Data-Driven Signal Separation in Radio Spectrum

Tejas Jayashankar, Binoy Kurien, Alejandro Lancho, Gary C.F. Lee,
Yury Polyanskiy, Amir Weiss and Gregory W. Wornell

MIT (e-mail: rfchallenge@mit.edu)
Github Repository: ICASSP 2024 RFChallenge
Discord server: ICASSP 2024 RFChallenge

Abstract

This challenge will require developing an engine for signal separation of radio-frequency (RF) waveforms. At inference time, a superposition of a signal of interest (SOI) and an interfering signal will be fed to the engine, which should recover the SOI by performing a sophisticated interference cancellation. SOI is a digital communication signal whose complete description is available (modulation, pulse-shape, timing, frequency, etc). However, the structure of the interference will need to be learned from data. We expect successful contributions to adapt existing machine learning (ML) methods and/or propose new ones from the areas of generative modeling, variational auto-encoders, U-Nets and others.

I. INTRODUCTION

With the proliferation of wireless technologies, communication systems of different types share the same part of the radio frequency (RF) bands. When active simultaneously, they create *co-channel* interference. In such cases, given a snapshot of the mixed (superimposed) different signals operating simultaneously, one is interested in separating them into their respective components. This may be the case when wireless devices operating in the same frequency band give rise to disruptions to their operation due to interference (be it intentional or not). For example, noise from a microwave oven occupies the same 2.4 GHz ISM band as several classes of wireless signals (e.g., 802.11 WiFi, Bluetooth and ZigBee), and therefore may interfere with such communication systems. Other examples include 5G vs radar vs satellite in the C- and Ka-bands.

The goal of *signal separation* (or *interference mitigation* in this specific context) is to extract the signal-of-interest (SOI) with highest fidelity, thereby improving downstream task's performance (demodulation and decoding).

Recent efforts in source separation have demonstrated how machine learning techniques could be used in domains such as computer vision and audio. However, the RF signal space posesses several specifics: on one hand intereference is superimposed approximately linearly (unlike in vision, where it results in hard occlusion), on the other hand on small time-frequency scales interference does not exhibit much structure (unlike audio). Note that in the context of RF signals, traditional separation architectures (multiplexing, filtering) are applicable only when signals are separable in time or frequency, in which case the components can be identified via spectrogram or a simple DFT.

The key signal processing challenge that we put forth is separation of co-channel signals (on a single, SISO, channel), in which the energy content of the independent components is overlapping in both time *and* frequency, partly or fully, and where standard separation methods perform relatively poorly. In particular, we are interested in a) the separation of a single co-channel signals; and b) demodulation of a SOI in the presence of a non-Gaussian, non-white co-channel interference.

The distinct characteristics of co-channel RF signals compared to the aforementioned computer vision and audio domains motivate the necessity for investigating and developing novel separation methods. Specifically, within the context of deep learning, there is a need for new neural network architectures. Notably, methods tailored for co-channel source separation of RF signals would likely need to capture additional, less trivial, and perhaps less intuitive features, which are not necessarily easily discernible through time and/or frequency domain analysis.

Creating such tools for signal separation and/or interference mitigation of RF signals holds potential applications across various fields. Providing a more accurate estimation of underlying components could aid downstream processes like anomaly detection or fine-grained classification. In the context of communications, this capability could function as an additional step for channel equalization, preceding standard demodulation and decoding steps.

II. DEMODULATION CHALLENGE

We consider scenarios where we know the generation process of one of the components, which serve as the SOI, and we are interested in demodulating such a SOI in the presence of co-channel interference.

Mixture Signals Generation

We consider 40960-samples long mixture signals of the form,

$$y = s + b \in \mathbb{C}^{40960}, \tag{1}$$

where s is the SOI, a digital communication signal whose generation process is known, and b is an interference signal, which is a time-series segment from one of the frames of the EMISignal1, CommSignal2, CommSignal3, or CommSignal5G1 dataset, to be described in detail in Section III. In this challenge, we focus on two different types of SOI, whose generation processes are described next.

