Single-Channel RF Challenge

MIT (e-mail: rfchallenge@mit.edu)

I. Introduction

With the proliferation of wireless technologies, many different types of devices might occupy the same part of the radio frequency spectrum. Here, given a snapshot of different signals operating simultaneously, one might be interested in separating them into their respective components for further analysis and characterization. Further, wireless devices operating in the same frequency band may 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 therefore may cause interference to such communication systems. Thus, either signal separation or interference mitigation would help in extracting the signal-of-interest with potentially higher fidelity, thereby improving the performance in its 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. In those domains, however, the algorithms could exploit the fact that components may be separable in space (e.g. images), time or frequency (e.g. audio). In the context of RF signals, if components are separable in time and/or frequency, one could potentially identify the different components through the spectrogram of the measured mixture signal, and separate the components through appropriate masking and conventional filtering methods.

A. Challenge Statement

The key challenge we identified is the processing of co-channel signals, for which components are overlapping (possibly partially/fully) in time and frequency – particularly, we are interested in a) the separation of co-channel signals; and b) demodulation of a signal-of-interest in the presence of other interference signals.

The differences in characteristics of co-channel RF signals from the other aforementioned domains, such as images and audio signals, motivate the need to investigate and develop new methods. In particular, the methods for co-channel RF signals would likely have to capture other features that are not necessarily easily identified in a standard time and/or frequency domain analysis—of the underlying components to aid in the signal separation process.

Can data-driven methods capture signal structures useful for the separation and/or demodulation?

Designing such a signal separation and/or interference mitigation tool could find uses in various applications. Having a higher fidelity estimate of the underlying component may assist in downstream processing, such as anomaly detection or finer-grained classification. From the perspective of communications, such a capability could potentially serve as an add-on to channel equalization steps before standard demodulation and decoding steps, by exploiting knowledge about the likely interference present (as trained from examples).

B. Recommended Design Considerations

In the spirit of scalable development, participants are encouraged to consider the following in their architecture design:

- The pipeline should depend as little as possible on particular specifications of the signals; the particular specifications of different signal types should be encapsulated by the corresponding model used in the signal separation/interference mitigation step.
- A suggested guiding principle is the following —- we are analyzing a new mixture with two different classes of signals (which we might be able to identify using other techniques); a different model may be used that corresponds to the classes of input components expected, but the algorithmic pipeline should remain largely unchanged.
- Additionally, we are interested in solutions that accounts for the following situation – consider that a new signal type has been introduced, and we would like to analyze a mixture between this new signal and one of the older types; can we still separate or demodulate this new signal with minimal to no re-training on the previously seen and trained signal types?

II. Dataset

The dataset comprises recordings from three types of signals – an electromagnetic interference due to unintentional radiation from a man-made source (EMISignal1), and two types of digital communication signals from two different sources (CommSignal2 and CommSignal3).

A representative time-series plot of the three signal types are shown in Figure 1.

