{i}_{0|t}$  can be described by exploiting Tweedie's approximation:

$$\mathbf{i}_{0|t} \approx \frac{1}{\sqrt{\bar{\alpha}_t}} \left( \mathbf{i}_t + (1 - \bar{\alpha}_t) \mathbf{s}_\theta(\mathbf{i}_t, t) \right), \quad (11)$$

which implies that for a perturbed sample  $\mathbf{i}_t$ , the learned score allows to retrieve its noise-reduced version,  $\mathbf{i}_{0|t}$ . Analogous to the target signals, during inference, this can be used to obtain the perturbed estimate of the previous time step  $\tilde{\mathbf{i}}_{t-1}$ .

### 3.3. Joint Data Consistency Score

The score of the joint data consistency in (5) and (6) is typically intractable due to its dependence on the perturbed signals at time  $t$ . Instead, we employ a method known as Diffusion Posterior Sampling (DPS) as a surrogate, as described in [13]. Considering our measurement model, as described in (1), adheres to a Gaussian distribution with a zero mean and a standard deviation of  $\sigma_n$ , we can establish the following approximation:

$$p(\mathbf{y}|\mathbf{x}_t, \mathbf{i}_t) \approx p(\mathbf{y}|\mathbf{x}_{0|t}, \mathbf{i}_{0|t}) \sim \mathcal{N}(\mathbf{x}_{0|t} + \mathbf{i}_{0|t}, \sigma_n^2), \quad (12)$$

where  $\mathbf{x}_{0|t} = \mathcal{F}^H \mathbf{X}_{0|t}$  and  $\mathbf{i}_{0|t}$  are obtained using (9) and (11), respectively. To this end, the score of the joint data consistency turns tractable and can be formulated as follows:

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t, \mathbf{i}_t) \approx -\frac{1}{\sigma_n^2} \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{x}_{0|t} - \mathbf{i}_{0|t}\|_2^2 \quad (13)$$

$$\nabla_{\mathbf{i}_t} \log p(\mathbf{y}|\mathbf{x}_t, \mathbf{i}_t) \approx -\frac{1}{\sigma_n^2} \nabla_{\mathbf{i}_t} \|\mathbf{y} - \mathbf{x}_{0|t} - \mathbf{i}_{0|t}\|_2^2 \quad (14)$$

By integrating the score functions of the targets, interference and joint data consistency within an iterative framework, we establish a comprehensive flow for a joint conditional process. See the algorithm below for details. In this paper, the estimated target signal  $\hat{\mathbf{x}} = \mathbf{y} - \mathbf{i}_0$  is used for evaluation.