SOI 1: QPSK signal: The first SOI, which we call "QPSK", is a single-carrier signal, modulated by a root-raised cosine pulse shaping function. This signal bears an M-bit long message, which are mapped to L symbols from a QPSK constellation using Gray coding. Each bit is randomly generated (via a fair coin toss) in an independent and identically distributed fashion. The n-th sample of s for this signal type is expressed as

$$s[n] = \sum_{\ell=0}^{L-1} a_\ell g[n - kF - \tau_0], \tag{2}$$

whereby $F \in \mathbb{N}$ is the symbol interval (in discrete-time) and $\tau_0 \in \{0, \dots, F-1\}$ is the offset for the first symbol (we use $F = 16$ and $\tau_0 = 8$ in this challenge); and $g[n]$ is the discrete-time impulse response of the transmitter filter (pulse shaping function; Figure 1a). The filter corresponds to a root-raised-cosine filter with roll-off factor 0.5 and window length of 127 samples (8 symbols). Figure 1b shows a simplified diagram for the generation process of the QPSK signal.

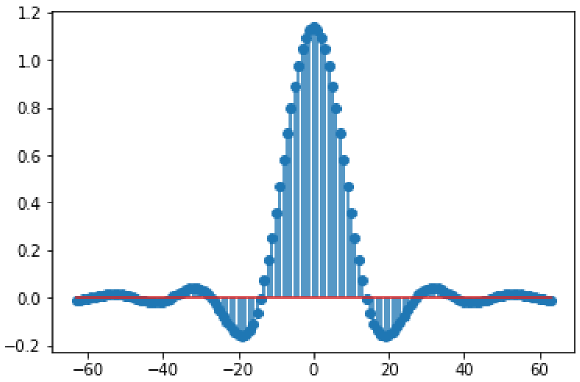

(a) Root-Raised-Cosine Pulse Shaping Function

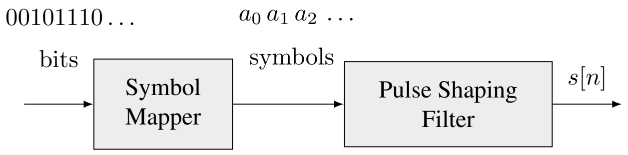

(b) Block diagram for the generation process of the QPSK signal (SOI 1) modulated using a pulse shaping filter.

SOI 2: OFDM QPSK signal: The second type of SOI is an orthogonal frequency division multiplexing (OFDM) signal—a multi-carrier signal that is comprised of K orthogonal subcarriers, each carrying a QPSK symbol. We call it "OFDM QPSK".

Again, the bits are randomly generated using a fair coin toss in an independent and identically distributed fashion, and they are subsequently mapped to symbols from a QPSK constellation using Gray coding. The n-th sample of this SOI is given by

$$s[n] = \sum_{p=0}^{P-1} \sum_{k=0}^{K-1} g_{k,p}\, r[n - p \cdot (K + T_{\text{cp}}) - T_{\text{cp}},\, k], \quad \forall n \in \{0, \dots, N-1\}, \tag{3}$$

where

$$r[n,\, k] \triangleq e^{\frac{j2\pi kn}{K}}\, \mathbb{1}\left\{-T_{\text{cp}} \le n < K\right\}, \tag{4}$$