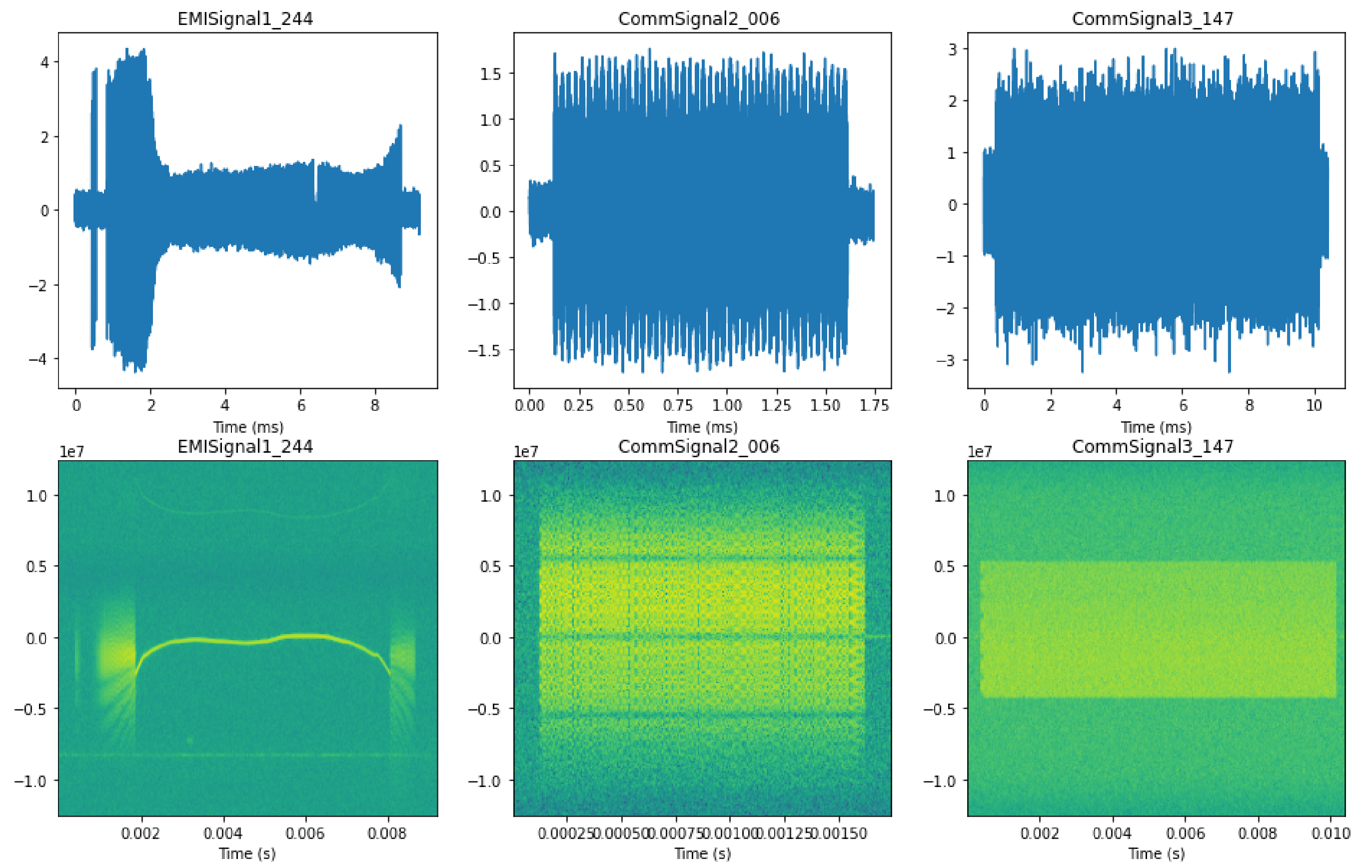

Fig. 1: Representative Frames of the three signal types in the dataset – Left: EMISignal1, Middle: CommSignal2 and Right: CommSignal3; Top: Plotting the real part of the waveforms; Bottom: Spectrogram of the respective frames

Within the dataset, each file contains the waveform of a characteristic frame. The length of each frame is consistent within a given signal type, but differs between the different types provided (9.2 ms for EMISignal1, 1.7424 ms for CommSignal2, and 10.4 ms for CommSignal3; all measurements are sampled at 25 MHz).

The respective frames were extracted from over-the-air recordings, and scaled to unity power within each frame. Additionally, signals in EMISignal1 has been shifted in frequency such that the majority of its energy lies in baseband frequencies. Specifically, this is achieved by finding the peak frequency in each frame's power spectral density (using Welch's method), and shifting it to the origin.

Each frame is saved in SigMF format [1]. The corresponding metadata of each file contains the respective scaling and shift factors applied on the extracted frames from the raw recordings

These frames are used in generating the mixture signals for the separation and demodulation challenge; details on these mixtures are provided in the next section. For each signal type, 50 of such frames have been set aside to create the testing set for final evaluation. The remainder have been collated to form the training set (`train_frame`), which consists of examples of the individual signal types. We have also provided a 'validation set', which are generated using the last 50 frames of each signal type from the training set; the generation of this validation set shares the same pipeline as that for the testing set.

The testing set will not be released at the launch of the challenge; instead, more details will be provided closer to the evaluation date.