---

#### Algorithm 1 Joint Conditional Posterior Sampling

---

**Input:**  $N, \mathbf{y}, \mathbf{s}_\theta, \bar{\sigma}_L, \varepsilon_1, \varepsilon_2$   
**Initialization:**  $\mathbf{x}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \mathbf{i}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$   
1: **for**  $i = N - 1$  **to** 0 **do**  
2:    $t \leftarrow \frac{i+1}{N}$   
3:    $\lambda \leftarrow \frac{\sqrt{2(1-\bar{\alpha}_t)}}{\sqrt{\bar{\alpha}_t} \bar{\sigma}_L}$   
4:    $c_1 \leftarrow \frac{\sqrt{\bar{\alpha}_t(1-\bar{\alpha}_{t-1})}}{1-\bar{\alpha}_t}, c_2 \leftarrow \frac{\sqrt{\bar{\alpha}_{t-1}\beta_t}}{1-\bar{\alpha}_t}, c_3 \leftarrow \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}$   
5:    $\mathbf{z} \leftarrow \mathcal{N}(\mathbf{0}, \mathbf{I})$   
   ▷ Calculate Posterior mean  
6:    $\mathbf{X}_{0|t} \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}} \text{prox}_{\lambda \|\cdot\|}(\mathcal{F} \mathbf{x}_t)$   
7:    $\mathbf{x}_{0|t} \leftarrow \mathcal{F}^H \mathbf{X}_{0|t}$   
8:    $\mathbf{i}_{0|t} \leftarrow \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{i}_t + (1 - \bar{\alpha}_t) \mathbf{s}_\theta(\mathbf{i}_t, t))$   
9:    $\tilde{\mathbf{x}}_{t-1} \leftarrow c_1 \mathbf{x}_t + c_2 \mathbf{x}_{0|t} + c_3 \mathbf{z}$   
10:    $\tilde{\mathbf{i}}_{t-1} \leftarrow c_1 \mathbf{i}_t + c_2 \mathbf{i}_{0|t} + c_3 \mathbf{z}$   
   ▷ Data consistency steps  
11:    $\mathbf{x}_{t-1} \leftarrow \tilde{\mathbf{x}}_{t-1} - \varepsilon_1 \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{x}_{0|t} - \mathbf{i}_{0|t}\|_2^2$   
12:    $\mathbf{i}_{t-1} \leftarrow \tilde{\mathbf{i}}_{t-1} - \varepsilon_2 \nabla_{\mathbf{i}_t} \|\mathbf{y} - \mathbf{x}_{0|t} - \mathbf{i}_{0|t}\|_2^2$   
13: **end for**  
14: **return**  $\hat{\mathbf{x}} = \mathbf{y} - \mathbf{i}_0$

---

## 4. RESULTS


### 4.1. Experimental Setup

The radar data  $\mathbf{y}$  is simulated for  $K \in \{1, 2, \dots, 30\}$  targets, each target having a Signal-to-Noise Ratio (SNR) between  $-30$  dB and  $20$  dB. The number of interferers is limited to one, which is configured to have an Interference-to-Noise Ratio (INR) between  $0$  dB and  $20$  dB and randomly configured interference chirp settings, i.e., chirp slope, chirp rate, RF chirp bandwidth, chirp duration, carrier frequency, etc., such that a broad range of interference samples are achieved.

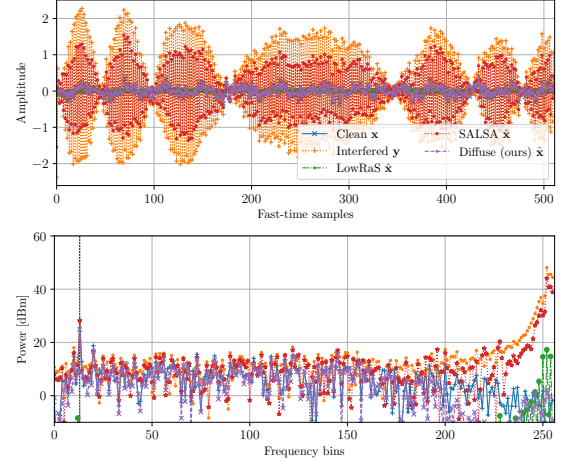
First, we trained the interference score function using a training dataset composed of 7,000 clean interference samples. The INR was randomly sampled from a range of  $0$  dB to  $20$  dB. For training, the batch size was set to 100 and training was stopped after convergence, which generally required a minimum of 500 epochs. For optimization, we employ the Adam optimizer with a learning rate set at  $10^{-4}$ . Subsequently, the score functions were integrated into an iterative framework. The hyperparameters  $\lambda, \varepsilon_1$ , and  $\varepsilon_2$  were fine-tuned based on the validation set.

### 4.2. Evaluation Metrics and Discussion

To emphasize the effectiveness of the proposed technique compared to state-of-the-art methods, we estimate the reconstruction performance in the time domain using the Normalized Mean Squared Error (NMSE) of  $\hat{\mathbf{x}}$  by  $10 \log_{10} \left( \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2}{\|\mathbf{x}\|_2^2} \right)$ ,



**Fig. 2:** Reconstruction for  $\text{SNR} = -20$  dB and  $\text{INR} = 0$  dB (top) Time-domain signal  $y$ , (bottom) Distance spectrum  $Y$ .



