

# Turbo Autoencoder with a Trainable Interleaver

Karl Chahine\*, Yihan Jiang<sup>†</sup>, Pooja Nuti\*, Hyeji Kim\*, and Joonyoung Cho<sup>‡</sup>

\* University of Texas at Austin, <sup>†</sup> Aira Technology, <sup>‡</sup> Samsung Research America

**Abstract**—A critical aspect of reliable communication involves the design of codes that allow transmissions to be robustly and computationally efficiently decoded under noisy conditions. Advances in the design of reliable codes have been driven by coding theory and have been sporadic. Recently, it is shown that channel codes that are comparable to modern codes can be learned solely via deep learning. In particular, Turbo Autoencoder (TURBOAE), introduced by Jiang et al., is shown to achieve the reliability of Turbo codes for Additive White Gaussian Noise channels.

In this paper, we focus on applying the idea of TURBOAE to various practical channels, such as fading channels and chirp noise channels. We introduce TURBOAE-TI, a novel neural architecture that combines TURBOAE with a trainable interleaver design. We develop a carefully-designed training procedure and a novel interleaver penalty function that are crucial in learning the interleaver and TURBOAE jointly. We demonstrate that TURBOAE-TI outperforms TURBOAE and LTE Turbo codes for several channels of interest. We also provide interpretation analysis to better understand TURBOAE-TI.

**Index Terms**—neural channel coding, Turbo autoencoder, deep learning, CNN, fading channels, chirp signal, bursty noise

## I. INTRODUCTION

The discipline of coding theory has made significant progress in the past seven decades since Shannon’s celebrated work [1]. Several codes have been invented for reliable communications, featuring convolutional codes, turbo codes, Low-Density Parity Check (LDPC) codes, and polar codes.

Deep learning is an emerging and powerful tool that has demonstrated promising success in channel coding [2]–[4] as well as other aspects of physical layer communications, such as beamforming [5], modulation [6], channel estimation [7] and channel feedback [8]. In channel coding, on the one hand, it is shown that introducing learnable components to the existing codes and decoders leads to the improvement in reliability and complexity (e.g., weighted belief propagation decoder for linear codes [4], [9] and [10], [11]). On the other hand, it is shown that end-to-end neural network based codes can perform comparably to the state-of-the-art codes for Additive White Gaussian Noise (AWGN) channels [12], [13]. As depicted in Figure 1, a pair of channel encoder and decoder modeled as neural networks and jointly trained via backpropagation achieves the reliability of modern codes.

While the training framework in Figure 1 might look straightforward, training a reliable code is shown to be highly challenging [12]. The generic neural architectures, such as Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN), are empirically shown to perform poorly [12] as they tend to learn codes with a very short

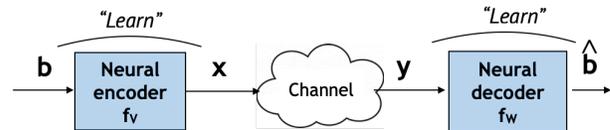


Fig. 1. Autoencoder framework for channel coding: channel encoder and decoder are modeled as neural networks and are trained jointly.

memory. To address these challenges, [12] introduces a custom neural architecture called TURBOAE, which concatenates parity bits generated from three CNNs; as depicted in Figure 2 (left), one of the CNNs takes an interleaved bit stream, which creates a long-range memory. This architecture is similar to Turbo codes, shown in Figure 2 (right). TURBOAE, trained via a carefully designed training methodology, mimics the reliability of Turbo codes for AWGN channels [12].

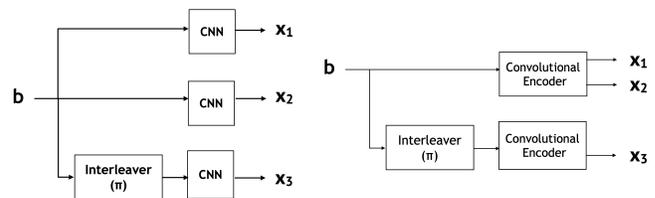


Fig. 2. Encoder of TURBOAE vs. Turbo codes for a rate-1/3 code.

Since [12], TURBOAE architecture has gained huge interest, and similar approaches for designing channel codes have been proposed. Neural codes inspired by serially concatenated Turbo codes and Reed-Muller/Polar codes are introduced in [14] and in [15], respectively. TURBOAE has been applied to channels with feedback [16] and joint coding and modulation [17]. Methods to speed up training are presented in [14].

Despite the success, there are two important open problems. First, how TURBOAE performs on various *practical* non-AWGN channels, such as fading channels and channels with bursty interference, is vastly unknown. Can we leverage the *trainability* of TURBOAE for challenging communications?