III. Challenges

Summary of Proposed Challenges

- Separation Challenge
  - Goal: Develop a machine learning algorithm that can recover high-fidelity reconstructions of individual co-channel signals given mixtures of them
  - Evaluation Metric: mean square error (MSE) from 'ground truth' of the raw waveforms
  - Mixture: CommSignal2 + Interference (either EMISignal1 or CommSignal3)
  - Supplementary Metric: Packet Error Rate (PER) obtained by passing the CommSignal2 component estimate into a standard demodulator and a Cyclic Redundancy Check (CRC) integrity checker
- Demodulation Challenge
  - Goal: Develop a machine learning algorithm that can reject interference in order to protect a communications receiver
  - Evaluation Metric: Bit Error Rate (BER) on the QPSK signal as a function of its Signal-to-Interference-Plus-Noise Ratio (SINR)
  - Mixture: QPSK Signal + Interference (either EMISignal1, CommSignal2 or CommSignal3)
  - Supplementary Metric: Binary Cross-Entropy (Log-Loss) could be considered, especially to assess 'harder' conditions

IV. Task 1: Separation Challenge

In this challenge, we consider situations where the received waveform is a superposition of multiple components. In such a setting, we may be interested in separating the underlying sources, thereby obtaining an estimate of the waveform for the respective components. A high fidelity reconstruction of the respective components could potentially help in subsequent processing or analysis of the individual sources (e.g. decoding through a standard demodulator tool).

Generating Test Mixture Signals for Separation

We consider a mixture signal, $y \in \mathbb{C}^{40960}$, which is a sum of 2 components –

$$y = s + b, \tag{1}$$

where s is a digital communication signal (CommSignal2), and b is an interference signal (either EMISignal1 or CommSignal3).

For each of these components, a frame of the respective signal type is first selected at random from the dataset (outlined in Section II), and a random block of 40,960 samples is chosen from it. Each component is then scaled to achieve a target signal-to-interference-and-noise-ratio (SINR). (Figure 2)

Specifically, the CommSignal2 block is scaled such that $\|s\|_2^2 = 1$, and the interference block (EMISignal1 or CommSignal3) is scaled to attain a target SINR level, i.e., $\|s\|_2^2 / \|b\|_2^2 =$ Target SINR.

Note 1: For ease of computation, we scaled s, which inherently has some noise, to have a signal power of 1. Additionally, the 'SINR' here refers to the signal power ratio between s and b, despite s also containing background noise.

Empirically, the recordings of s has a signal-to-noise ratio of around 17 dB. To account for the noise present in s, the actual SINR can be approximated as

$$\textit{Actual SINR} \approx \frac{P_s - P_n}{P_n + P_b} \approx 0.98 \cdot \textit{SINR}$$

Note 2: In the testing set, at least one of the components is selected from the unseen portion of the dataset that has been set aside. The other component could potentially come from a frame that is in the training set; however, the mixture would still be an unseen realization.

Separation Metric

Given y, participants are to design algorithms that can extract estimates of the two respective components $\hat{s}$ and $\hat{b}$.

Based on these estimates, the performance metric is the mean square error (MSE) of the components. In particular, this challenge will focus on the reconstruction of the s component (CommSignal2 waveform) –

$$\mathrm{MSE}(\hat{s}, s) = \|\hat{s} - s^*\|_2^2. \tag{2}$$

We also encourage participants to include the MSE of the other component, particularly if it differs from the MSE reported above –