**Fig. 3:** Reconstruction for  $\text{SNR} = -20$  dB and  $\text{INR} = 20$  dB (top) Time-domain signal  $y$ , (bottom) Distance spectrum  $Y$ .


**Table 1:** Reconstruction Performance, measured in NMSE

Method \ INR	SNR = -20 dB			SNR = 0 dB		
	0 dB	10 dB	20 dB	0 dB	10 dB	20 dB
No mitigation	6.381	66.97	651.8	0.412	4.487	44.82
Zeroing	0.611	1.375	8.621	0.322	0.454	2.789
SALSA [17]	2.179	20.89	198.4	<b>0.247</b>	1.425	13.86
LowRaS [18]	4.143	39.82	402.8	0.677	2.953	25.13
<b>Diffuse (ours)</b>	<b>0.159</b>	<b>0.252</b>	<b>0.286</b>	0.394	<b>0.361</b>	<b>0.337</b>

and the detection performance using ROC curves, implemented with Ordered-Statistic Constant False Alarm Rate (OS-CFAR) on the distance spectrum  $\hat{\mathbf{X}}$ . Our method is compared to other source separation solutions that exploit sparsity in transformed domains (frequency, and STFT resp.) [17], or in Hankelized matrices [18].

In Fig. 2 and 3, we present two examples of real-valued time-domain signal  $\hat{\mathbf{x}}$  and the positive frequency components of  $\mathbf{F}\hat{\mathbf{x}}$  to substantiate the reconstruction performance of the proposed method in realistic interference scenarios, an uncorrelated and semi-correlated case, respectively. Despite the low SNR-valued target signal, we can accurately remove the interference and recover the distance spectrum. In Tab. 1, statistically grounded results for various SNR and INR values are provided. Hereby, we prove that our method significantly outperforms other methods for all cases, except when  $\text{SNR} = 0$  dB and  $\text{INR} = 0$  dB. Here, the learned interference score experiences difficulties in separating low-amplitude interference signals from target signals with similar power, which we intend to address in future work.

In Fig. 4, we show the ROC curves to compare various algorithms for a single target and single interferer by fixing SNR and INR. The detection performance of our method is close to the performance of the clean signal, but it significantly outperforms the other methods for all preset false-alarm rates  $P_{FA}$  of the detector.



**Fig. 4:** ROC curves for  $\text{SNR} = -20$  dB with 1 interferer with (top)  $\text{INR} = 0$  dB and (bottom)  $\text{INR} = 20$  dB.

## 5. CONCLUSIONS

We present a novel and flexible algorithm to address the interference mitigation problem as an inverse problem. Our method leverages a complex-valued proximal score function that implicitly serves as a target signal prior. In addition, we utilize a data-driven score, to accurately represent the signal structure of the interference. Moreover, leveraging the core of the diffusion process, we integrate the score functions to enable joint conditional posterior sampling. Empirically, we demonstrate that this algorithm surpasses other leading reconstruction methods utilized in FMCW automotive radar. While the paper emphasizes removing FMCW radar interference, the proposed approach is flexible and thus may be adapted to train on various interfering waveforms, i.e., Phase-Modulated Continuous Wave, Orthogonal Frequency Domain Multiplexing, and more.