Second, can we improve TURBOAE even further? In particular, a *pseudo-random* interleaver is deployed in the encoder of TURBOAE in [12], [14]. Is this an optimal choice? Does interleaver play an important role in the reliability? If so, how can we choose an interleaver that leads to the high reliability? These are challenging problems as the space of interleaver grows faster than exponential in the block length. In this paper, we make progress on both aforementioned problems. Our main contributions are as follows.

- *Development of Turbo Autoencoder with a Trainable Interleaver (TURBOAE-TI):* We introduce a penalty func-

tion to make the interleaver *trainable* and provide a well-thought-out optimization methodology to *learn* the interleaver together with the rest of TURBOAE.

- *Adaptivity and robustness:* We demonstrate that TURBOAE-TI outperforms TURBOAE and LTE Turbo codes on several channels of practical interest, such as fading channels and bursty noise channels. We show that we gain in reliability up to 1dB.
- *Performance on AWGN channels:* We show that TurboAE-TI outperforms TurboAE and LTE turbo codes for the classical AWGN channels. We develop a novel training methodology, called *Rician training*, which is crucial in achieving the high reliability.
- *Interpretation:* We provide analysis to interpret the behavior of learned codes and interleavers. Our analysis implies that TURBOAE-TI has learned a sophisticated pattern that is beyond maximizing the codeword distance.

In the rest of this paper, we refer to TURBOAE by TURBOAE-UI (TurboAE with a Uniform Interleaver), to be clear about their interleaver design, which is generated uniformly at random.

## II. BACKGROUND: TURBO AUTOENCODER

In this section, we review Turbo Autoencoder (TURBOAE-UI), one of the state-of-the-art neural network-based channel codes introduced in [12]. TURBOAE-UI is inspired by Turbo codes, which concatenates a convolutional code and another convolutional code generated from the *interleaved* bit stream.

**Encoder.** As shown in Figure 2, the neural encoder of TURBOAE-UI combines the interleaver and convolutional neural networks (CNNs); it consists of three trainable CNN encoding blocks followed by a power normalization layer. The input to the first two CNN encoding blocks is the original bit sequence  $\mathbf{b}$ , while the input to the last CNN block is an interleaved bit sequence  $\pi(\mathbf{b})$ . Each of the 3 CNN blocks consists of 2 layers 1-D CNN followed by a linear layer. We use 100 filters, a kernel of size 5, and a stride of 2. The output  $\mathbf{x}$  of each CNN encoder block is of shape  $(L, 1)$ .

**Interleaver.** One of the key components of TURBOAE-UI encoder is the interleaver  $\pi$ . [12], [14] deploy a *pseudo-random* interleaver, generated uniformly at random. In Section IV, we empirically show that this selection is *sub-optimal*.

**Decoder.** The neural decoder of TURBOAE-UI is depicted in Figure 3. It consists of several layers of CNNs with interleavers and de-interleavers in between. As received codewords are encoded from original message  $\mathbf{b}$  and interleaved message  $\pi(\mathbf{b})$ , decoding interleaved code requires iterative decoding on both interleaved and de-interleaved order shown in Figure 3. Let  $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3 \in \mathbb{R}^L$  denote noisy versions of  $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 \in \mathbb{R}^L$ , respectively. The decoder generates an estimate  $\hat{\mathbf{b}} \in \{0, 1\}^L$  based on  $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3$  via multiple iterations of CNNs. Two types of decoders are alternatively applied for  $P$  iterations,  $\text{CNN}_i^j$  for  $j \in \{1, 2\}$  and  $i \in \{1, 2, \dots, P\}$ , with an interleaver and de-interleaver in between. Each  $\text{CNN}_i^j$  block consists of 5 layers 1-D CNN followed by a linear layer. We use 100 filters, a kernel of size 5, and a stride of 2. For  $i \in \{1, 2, \dots, 11\}$ , the output posteriors  $\mathbf{q}_i^1$  and  $\mathbf{q}_i^2$  of each decoder block  $\text{CNN}_i^j$  are of shape  $(L, 5)$ , and for  $i = 12$ , the output posterior  $\mathbf{q}_P^2$  is of shape  $(L, 1)$ . The last iteration generates  $\hat{\mathbf{b}} = \text{sigmoid}(\pi^{-1}(\mathbf{q}_P^2))$ .