and we recall that $N = 40960$ denotes the frame length. Here, K is the total number of orthogonal complex sinusoid terms,[1] also termed as subcarriers, that corresponds to the FFT size of the inverse discrete Fourier transform (DFT) invlolved in the generation process of an OFDM signal (see Figure 2). The coefficients $g_{k,p} \in \mathcal{G}$ are the information modulating symbols, where $\mathcal{G}$ stands for the alphabet (constellation), which is again QPSK. A cyclic prefix (CP) is typically added before an OFDM symbol. Hence, each OFDM symbol is described within the interval $[-T_{\text{cp}}, K]$, where T_{cp} is the CP length. In this challenge, $K = 64$, $T_{\text{cp}} = 16$, and $\bar{K} = 56$ out of the 64 available subcarriers are active (namely, the 8 inactive subcarriers "carry" the zero symbol). The signals then span $P = N/(\bar{K} + T_{\text{cp}}) \in \mathbb{N}$ OFDM symbols, and their individual finite support is reflected by the finitely-supported function $r[n,\, k]$ (4). Figure 2 shows the block diagram of the generation process of an OFDM symbol.

In this particular challenge, in order to focus on separation aspects, we make an additional simplifying assumption that synchronization and channel estimation has been successfully accomplished; hence, phase offsets have been compensated, and the underlying SOI is aligned such that the sampling points for the symbol of the corresponding ground truth component are the same across the different samples. This way, the focus is on the mitigation of the interference component from the SOI.

---

[1] Which are nonetheless not all necessarily active.

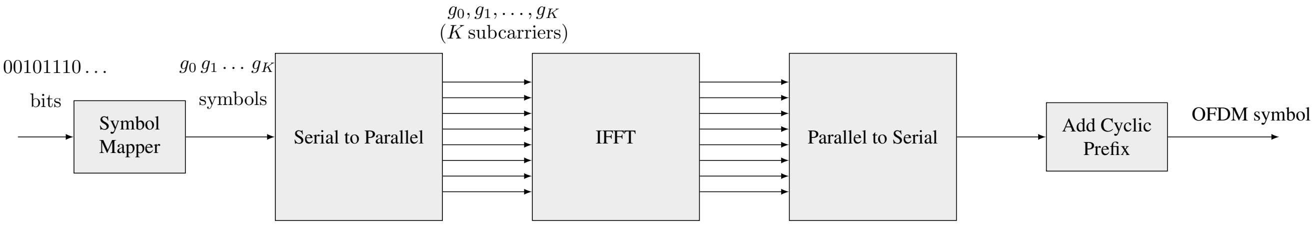

Fig. 2: Block diagram for the generation process of an OFDM symbol carrying QPSK symbols in each active subcarrier.

Interference: A frame of the respective signal type is first selected at random (uniformly) from the dataset (described in Section III), and a random window of $N = 40960$ samples is chosen from it. Each component is then scaled to achieve a target (empirical) signal-to-interference-and-noise ratio (SINR). Specifically, the selected window of interference signal (EMISignal1, CommSignal2, CommSignal3, or CommSignal5G1) is scaled based on the target SINR level. Thus, for a target SINR level $\kappa^2 = 10^{(\text{SINR in dB})/10}$, the interference signal is scaled by $1/\kappa$. Each signal dataset is normalized to have unit power (i.e., empirical variance). Note that this normalization implies that the *instantaneous* SINR in a given interference frame may not necessarily have unit power. Still, we refer to κ^2 as the target SINR.

Additionally, each interference frame b incorporates a random phase rotation before being added to the SOI s to create a mixture example y. Consequently, every interference frame b encompasses the interference frame from the dataset (not necessarily with a *per-frame* unit power), with a scaling of $1/\kappa$ and a random phase rotation. The scaling factor κ and the instantaneous SINR of every interference frame added to the mixture are provided as metadata and can be exploited at training.

Demodulation Metric

Participants are tasked with designing algorithms to estimate both the SOI waveform and its bit-sequence message from the given signal y. The anticipated output should include an estimate of the SOI waveform $\hat{s} \in \mathbb{C}^{40960}$ and the associated bit message $\hat{m}$. For both QPSK and OFDM QPSK SOIs, the estimate $\hat{s}$ should consist of 40960 complex-valued samples. However, for SOI 1 (QPSK) the estimated message $\hat{m}$ should encompass 5120 bits (two bits per symbol, and recall that $F = 16$), and for SOI 2 (OFDM QPSK) it should encompass 57344 bits ($(40960/80) = 512$ OFDM symbols $\times 56$ active subcarrieris $\times 2$ bits per QPSK-symbol).