## 6. REFERENCES

- [1] Lizette Lorraine Tovar Torres, Timo Grebner, and Christian Waldschmidt, "Automotive radar interference avoidance strategies for complex traffic scenarios," in *2023 IEEE Radar Conference (RadarConf23)*, 2023, pp. 1–6.
- [2] Francesco Laghezza, Feike Jansen, and Jeroen Overdevest, "Enhanced interference detection method in automotive FMCW radar systems," in *2019 20th International Radar Symposium (IRS)*, 2019, pp. 1–7.
- [3] Jonathan Bechter, Fabian Roos, Mahfuzur Rahman, and Christian Waldschmidt, "Automotive radar interference mitigation using a sparse sampling approach," in *2017 European Radar Conference (EURAD)*, 2017, pp. 90–93.
- [4] Sharef Neemat, Oleg Krasnov, and Alexander Yarovoy, "An interference mitigation technique for FMCW radar using beat-frequencies interpolation in the STFT domain," *IEEE Transactions on Microwave Theory and Techniques*, vol. 67, no. 3, pp. 1207–1220, 2019.
- [5] C. Fischer, M. Goppelt, H.-L. Blöcher, and J. Dickmann, "Minimizing interference in automotive radar using digital beamforming," *Advances in Radio Science*, vol. 9, pp. 45–48, 2011.
- [6] Sian Jin, Pu Wang, Petros Boufounos, Philip V. Orlik, Ryuhei Takahashi, and Sumit Roy, "Spatial-domain interference mitigation for slow-time MIMO-FMCW automotive radar," in *2022 IEEE 12th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, 2022, pp. 311–315.
- [7] Shengyi Chen, Wangyi Shangguan, Jalal Taghia, Uwe Kühnau, and Rainer Martin, "Automotive radar interference mitigation based on a generative adversarial network," in *2020 IEEE Asia-Pacific Microwave Conference (APMC)*, 2020, pp. 728–730.
- [8] Nicolae-Cătălin Ristea, Andrei Anghel, and Radu Tudor Ionescu, "Fully convolutional neural networks for automotive radar interference mitigation," in *2020 IEEE 92nd Vehicular Technology Conference (VTC2020-Fall)*, 2020, pp. 1–5.
- [9] J. Overdevest, A.G.C. Koppelaar, M.J.G. Bekooij, J. Youn, and R.J.G. van Sloun, "Signal reconstruction for fmcw radar interference mitigation using deep unfolding," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [10] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole, "Score-based generative modeling through stochastic differential equations," *ICLR 2021*, 2021.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [12] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon, "Solving inverse problems in medical imaging with score-based generative models," *arXiv:2111.08005*, 2021.
- [13] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye, "Diffusion posterior sampling for general noisy inverse problems," in *The Eleventh International Conference on Learning Representations*, 2023.
- [14] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz, "Pseudoinverse-guided diffusion models for inverse problems," in *International Conference on Learning Representations*, 2022.
- [15] Tristan S. W. Stevens, Hans van Gorp, Faik C. Meral, Jun Seob Shin, Jason Yu, Jean-Luc Robert, and Ruud J. G. van Sloun, "Removing structured noise with diffusion models," *arXiv:2302.05290*, 2023.
- [16] Tristan S. W. Stevens, Faik C. Meral, Jason Yu, Iason Z. Apostolakis, Jean-Luc Robert, and Ruud J. G. van Sloun, "Dehazing ultrasound using diffusion models," *arXiv:2307.11204*, 2023.
- [17] Faruk Uysal, "Synchronous and asynchronous radar interference mitigation," *IEEE Access*, vol. 7, pp. 5846–5852, 2019.
- [18] Jianping Wang, Min Ding, and Alexander Yarovoy, "Interference mitigation for fmcw radar with sparse and low-rank hankel matrix decomposition," *IEEE Transactions on Signal Processing*, vol. 70, pp. 822–834, 2022.
- [19] Bradley Efron, "Tweedie's formula and selection bias," *Journal of the American Statistical Association*, vol. 106, no. 496, pp. 1602–1614, 2011.
- [20] Simon Rouard and Gaëtan Hadjeres, "Crash: Raw audio score-based generative modeling for controllable high-resolution drum sound synthesis," *ArXiv*, vol. abs/2106.07431, 2021.
- [21] Pascal Vincent, "A connection between score matching and denoising autoencoders," *Neural Computation*, vol. 23, no. 7, pp. 1661–1674, 2011.