## III. PROBLEM SETUP

We design channel codes and their corresponding decoders via deep learning. As illustrated in Figure 1, we model the channel encoder and channel decoder as neural networks and train them jointly via an autoencoder training. For concreteness, we focus on rate-1/3 codes with  $L = 40$  information bits (We consider  $L = 100$  in Section VIII). We aim to minimize the Bit Error Rate (BER), defined as  $\text{BER} = \frac{1}{L} \sum_{l=1}^L P(\hat{\mathbf{b}}_l \neq \mathbf{b}_l)$ . We consider various fading channels and bursty noise channels. In particular, we focus on the following models:

- **Rician** ( $K, \sigma$ ): A general description of a Rician fading channel is considered in which a channel is comprised of both a line-of-sight (LOS) and non line-of-sight (NLOS) component. The amplitude of the LOS and NLOS components is dictated by  $K$ . The channel is defined as  $y_l = h_l x_l + z_l$ , where  $z_l \sim \mathcal{N}(0, \sigma^2)$  is Gaussian, and

$$h_l = \left| \sqrt{\frac{K}{(K+1)}}(1 + \text{li}) + \sqrt{\frac{1}{(K+1)}} h_l^{\text{NLOS}} \right|,$$

and  $h_l^{\text{NLOS}}$  is distributed as  $\mathcal{CN}(0, I)$ .

- **Rician with bursty noise** ( $\sigma, \sigma_b$ ): In this bursty channel model, we proceed in two successive steps. In step 1, we apply fading and add AWGN noise to our coded bits:  $y_l = h_l x_l + z_l$ , where  $h_l$  and  $z_l$  are defined as above. In step two, we generate a bursty sequence  $\mathbf{n} \sim \mathcal{N}(0, \sigma_b^2 I) \in \mathbb{R}^S$ , where  $S < L$ . More specifically, we uniformly pick an index  $j \in \{0, 1, \dots, 3L - S - 1\}$  and add bursty noise to our noisy sequence:  $\mathbf{y}[j : j + S] = \mathbf{y}[j : j + S] + \mathbf{n}$ .
- **Rician with linear chirp noise** ( $\sigma_b, f_0, f_1, T, \phi_0$ ): We also evaluate our model on the linear chirp channel, which consists of applying a high-power frequency-increasing sinusoid to our transmitted bits. The jamming signal in function of time  $t$  is given by  $q(t) = \sigma_b^2 [\sin(\phi_0) + 2\pi(\frac{c}{2}t^2 + f_0 t)]$ , where  $\phi_0$  is the initial phase (at time  $t = 0$ ) and  $c = \frac{f_1 - f_0}{T}$  is the chip rate.  $f_0$  is the starting

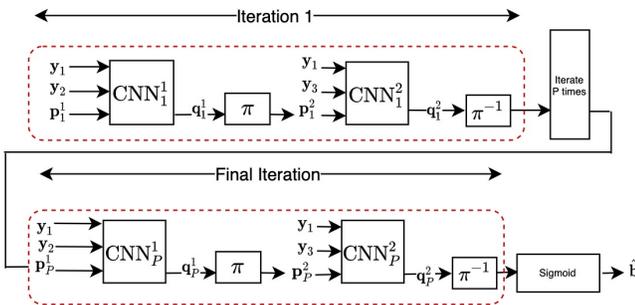


Fig. 3. The decoder architecture in Turbo Autoencoder for a rate-1/3 code.

frequency at time  $t = 0$  and  $T$  is the time it takes to sweep from  $f_0$  to  $f_1$ .

Similarly to the previously described channel, we proceed in two successive steps. In step 1, we apply fading and add AWGN noise to our coded bits:  $y_l = h_l x_l + z_l$ , where  $h_l$  and  $z_l$  are defined as above. In step two, using  $q(t)$ , we generate a jamming sequence  $\mathbf{n} \sim \mathcal{N}(0, \sigma_b^2 I) \in \mathbb{R}^S$ , where  $S < L$ . More specifically, we uniformly pick an index  $j \in \{0, 1, \dots, 3L - S - 1\}$  and add the bursty noise to our noisy sequence:  $\mathbf{y}[j : j + S] = \mathbf{y}[j : j + S] + \mathbf{n}$ .

#### IV. TURBOAE-TI: TURBOAE WITH A TRAINABLE INTERLEAVER

We introduce a neural architecture, called TURBOAE-TI, which combines *interleaver training* and TURBOAE-UI architecture. Learning an interleaver is challenging as the space of interleaver scales in  $L!$ , where  $L$  denotes the block length. To address this challenge, we introduce a carefully-designed training procedure, as well as an interleaver penalty function, which will be optimized in conjunction with our model's loss function. The details are described in Section IV. A-B. As we show in Section V, the *interleaver training* leads to a noticeable improvement upon TURBOAE-UI. For concreteness, we focus on rate-1/3 codes with  $L = 40$  information bits.