The performance metrics based on which the participants' score will be calculated are the **mean-square error (MSE)** between the estimated signal $\hat{s}$ and the true SOI waveform s^*, and the **bit error rate (BER)** between the estimated bits $\hat{m}$ and the true message m^*.

Test mixtures are provided for various target SINR levels, ranging from -30dB to 0dB at 3dB steps (a total of 11 SINR levels), with 100 test cases per target SINR level. The SINR of each test case will be provided for possible exploitation.

In this iteration of the RF challenge, the signal types of the components (i.e., the SOI and interference types) will be disclosed, allowing participants to evaluate the signal mixtures separately. This information can be effectively leveraged in participants' proposed solutions.

At intermediate validation runs, participants are encouraged to evaluate their $2 \times 4 \times 2$ performance curves (four mixture configurations vs. (MSE, BER)) as a function of the strength of the interference.

Final scoring: There will be two rankings, one based on BER and one based on MSE.

1) BER score: The (scalar) score for each possible mixture (i.e., each combination of SOI and interference) will correspond to the sum of smallest target SINR levels from the 11 provided values, ranging from -30 dB to 0 dB, at which the BER is smaller than or equal to 10^{-2}. In case of not reaching a BER of 10^{-2} for any value of the target SINR set, a zero will be assigned for that specific mixture case.
   Specifically, we denote $\text{BER}_{\text{SOI}_i\text{\_Interference}_j}(\text{SINR}_k)$ as the BER of the i-th SOI and the j-th interference at the k-th SINR level, where $(i, j, k) \in \{1, 2\} \times \{1, 2, 3, 4\} \times \{1, \ldots, 11\}$, $\text{SINR}_k = -30 + 3(k-1)$ dB, and

$$i = 1 \iff \text{QPSK},\ i = 2 \iff \text{OFDM QPSK},$$
$$j = 1 \iff \text{EMISignal1},\ j = 2 \iff \text{CommSignal2},\ j = 3 \iff \text{CommSignal3},\ j = 4 \iff \text{CommSignal5G1}.$$

   We also define

$$\overline{\text{SINR}}_{ij} \triangleq \left\{\min(\text{SINR}_\kappa), \kappa \in \{1, \ldots, 11\} : \text{BER}_{\text{SOI}_\text{i}\text{\_Interference}_\text{j}}(\text{SINR}_\kappa) \leq 10^{-2}\right\}. \tag{5}$$

   With these notations, the final BER score is calculated as

$$\text{Final BER score} = \sum_{i=1}^{2} \sum_{j=1}^{4} \overline{\text{SINR}}_{ij}. \tag{6}$$

2) MSE score: For each mixture, the (scalar) score will correspond to the truncated average of MSE values (in dB scale) over the 11 provided SINR levels, which range from $-30\,\text{dB}$ to $0\,\text{dB}$, where the truncation level is at $-50\,\text{dB}$. Specifically, we denote $\text{MSE}_{\text{SOI}_i\text{\_Interference}_j}(\text{SINR}_k)$ as the MSE of the i-th SOI and the j-th interference at the k-th SINR level, where the association of the triplet (i, j, k) as is described above for the BER score. For each mixture, we define

$$\overline{\text{MSE}}_{ij} \triangleq \frac{1}{11}\sum_{k=1}^{11} \max\left\{-50, \text{MSE}_{\text{SOI}_i\text{\_Interference}_j}(\text{SINR}_k)\right\}. \tag{7}$$

The final MSE score is then calculated as

$$\text{Final MSE score} = \sum_{i=1}^{2}\sum_{j=1}^{4} \overline{\text{MSE}}_{ij}. \tag{8}$$

Ties will only be resolved during the final evaluation of the challenge (refer to Section IV-B for the dates of the partial and final evaluations). If a tie occurs in the final evaluation, it will be resolved based on the stability and generazability of the trained models, e.g., by using additional mixtures generated with a finer SINR grid or by using more examples per SINR level.

