DEEP LEARNING PROVIDES A PROMISING APPROACH TO DENOISING MARS RADAR DATA.

C. C. Amos[1], M. R. Perry[1], and N. E. Putzig[1], [1]Planetary Science Institute (camos@psi.edu).

Introduction: Electromagnetic interference (EMI) is one of several factors that contribute noise to radar sounding data collected by the Shallow Radar (SHARAD) sounder aboard the Mars Reconnaissance Orbiter, and EMI is currently addressed via automated notch filtering [1]. Various deep learning (DL) methods have been successfully applied to a wide range of denoising problems including speech recognition [2], electrocardiograms [3], and gravitational wave detection [4]. In this work we use the Pix2Pix conditional generative adversarial network (GAN) DL architecture as described by Isola et al. [5] to translate noisy radar spectrograms into noise-free spectrograms. This effort provides the first steps in determining if DL denoising techniques could provide improved EMI noise reduction as compared to notch filtering for SHARAD data processing, allowing more accurate interpretation of Mars' subsurface properties.

Methods: The project workflow is: 1) add arbitrary synthetic noise to a set of fully processed SHARAD radar frames, 2) create spectrograms of the original and noisy radar frames, 3) train a DL model to translate a noisy spectrogram to a noise-free spectrogram, and 4) convert the denoised spectrogram back to a time-domain radar frame. Data for this study consist of >30,000 radar frames taken from a ~250 km x 250 km area over the north polar layered deposits (NPLD) on Mars. This area contains a vertical sequence of subsurface reflectors [6] enabled by low radar attenuation in ice, which allows us to evaluate denoising methods over greater depths than in dominantly ice-free areas.

As this work represents a proof-of-concept study, we chose arbitrary synthetic noise parameters rather than performing a time-intensive detailed noise analysis of SHARAD data to inform the model training dataset. We added two synthetic random noise bands to radar frames at 3.00-3.01 MHz and 6.0-6.1 MHz, allowing slight amplitude and phase variation between individual radar frames. After adding the synthetic noise, we took the absolute values of short-term fourier transforms (STFT)[7] to produce magnitude spectrograms of both the original and noisy signals. We chose these frequency-domain representations of the signals under the assumption that band-limited noise will be easier for a DL model to characterize in the frequency domain than in the time domain. Our choice of magnitude spectrograms as DL model inputs avoids the introduction of complex numbers by STFT and allows the use of well-established DL network architectures such as Pix2Pix. Figure 1 presents an example of these magnitude spectrograms.

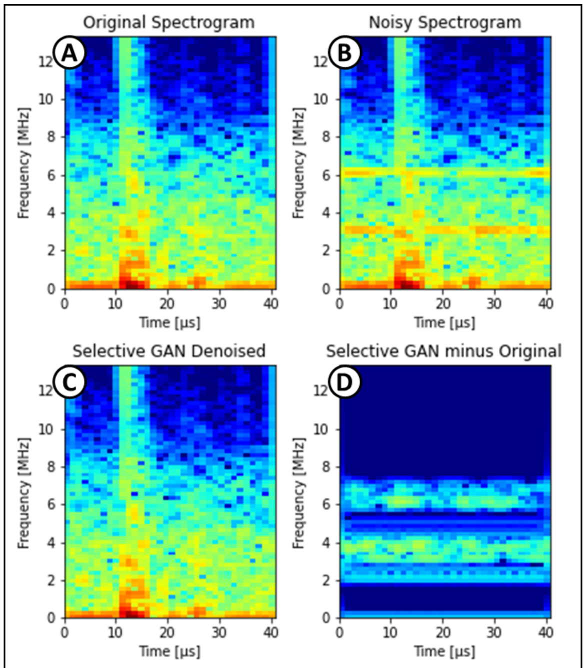

Figure 1. A) Example magnitude spectrogram from a single radar frame, taken from Observation 1308201. B) With synthetic noise added. C) Denoised magnitude spectrogram from the selective Gan workflow. D) Difference between the denoised and original spectrograms.

Once the dataset was prepared, we used 30,000 original and noisy spectrogram pairs to train a DL model, withholding an additional set of 1000 data pairs from training for blind validation. After 500 epochs (the number of times the model processes the entire dataset) of model training, we evaluated model results every 50 epochs by predicting denoised spectrograms from the blind dataset and comparing to the original noise-free data. Our evaluation consisted of computing both the mean squared error (MSE) and the cosine similarity between the time-domain signals created from the magnitude spectrograms (method described below), and we stopped training at 1000 epochs when these metrics plateaued, which indicates that no further accuracy improvement was being achieved by the DL model.

To create a time domain signal from a magnitude spectrogram, the phase of the signal must be estimated. Here we use a variant of the Griffin-Lim Algorithm (GLA)[8] commonly used in speech and audio recognition workflows to estimate the phase of the denoised model spectrograms and create time-domain denoised radar frames. As a modification, we used the phase information from the noisy signal as initial input to the iterative GLA rather than the typical method of using a random phase as initial input. This modification significantly improved the accuracy and stability of the GLA output signals.