##### A. Trainable Interleaver

We propose to learn an interleaver via deep learning, hence getting rid of its manual design, resulting in a more efficient matrix shuffle. The interleaver operation can be represented as multiplying the message sequence  $\mathbf{b}$  by a permutation matrix  $T \in [0, 1]^{L \times L}$ , a square binary matrix with exactly one entry of one in each row and each column and zeros elsewhere. Since it's too costly to enumerate all possibilities, we propose to approach it by adding a matrix generalization of the  $l_1 - l_2$  penalty [18] to our model's loss function during training. The proposed interleaver penalty  $P(T)$  is given by:

$$P(T) = \sum_{i=1}^L \left[ \sum_{j=1}^L |t_{ij}| - \left( \sum_{j=1}^L t_{ij}^2 \right)^{1/2} \right] + \sum_{j=1}^L \left[ \sum_{i=1}^L |t_{ij}| - \left( \sum_{i=1}^L t_{ij}^2 \right)^{1/2} \right]$$

Moreover, to satisfy the permutation matrix's properties, we have to abide by the following conditions:

$$t_{ij} \geq 0, \forall (i, j); \sum_{i=1}^L t_{ij} = 1, \forall j; \sum_{j=1}^L t_{ij} = 1, \forall i \quad (1)$$

The training details are discussed in the following subsection.

##### B. Training Methodology

Let  $L(w_{enc}, w_{dec}, T)$  denote our model's loss function. The training thus minimizes:

$$f = L(w_{enc}, w_{dec}, T) + \lambda \cdot P(T)$$

Where  $w_{enc}, w_{dec}$  are the encoder's and decoder's weights, respectively, and  $\lambda \in \mathbb{R}$  is a regularization constant used to balance the contribution of each component in the loss function. The details are described in Algorithm 1. The main differences from conventional training are as follows:

- At every epoch, encoders and decoders are trained separately. The number of iterations for the encoders and decoders are  $S_{enc}$  and  $T_{Sec}$  respectively. This prevents the training from converging to a local optimum.
- Empirically, our results showed that using different training noise levels for the encoders ( $\sigma_{enc}$ ) and decoders ( $\sigma_{dec}$ ) resulted in better performance.
- The used batch size was very large (500). This is important to average out the channel noise effects.
- To satisfy the three constraints in (1), we first clip the values of  $T$ :  $T \leftarrow \max(T, 0)$ . Then, we normalize each column of  $T$  by dividing each column element by the sum of the entries in this particular column. Finally, we normalize each row of  $T$  by dividing each row element by the sum of the entries in this particular row.

---

#### Algorithm 1: Training of TURBOAE-TI

---

**Inputs:** Number of Epochs  $E$ , Encoder Training Steps  $S_{enc}$ , Decoder Training Steps  $S_{dec}$ , Encoder Training noise  $\sigma_{enc}$ , Decoder Training noise  $\sigma_{dec}$ , Learning Rate  $\alpha$ , Regularizer  $\lambda$

**Outputs:** Encoder and Decoder Weights  $w_{enc}, w_{dec}$ , Interleaver Matrix  $T$

**for**  $e \leq E$  **do**

**for**  $k \leq S_{enc}$  **do**

        Generate examples  $\mathbf{b}$ , noise  $\mathbf{z} \sim N(0, \sigma_{enc}^2)$

        Compute  $f$  using  $\mathbf{b}, \mathbf{z}$

        Generate the gradients  $\nabla_{w_{enc}} f, \nabla_T f$

        Update  $w_{enc}$ :  $w_{enc} \leftarrow w_{enc} - \alpha \cdot \nabla_{w_{enc}} f$

        Update  $T$ :  $T \leftarrow T - \alpha \cdot \nabla_T f$

        Non-negativity constraint:  $T \leftarrow \max(T, 0)$

        Normalize each row and column of  $T$

**end**

**for**  $k \leq S_{dec}$  **do**

        Generate examples  $\mathbf{b}$ , noise  $\mathbf{z} \sim N(0, \sigma_{dec}^2)$

        Compute  $f$  using  $\mathbf{b}, \mathbf{z}$

        Generate the gradients  $\nabla_{w_{dec}} f, \nabla_T f$

        Update  $w_{dec}$ :  $w_{dec} \leftarrow w_{dec} - \alpha \cdot \nabla_{w_{dec}} f$

        Update  $T$ :  $T \leftarrow T - \alpha \cdot \nabla_T f$

        Non-negativity constraint:  $T \leftarrow \max(T, 0)$

        Normalize each row and column of  $T$

**end**

**end**

---

## V. RESULTS AND ANALYSIS

We show that TURBOAE-TI outperforms TURBOAE-UI and LTE Turbo codes for fading channels, bursty noise channels, and AWGN channels. We consider various SNRs, defined as  $\text{SNR} = -10 \log_{10}(\sigma^2)$ . To measure performance, we plot BER vs  $E_b/N_0$ , where  $E_b$  is the energy per bit, and  $N_0$  is the noise power. The relation between SNR and  $E_b/N_0$  is given by  $E_b/N_0 = 10 \log_{10}(\text{SNR}) - 10 \log_{10}(\text{rate})$ . The rate is fixed as  $1/3$ . Interpretation analysis is followed in Section VI.