$$\mathrm{MSE}(\hat{b}, b) = \|\hat{b} - b^*\|_2^2. \tag{3}$$

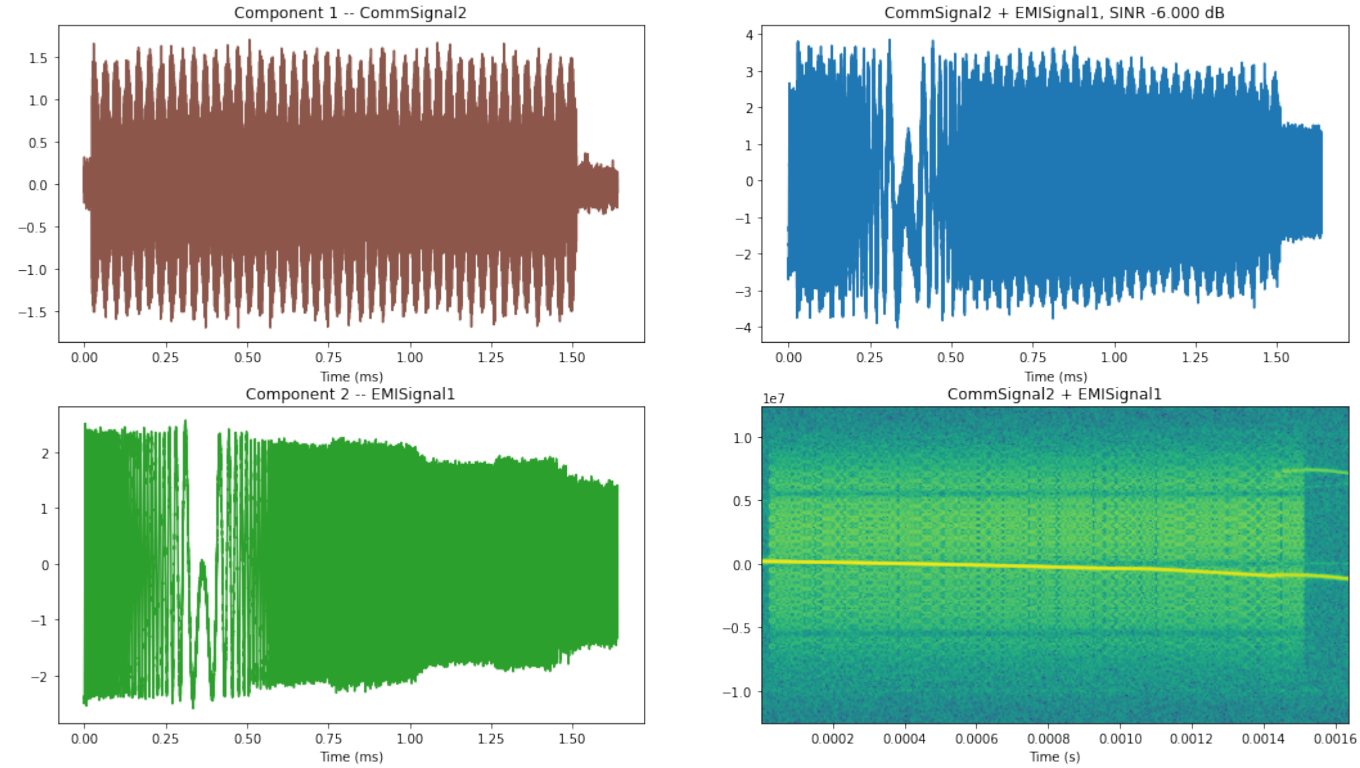

Fig. 2: Example of a mixture signal of CommSignal2 and EMISignal1

Test mixtures for each type of 11 different SINR, ranging from -12 dB to +3 dB, and 100 test cases per SINR level. Note that the SINR of each test case will not be provided.

In the current iteration of the RF challenge, the signal types of the components (i.e., the type of interference) present will be provided, and this information can be utilized in the participants' solution.

Supplementary Metric based on Packet Integrity

Given the estimated CommSignal2 waveform, $\hat{s}$, we can also evaluate the integrity of the reconstructed component. This is done through the use of a standard demodulator for that signal type, after which we utilize the cyclic redundancy check (CRC) to assess the integrity of the frame and thereby obtain a packet error rate (PER) metric across the testing set.

We will be using this as a bonus metric to assess the fidelity of the extracted component, in addition to the 'closeness' in reconstruction as captured by the MSE metric.

(The demodulation tool for CommSignal2 is currently unavailable; however, its inclusion is in our planned future release. We will announce more details pertaining to this supplementary metric when such a tool is made available to the public.)

V. Task 2: Demodulation Challenge

In this challenge, we consider situations where we possess information about one of the components, which serve as the signal-of-interest. In such a setting, we may be interested in demodulating the signal-of-interest in the presence of co-channel interference.

Generating Test Mixture Signals for Demodulation

We consider a mixture signal, $y \in \mathbb{C}^{40960}$, which is a sum of 2 components –

$$y = s + b, \tag{4}$$

where s is a QPSK modulated signal, and b is an interference signal (either EMISignal1, CommSignal2 or CommSignal3).