Summary of Demodulation Challenge

- Goal: Develop a machine learning algorithm aimed at rejecting interference
- Mixture: SOI (QPSK or OFDM QPSK) + Interference (EMISignal1, CommSignal2, CommSignal3 or CommSignal5G1)
- Evaluation Metric: Final BER score (6) and Final MSE score (8), as a function of target SINR. (The lower, the better).

III. DATASET

The relevant datasets described below in Section III-A were all created from a "global" dataset that contains all the examples of all the interference signal type (recall that the SOIs' generation processes are known, hence saved dataset are not required for them, as they can be generated locally by participants). The global dataset contains examples of **four types** of interference:

1) EMISignal1: an electromagnetic interference due to unintentional radiation from a man-made source;
2) CommSignal2: a digital communication signal from commercially available wireless device;
3) CommSignal3: yet another, but different, digital communication signal from commercially available wireless device; and
4) CommSignal5G1: a 5G-compliant waveform;

The examples in the dataset of the first three types (EMISignal1, CommSignal2 and CommSignal3) have been recorded over-the-air, and the last one (CommSignal5G1) was generated and recorded within a controlled wired laboratory environment, where impairments inherent to wireless communication mediums were introduced through designated simulators.

A. Detailed Structure of the Datasets for the Challenge

For any of the four interference types, the global dataset was divided into three main (sub-)datasets—INTERFERENCESET, TESTSET1 and TESTSET2. The composition of these three main (sub-)datasets are described in detail next.

1) INTERFERENCESET: Each sub-dataset, corresponding to one of the four types of interference, is a single file. Each such file contains a number of *frames*. The length (in complex-valued samples) of each frame may differ for each interference type: 230000 for EMISignal1, 43560 for CommSignal2, 260000 for CommSignal3, and 230000 for CommSignal5G1. As mentioned above, all the frames were extracted from real-world recordings (either over-the-air or a wired laboratory setup with channel emulators). Each frame is saved in the same format (details below). Representative visualizations of the four signal types are shown in Figure 3. Information on how to generate training frames using the INTERFERENCESET can be found in Github repository README ("Helper Functions for Training").
2) TESTSET1: Contains 50 frames with the same specifications as the INTERFERENCESET described above, but that are NOT part of the INTERFERENCESET.

2.1) TESTSET1EXAMPLE: By executing the file `sampletest_testmixture_generator.py` (with the relevant input arguments; see the starter code in the challenge Github repository), a dataset consisting of the specified SOI and interferencetypes are created. Each such generated sub-dataset consists of 1100 waveforms + metadata, wherein each waveform consists of 40960 complex-valued samples that corresponds to a sum of an SOI and interference. The SOI is perfectly frame- and symbol-synchronized. The interference is computed by first selecting a (uniformly) random frame from TESTSET1, then cropping it to random 40960-samples long segments, and then multiplying by the inverse square root target SINR and the random phase rotation. Metadata includes the exact values for these transformations. Additionally, the ground truth values of the SOI and interference signals are also provides for this dataset.

2.2) TESTSET1MIXTURE: Same generation process as of TESTSET1EXAMPLE, but the ground truth values of the SOI and interference signals *are not provided*. **Note:** This is the test set that will be used for evaluation of the intermediate submissions (See Section IV-B for details on submission deadlines).

3) TestSet2Mixture: **will be released only 5 days prior to the final deadline**. It will have *exactly* the same format as TestSet1Mixture and will be statistically equivalent with respect to the generation process of TestSet1Mixture. However, the interference segments will be drawn from 50 frames which *were not* part of either InterferenceSet or TestSet1, and will not be released. The scores attained on this testset will be used for the final ranking (See Section IV-B).