### A. Fading channels

We consider Rician channels described in Section III (for  $K = 10$ ) in Figure 4 (top), and Rayleigh channels, for which there is no line of sight component (i.e., Rician channel with  $K = 0$ ) in Figure 4 (bottom). TURBOAE-TI, coupled with a neural interleaver and a personalized training, shows a clearly better performance when compared to TURBOAE-UI and LTE Turbo in the low-to-moderate  $E_b/N_0$  regime.

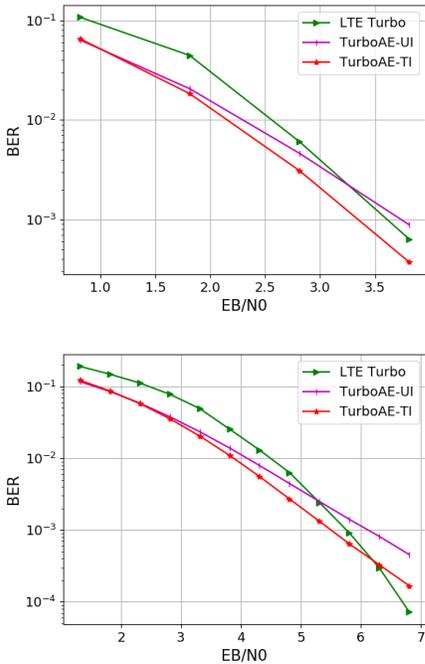


Fig. 4. TURBOAE-TI outperforms TURBOAE-UI and LTE Turbo codes for Rician channels ( $K = 10$ ) (top) and Rayleigh channels (bottom) on low-to-moderate  $E_b/N_0$  regimes.

### B. Adaptivity and Robustness

In this subsection, we analyze the *robustness* and *adaptivity* of TURBOAE-TI and TURBOAE-UI. *Robustness* refers to the ability of a network trained for a particular channel model to work well on a differing channel model without re-training. *Adaptivity* refers to the ability of the network to adapt and retrain for differing channel models with minimal re-training.

To measure these two metrics, we conduct two experiments. In the first experiment, we consider applying a bursty noise to the Rician channel. The setup is described in Section III (Rician with bursty noise). The results are shown in Figure 5.

First, to measure *robustness*, we consider training TURBOAE-TI and TURBOAE-UI on Rician channels ( $K = 10$ ), and testing those models on Rician with bursty noise. Although the models show an improved performance to LTE Turbo, the gap is not significant. Moreover, TURBOAE-UI outperforms TURBOAE-TI, which is not a satisfactory result. To address those issues, we propose *fine-tuning* our models, by training them on Rician channels coupled with bursty noise, for 100 epochs. (TURBOAE-TI finetuned and TURBOAE-UI finetuned). After *fine-tuning* for as little as 100 epochs, (a) the performance of TURBOAE-UI and TURBOAE-TI improves dramatically compared to LTE Turbo, and (b) TURBOAE-TI performs better than TURBOAE-UI. Those results highlight the *adaptivity* of TURBOAE-TI.

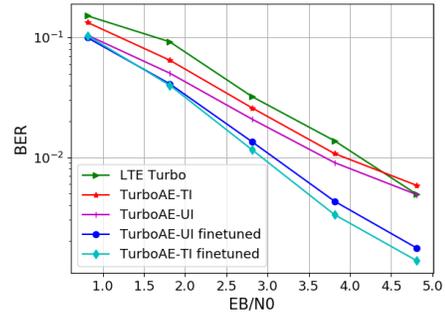


Fig. 5. Fine-tuned TurboAE-TI outperforms TURBOAE-UI and LTE Turbo for Bursty Rician channels ( $K = 10$ ) on low-to-moderate  $E_b/N_0$  regimes.

In the second experiment, we aim to measure the *robustness* of our models to a different type of bursty noise: chirp jamming signals (described in Section III). We use our previously described models (TURBOAE-TI finetuned and TURBOAE-UI finetuned), and test them on the Rician channel with chirp jamming. The results are shown in Figure 6. We notice that both TURBOAE-TI and TURBOAE-UI show good *robustness* to the modified channel when compared to LTE Turbo, and TURBOAE-TI finetuned has a slightly better performance in the high  $E_b/N_0$  regime compared to TURBOAE-UI.