As a first analysis step, we compared the blind GAN-denoised radar frames to those produced from a notch-filtering approach designed to remove the synthetic noise bands described earlier (Figure 2). The performances in reconstructing the noise-free signal of the GAN and notch filtering approaches are quite similar based on the evaluation metrics. To improve results of the DL approach, we developed a selective GAN (sGAN) workflow. In this approach, we first compute fast fourier transforms of the input noisy signal and the output GAN-denoised signal. For the resulting two spectra, we set a threshold amplitude difference to identify noisy frequencies. We then combine the two spectra by selecting the output GAN-denoised spectrum only for the noisy frequencies, with the remainder of the spectrum taken from the input signal. Finally, we take the inverse fast fourier transform of the combined spectrum to obtain the sGAN time domain denoised radar frame.

Discussion: As demonstrated in Figure 2, the sGAN workflow described here almost perfectly reconstructs a noise-free time-domain signal within the parameters of this proof-of-concept study. The modified GLA method produces accurate phase estimations from magnitude spectrograms, enabling the use of well-established DL networks by avoiding the use of complex numbers. Selectively choosing GAN denoised data in the frequency domain reduces the pressure for a DL network to accurately predict every frequency at every time step in the denoised signal, yielding sGAN-denoised radar frames with lower MSE compared to notch filtering.

Future work should include a detailed evaluation of EMI noise within SHARAD data. While this study presents promising results by removing relatively stationary arbitrary synthetic noise from radar frames, a more realistic model training dataset would allow the sGAN approach to be compared against the automated notch filtering approach [1], which is designed to identify and remove variable noisy frequencies and is the current standard for SHARAD data products available in NASA's Planetary Data System. Additional DL model architectures should also be tested for improved performance or robustness, such as complex-valued networks which would avoid the phase estimation step presented in this study. Lastly, we note that the sGAN approach may have applications beyond EMI noise removal and that it may be applied to any definable band-limited noise which could be modelled into a robust training dataset.

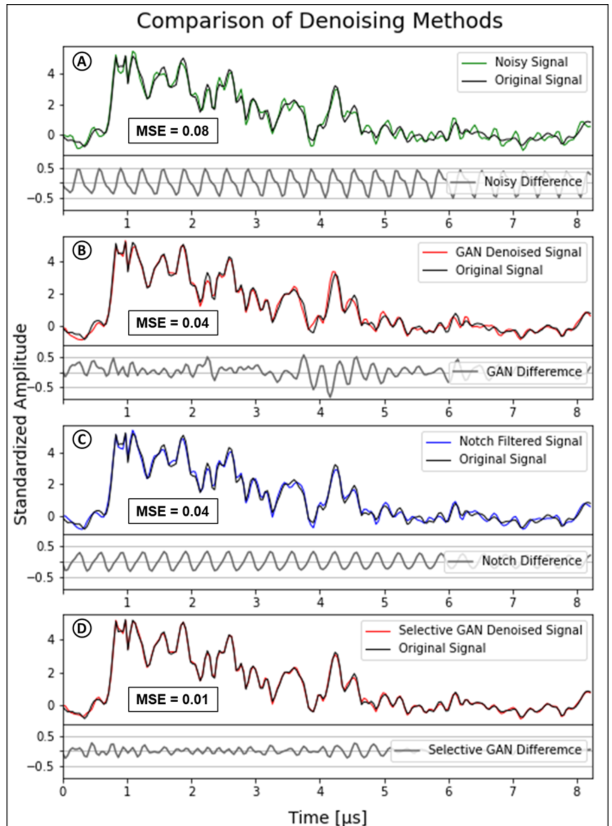

Figure 2. Comparing denoising methods for a single radar frame, surface return of the NPLD is shown at ~0.9μs. Mean squared error (MSE) values indicates similar performance between the GAN (B) and notch filter (C) results. The selective GAN (D) shows the best results, with the smallest amplitude difference compared to the original signal

Acknowledgments: This work is supported by MRO Project funding to SHARAD DTL Putzig

References: [1] Campbell et al. (2021) Icarus, 360, 114358. [2] Lu et al. (2013) Interspeech, 2013, 436-440. [3] Xiong et al. (2016) Engineering Applications of Artificial Intelligence, 52, 194-202. [4] Wei and Huerta (2020) Physics Letters B, 800, 135081. [5] Isola et al. (2017) Proceedings of the IEEE conference on computer vision and pattern recognition, 1125-1134. [6] Phillips et al. (2008) Science, 320(5880), 1182-1185 [7] Allen (1977) IEEE Transactions on Acoustics, Speech, and Signal Processing, 25(3), 235-238 [8] Griffin and Lim (1984) IEEE Transactions on Acoustics, Speech, and Signal Processing, 32(2), 236-243.