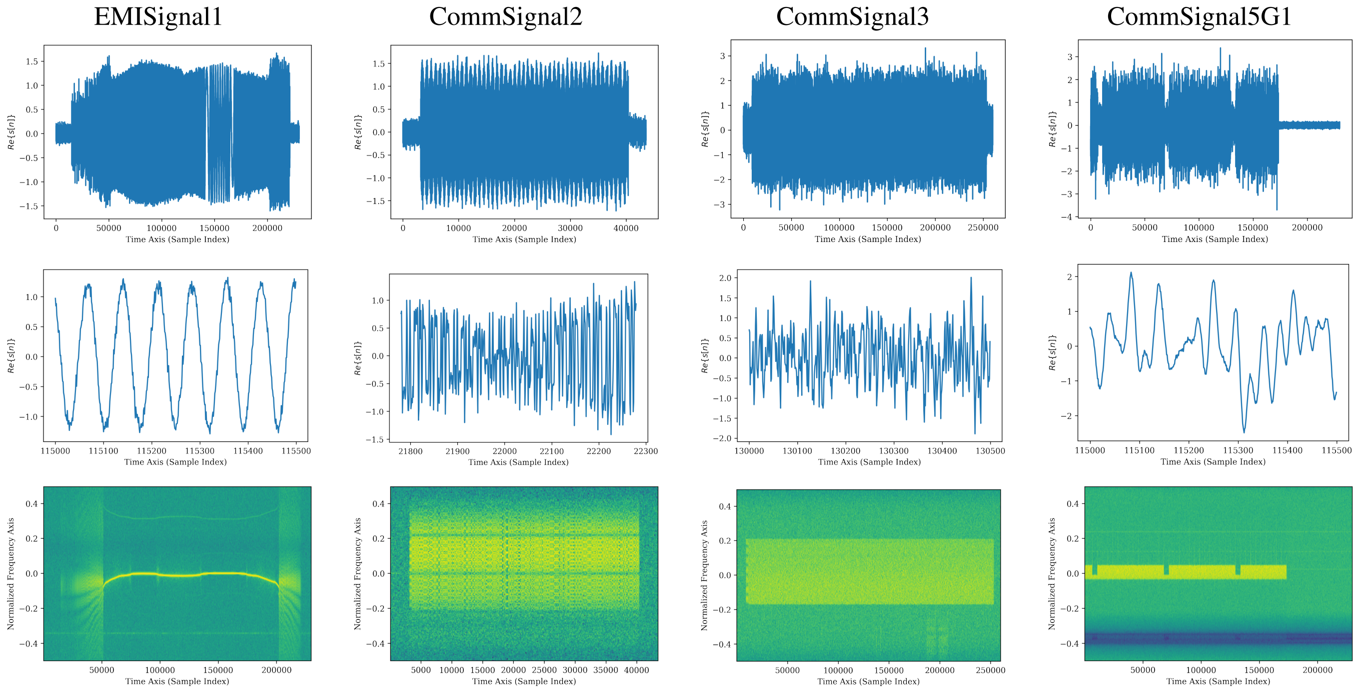

Fig. 3: Representative frames of the four signal types in the dataset: EMISignal1, CommSignal2, CommSignal3 and CommSignal5G1. Top: Plotting the real part of the waveforms; Middle: A zoomed-in segment of the signal; Bottom: Spectrogram of the respective frames.

Some additional details and considerations about the datasets are the following:

- The length of each frame is consistent within a given signal type, but may differ between the different types provided ($230,000$ complex-valued samples for EMISignal1, $43,560$ for CommSignal2, $260,000$ for CommSignal3, and $230,000$ for CommSignal5G1). At any rate, the length of all the input mixtures is 40960-samples long.
- The signals in EMISignal1 and CommSignal5G1 have been shifted in frequency such that the majority of their spectral energy content lies in baseband frequencies.
- The frames are saved in an HDF5 file.
- These frames are used in generating the mixture signals for the challenge. For each signal type, 50 of such frames have been set aside to create the testing set (TestSet2Mixture) for final evaluation.
- Details regarding technical coding aspects and examples on how to handle these datasets to create mixtures for training and inference (and more utility functions) are provided in the README file and in the starter code files within the Challenge's Github repository.