The QPSK signal (Figure 3) is a single carrier modulated by a 5120-long bit message, corresponding to 2560 symbols (i.e., one of four possible symbols, $a \in \{+1, -1, +j, -j\}$) encoding 2 bits each. Each bit is randomly generated in an independent and identically distributed fashion. We can express the n-th element of the QPSK signal s as

$$s[n] = \sum_{k=0}^{K-1} a_k g[n - kN - \tau_0], \tag{5}$$

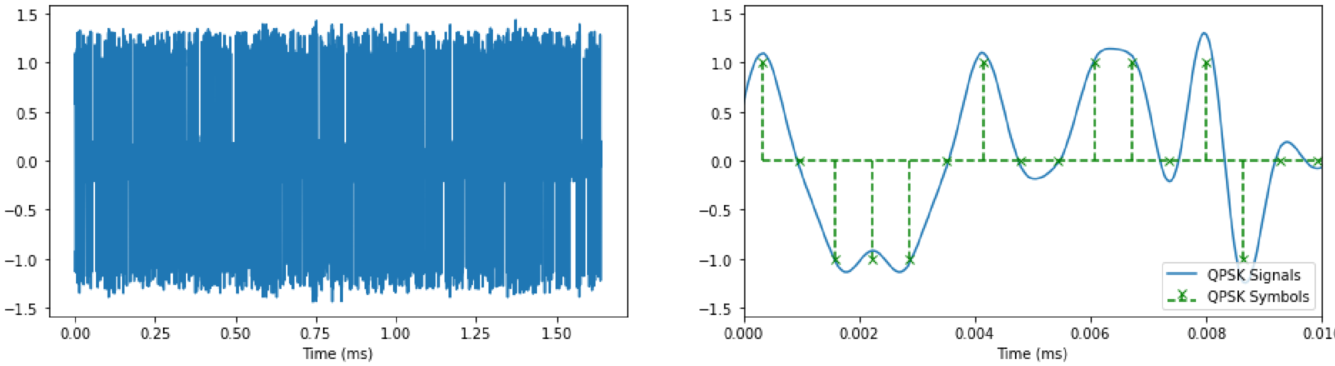

Fig. 3: QPSK Signal (real part of the waveform) – Left: Block of 40960 samples used in the challenge; Right: Zoom-in on the first 10μs of the signal

whereby N is the symbol interval and τ_0 is the offset for the first symbol (we use $N = 16$ and $\tau_0 = 8$ for this challenge); and $g[n]$ is the discrete-time impulse response of the transmitter filter (pulse shaping function; Figure 4). The filter corresponds to a root-raised-cosine filter with roll-off factor 0.5 and window length of 127 (8 symbols).

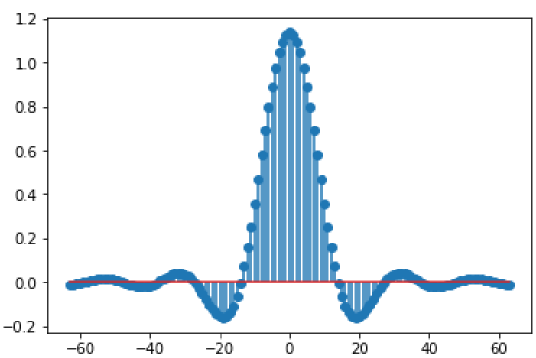

Fig. 4: Root-Raised-Cosine Pulse Shaping Function

In this particular challenge, we make a simplifying assumption that synchronization and channel estimation has been done; hence, phase offset has been compensated, and the underlying QPSK signal is aligned such that the sampling points for the symbol for the corresponding ground truth component is the same across the different samples. The focus is hence placed on the source separation and/or mitigation of the interference component from the QPSK signal-of-interest.

For the interference, a frame of the respective signal type is first selected at random from the dataset (outlined in Section II), and a random window of 40,960 samples is chosen from it. Each component is then scaled to achieve a target signal-to-interference-and-noise-ratio (SINR).

Specifically, the interference block (EMISignal1, CommSignal2 or CommSignal3) is scaled to attain a target SINR level, i.e., $\|s\|_2^2/\|b\|_2^2 =$ Target SINR.

Note: In creating the testing set, only the 50 unseen frames from the dataset are used. This is to ensure that the mixtures in the test set have not been encountered through the training set.

Demodulation Metric

Given y, participants are to design algorithms that can estimate the bit message of the QPSK signal in the presence of interference; the expected output should be an estimate of the bits, $\hat{m}$. Hence, 5120 bits (1's and 0's) are expected in $\hat{m}$.

The performance metric is Bit Error Rate (BER) between the estimate bits $\hat{m}$ and true message m^*.

Test mixtures for each type of 11 different SINR, ranging from -12 dB to +3 dB, and 100 test cases per SINR level. Note that the SINR of each test case will not be provided.