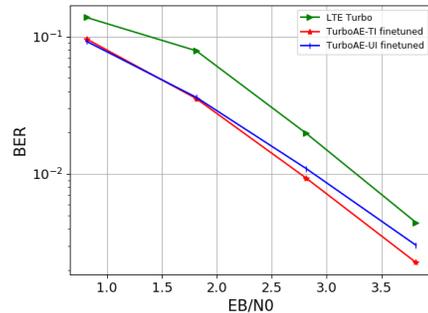


Fig. 6. Robustness of TURBOAE-TI and TURBOAE-UI on Rician fading channels with chirp noise

### C. AWGN channels

Given the reliability improvement of TURBOAE-TI upon TURBOAE-UI on fading and bursty noise channels, a natural

question is whether TURBOAE-TI outperforms TURBOAE-UI on AWGN channels, where  $\mathbf{y} = \mathbf{x} + \mathbf{z}$ , where  $\mathbf{z} \sim \mathcal{N}(0, \sigma^2 I)$  is the IID Gaussian noise.

We train and test both TURBOAE-TI and TURBOAE-UI on AWGN channels in Figure 7, from which we observe that TURBOAE-TI (-▽-) performs worse than TURBOAE-UI (-\*-). Furthermore, both perform worse than LTE Turbo codes. We conjecture this is because training an interleaver along with the rest of TURBOAE-UI is challenging.

To mitigate such challenge in training TURBOAE-TI, we apply the idea of *Rician training* inspired by the promising results on fading channels. Instead of training TURBOAE-TI on AWGN channels, we train it on Rician fading channels (with  $K = 10$ ). As shown in Figure 7, TURBOAE-TI trained for Rician fading channels perform noticeably better than other TURBOAE-UI models. On the other hand, Rician training does not improve the reliability of TURBOAE-UI.

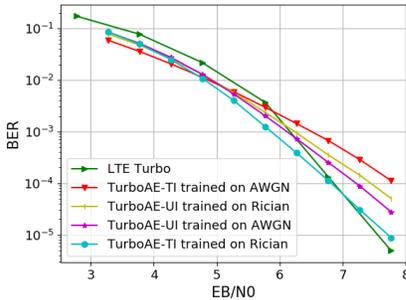


Fig. 7. TurboAE-TI trained on Rician fading channels outperforms TurboAE-UI and LTE Turbo codes for AWGN channels.

## VI. INTERPRETATION

A natural question is ‘what has TURBOAE-TI learned?’. Although interpreting the behavior of deep learning models is in general challenging, we run several experiments to better understand the behavior of our trained network. The results are described below.

**Visualization of the learned interleaver.** In Fig. 8, we visualize the  $40 \times 40$  learned permutation matrix  $T$  of TURBOAE-TI trained on the Rician channel. Yellow (purple) squares denote the positions of 1’s (0’s). We conclude we learned a legitimate interleaver, given that  $T$  satisfies constraints in (IV-A).

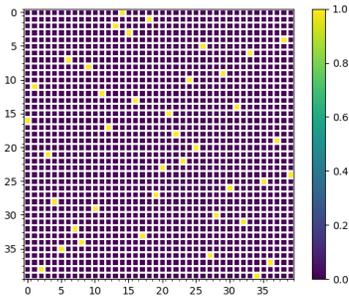


Fig. 8. Learned Interleaver  $T$  for Rician Channel Training

**Effect of Rician training on TURBOAE-TI.** To further understand the behavior of TURBOAE-TI trained on Rician

fading channels, in Figure 9, we plot the test BER performance on AWGN channel as a function of training epochs, comparing both training on Rician and AWGN channels when SNR = 1 dB. Rician training shows stable improvement while AWGN training tends to get stuck at local sub-optima.

Given the general channel form as  $\mathbf{y} = h\mathbf{x} + \mathbf{z}$ , the gradient of loss  $L(\cdot)$  with respect to the encoder weight  $\theta$  is:

$$\frac{\partial L}{\partial \theta} = \frac{\partial L}{\partial \mathbf{y}} h \frac{\partial \mathbf{x}}{\partial \theta} \quad (2)$$

When training with AWGN channel, the  $h$  is constant. For Rician  $K = 10$ , the realization of Rician fading channels is similar to the realization of AWGN channels, but the channel coefficient  $h$  includes more perturbation. Based on this result, we conjecture that Rician training helps TURBOAE-TI to escape the local sub-optima, via injecting gradient perturbations for training encoders.