IV. Evaluation and Deadlines

A. Dataset for evaluation

Each of the test files contain 100 mixtures across 11 target SINR levels ranging from -30 dB to 0 dB, corresponding to 1100 test cases altogether. The specifications of the mixtures can be found in Sections II and III.

B. Submission and Scoring

To evaluate the performance of your submission based on the signal mixtures from the corresponding TestSet, please upload the following numpy arrays (.npy files, and only in this format) to your preferred cloud server (e.g., Dropbox or Google Drive):

- The 1100 segments (11 SINR levels $\times 100$ examples) of the 40960-samples long signal estimates for the SOIs; and
- Their corresponding bit string estimates (see the starter code readme file for the demodulation process).

Submissions not abiding by the format requirements will not be evaluated, and consequently, they will not be able to receive their final scores.

You can use separate files for each type of mixture to simplify the upload process. Once you have uploaded these files to the cloud server, kindly share access with the Challenge organizers, and notify them via email or via Discord, to facilitate the evaluation process.

The organizers of the RF Challenge will run the test script to compute the MSE and BER outputs and share these summary metrics with participants.

The final score used to determine your overall ranking will be based on the two separated rankings according to (6) and (8).

Important Dates:

- **Oct. 4, 2023** - Submission 1 deadline: Initial submission containing outcomes on TESTSET1MIXTURE.
- **Nov. 1, 2023** - Submission 2 deadline: Second submission containing the outcomes on TESTSET1MIXTURE.
- **Dec. 1, 2023** - Final submission deadline: Last submission containing the outcomes on TESTSET2MIXTURE. The final ranking will be exclusively determined by the results of this ultimate submission.

We note that the intellectual property (IP) is not transferred to the challenge organizers, i.e., if code is shared/submitted, the participants remain the owners of their code (when the code is made publicly available, an appropriate license should be added).

V. BASELINE-METHODS EVALUATION

We trained baseline methods using the INTERFERENCESET[2] In Figure 4, we present performance curves on mixtures generated from TESTSET1EXAMPLE, hoping they will serve as a helpful reference for participants. The various plots depict BER and MSE against the target SINR, for the different combinations of SOI and interference used in this Challenge. You can access all these figures and their generation code in the Challenge Github repository. In particular, use the following notebook link.

ACKNOWLEDGMENT

The preparation of this challenge and its data was supported, in part, by the United States Air Force Research Laboratory and the Department of the Air Force Artificial Intelligence Accelerator under Cooperative Agreement Number FA8750-19-2-1000, and with resources from MIT SuperCloud and the Lincoln Laboratory Supercomputing Center.

---

[2] For training we found that using a continuous range of target SINR values from $-33\,\text{dB}$ to $3\,\text{dB}$ yields better results (in generalizability) than focusing on only the 11 test SINR values (see `example_generate_rfc_mixtures.py` within the Challenge Github repository).