In the current iteration of the RF challenge, the signal types of the components (i.e., the type of interference) present will be provided, and this information can be utilized in the participants' solution.

The output expected of the submitted algorithm is the estimated bits from the underlying QPSK signal; the estimate of the QPSK signal waveform or the interference waveform is not necessary in this section. Nonetheless, participants may adopt similar solutions from the separation challenge established earlier, and pass the estimated QPSK signal component through a standard matched filtering routine (provided in the starter code).

Supplementary Metric involving Log-Likelihood (Soft Demodulation)

As a bonus metric, we can also evaluate demodulation performance with estimated bit probabilities or log-likelihood ratios as the output. With bit probabilities (where p_j corresponds to the probability that j-th bit is 1), we can compute scores based on log loss (or binary cross entropy) –

$$\text{Log Loss} = -\frac{1}{M}\left(\sum_{j=0}^{M-1} \mathbb{1}_{m_j=1}\log_2(\hat{p}_j) + \mathbb{1}_{m_j=0}\log_2(1-\hat{p}_j)\right) \tag{6}$$

($\mathbb{1}_{m_j=1}$ is 1 if the j-th bit of the message m is 1, and 0 otherwise).

The pitfall with this metric is that the log-loss score is unbounded, and will be biased by an overly confident but wrong estimate of the probabilities on a few bits. Nevertheless, we recognize that a soft demodulation would be useful, particularly in future runs where coded bits are adopted. Furthermore, bit probabilities may be used to convey confidence on the estimates, which is useful for assessing harder situations where signal separation and/or interference mitigation may be difficult.

If you are developing a soft demodulation method, we encourage you to include such a metric in your results and discussion.

VI. Performance Evaluation of Submissions

A. Dataset for evaluation

During this preliminary stage of the RFChallenge, performance evaluation will be conducted over the test cases in the `sep_val` and `demod_val` folders contained within the dataset. When the official competition begins, we will release test cases in `sep_test` and `demod_test` folders.

Each of these folders contain 100 mixtures across 11 SINR levels ranging from -12 dB to +3 dB, corresponding to 1100 test cases altogther. The specifications of the mixtures have been described in the respective sections (IV. Separation Challenge for `sep_val`/`sep_test` and V. Demodulation Challenge `demod_val`/`demod_test`).

B. Performance Metric and Scoring

To evaluate the performance of your submission, please submit the 1100 values corresponding to the performance metric of the respective challenge (MSE of CommSignal2 component for separation and BER of the QPSK SOI for demodulation across the 1100 test cases from the corresponding folder). We recommend saving these values and submitting them to us in a CSV file.

The final score used to determine your ranking will be based on the average across the performance metric values. (The scoring mechanism is subject to changes as we finalize the details of the challenge and release of the test dataset.)

Pre-Competition Evaluation

In the ”validation dataset” provided in the initial round of this competition, the test cases are sorted with respect to their SINR (i.e., the first 100 corresponds to SINR +3 dB, the next 100 has SINR +1.5 dB, and so on). Hence, you would be able to plot the performance metric of the respective test cases against SINR. You may refer to the plots on the Jupyter notebook for reference (https://github.com/RFChallenge/rfchallenge_singlechannel_starter/blob/main/notebook/QuickStart.ipynb).

You are encourage to also provide these plots, where the results of your solutions are overlaid against the curves for the baseline methods provided for comparisons.

(For the actual test dataset for the competition, the actual SINR values of the individual test cases are not provided. We can provide this information upon receiving your submission if you wish to generate a similar plot. However, it should be noted that explicit knowledge about the SINR should not be used in your algorithm.)

C. Submitting solutions

If you wish to be included on the Single-Channel RFChallenge Leaderboard, please submit a write-up pertaining to your solution with details regarding the separation/demodulation algorithm and the metric values mentioned above. The organizers will use this write-up to verify that the solution is within the guidelines set forth by the problem specification, before including your submission to the Single-Channel Leaderboard.

At this stage of the competition, please direct all submissions (with evaluations and comparisons based on test cases in `sep_val` and/or `demod_val`) to rfchallenge@mit.edu.s

Acknowledgment

The authors acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing HPC resources that have contributed to the research results reported within this report.

References

[1] “The Signal Metadata Format (SigMF), v0.0.2,” https://sigmf.org, 2020.