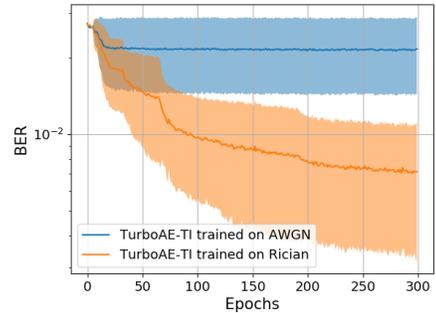


Fig. 9. Test BER along training epochs, averaged over 5 runs: Rician training leads to convergence to a lower BER.

**Partial Minimum Distance.** Additionally, we compute the minimum distance between codewords. Since empirically computing minimum distance is intractable even for moderate block length, we propose an alternative method to evaluate the partial minimum distance. The notation is:  $i$ -th message  $u_i$  of length  $L = 40$ , with  $M = 10$ , is composed as  $u_i = [a, u_i^M, b]$ , where  $a$  and  $b$  are the fixed random binary message of length  $(L - M)/2$  across each enumerations.  $u_i^M$  enumerates all possible message of length  $M$ . We enumerate all  $2^M$  messages, and compute their partial minimum Euclidean distance:

$$D_{min} = \min_{i \in \{1, \dots, 2^M\}, j \in \{1, \dots, 2^M\}, i \neq j} D(f(u_i), f(u_j))$$

The distance is also averaged over 100 instances. The results are shown in Figure 10. For TURBOAE-TI, Rician training with  $K = 10$  results in the best partial minimum distance, indicates that Rician training leads to better encoder. On the other hand, TURBOAE-UI trained on Rician channel lead to worse performance. We conjecture that Rician training is more beneficial for TURBOAE-TI, due to its sub-optima problem. Note that even the best TURBOAE-UI partial minimum distance is still worse than regular Turbo code. We conjecture TURBOAE-UIs learned sophisticated codes with patterns that are not fully captured in the codeword distance. On the other

hand, this result shows that a further improvement might be feasible by deploying a regularizer on the codeword distance.

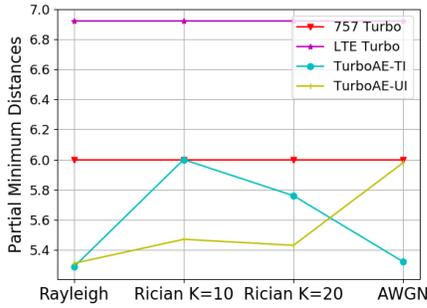


Fig. 10. Minimum distance of TURBOAE-TI and TURBOAE-UI trained on various channels: Rician training of TURBOAE-TI leads to a larger minimum codeword distance than AWGN training. However, the minimum distance of TURBOAE-TI is smaller than the one of LTE Turbo codes, which implies that TURBOAE-TI’s superior performance is not fully captured in the distance.

## VII. CONCLUSION

We introduce TURBOAE-TI, where we make the interleaver of TURBOAE-UI to be trainable by adding a matrix generalization of the  $l_1 - l_2$  penalty to our model’s loss function, and introducing a well-thought-out training methodology. We show that the trainable interleaver leads to the improved reliability on various channels, such as fading channels and bursty noise channels. The amount of gain varies from channel to channel and reaches up to 1dB.

TURBOAE-TI also outperforms TURBOAE-UI for AWGN channels, when combined with the novel Rician training methodology. Our analysis suggests that the perturbation of channel weights in the Rician training helps the TURBOAE-TI to escape the local optima. Moreover, Rician training also results in a better encoder, indicated by the improved partial minimum distance.

## VIII. OPEN PROBLEMS

There are several open problems that arise from this work. First, the extension of TURBOAE-TI to *multi-user scenarios* is an interesting future direction. In recent work [19], it is shown that designing a code that utilizes interleavers for interference channels is challenging. With the aid of a trainable interleaver, we conjecture that one can learn a code with a long range memory for such *multi-user scenarios* and hence achieve a better performance.

Moreover, developing deep learning methodologies driven by the communication theory is a promising direction. We conjecture that the Rician training framework can be applied to potentially enhance the performance of some state of the art applications (e.g., computer vision).

Finally, extending the TURBOAE-TI framework to longer blocklengths and higher SNRs is challenging, as the interleaver’s space grows in function of the blocklength, and the training becomes harder for higher SNRs. We illustrate our results on  $L = 100$  in Figure 11. Although TURBOAE-TI works better than TURBOAE-UI, it performs worse than

Turbo LTE. This indicates that the scalability of TURBOAE-TI to a long blocklength is an interesting future research direction.