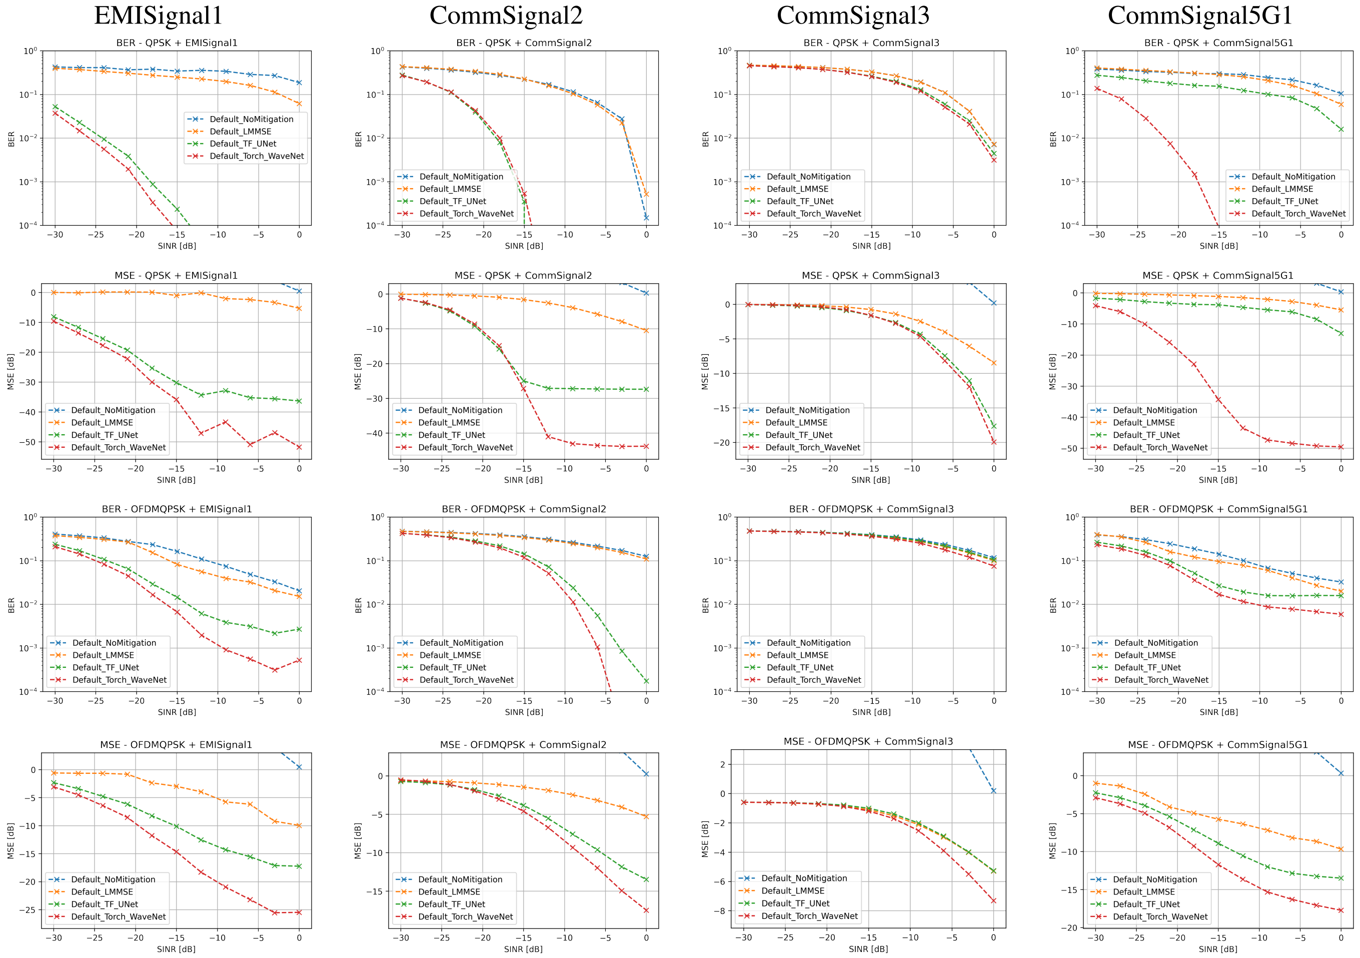

Fig. 4: BER and MSE as a function of the target SINR for all combinations of SOI and interferences considered in this challenge. Specifically, for the two SOI types: QPSK and OFDM QPSK; and the four signal types in the dataset: EMISignal1, CommSignal2, CommSignal3 and CommSignal5G1.