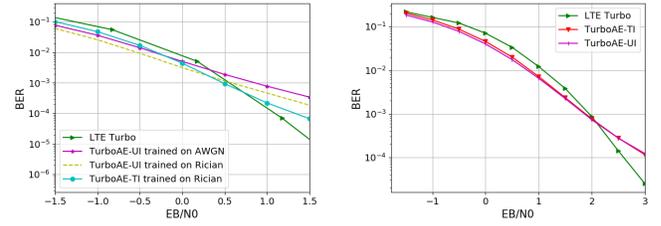


Fig. 11. Results on L=100 AWGN (left) and Rician fading channels (right)

## REFERENCES

- [1] C. E. Shannon, “A mathematical theory of communication, part i, part ii,” *Bell Syst. Tech. J.*, vol. 27, pp. 623–656, 1948.
- [2] T. J. O’Shea and J. Hoydis, “An introduction to machine learning communications systems,” *CoRR*, vol. abs/1702.00832, 2017.
- [3] S. Dörner, S. Cammerer, J. Hoydis, and S. t. Brink, “Deep learning-based communication over the air,” *arXiv*, 2017.
- [4] T. Gruber, S. Cammerer, J. Hoydis, and S. ten Brink, “On deep learning-based channel decoding,” in *Information Sciences and Systems (CISS), 2017 51st Annual Conference on*. IEEE, 2017, pp. 1–6.
- [5] A. Alkhateeb, S. Alex, P. Varkey, Y. Li, Q. Qu, and D. Tujkovic, “Deep learning coordinated beamforming for highly-mobile millimeter wave systems,” *IEEE Access*, vol. 6, pp. 37 328–37 348, 2018.
- [6] S. Park, H. Jang, O. Simeone, and J. Kang, “Learning how to demodulate from few pilots via meta-learning,” in *2019 IEEE SPAWC*, 2019, pp. 1–5.
- [7] H. Ye, G. Y. Li, and B.-H. Juang, “Power of deep learning for channel estimation and signal detection in ofdm systems,” *IEEE Wireless Communications Letters*, vol. 7, no. 1, pp. 114–117, 2017.
- [8] T. Wang, C.-K. Wen, S. Jin, and G. Y. Li, “Deep learning-based csi feedback approach for time-varying massive mimo channels,” *IEEE Wireless Communications Letters*, vol. 8, no. 2, pp. 416–419, 2019.
- [9] E. Nachmani, E. Marciano, L. Lugosch, W. J. Gross, D. Burshtein, and Y. Be’ery, “Deep learning methods for improved decoding of linear codes,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 119–131, 2018.
- [10] N. Shlezinger, N. Farsad, Y. C. Eldar, and A. J. Goldsmith, “Viterbinet: Symbol detection using a deep learning based viterbi algorithm,” *IEEE 20th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2019.
- [11] Y. He, J. Zhang, S. Jin, C.-K. Wen, and G. Y. Li, “Model-driven DNN decoder for turbo codes: Design, simulation, and experimental results,” *IEEE Transactions on Communications*, vol. 68, no. 10, 2020.
- [12] Y. Jiang, H. Kim, H. Asnani, S. Kannan, S. Oh, and P. Viswanath, “Turbo autoencoder: Deep learning based channel code for point-to-point communication channels,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [13] H. Kim, Y. Jiang, S. Kannan, S. Oh, and P. Viswanath, “Deepcode: Feedback codes via deep learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 9436–9446.
- [14] J. Clausius, S. Dörner, S. Cammerer, and S. ten Brink, “Serial vs. parallel turbo-autoencoders and accelerated training for learned channel codes,” vol. abs/2104.14234, 2021.
- [15] A. V. Makkuba, X. Liu, M. V. Jamali, H. Mahdaviyar, S. Oh, and P. Viswanath, “Ko codes: inventing nonlinear encoding and decoding for reliable wireless communication via deep-learning,” in *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.
- [16] Y. Jiang, H. Kim, H. Asnani, S. Oh, S. Kannan, and P. Viswanath, “Feedback Turbo autoencoder,” in *2020 IEEE ICASSP*, 2020.
- [17] Y. Jiang, H. Kim, H. Asnani, S. Kannan, S. Oh, and P. Viswanath, “Joint channel coding and modulation via deep learning,” in *2020 IEEE SPAWC*, 2020, pp. 1–5.
- [18] E. Esser, Y. Lou, and J. Xin, “A method for finding structured sparse solutions to nonnegative least squares problems with applications,” *SIAM Journal on Imaging Sciences*, vol. 6, no. 4, pp. 2010–2046, 2013.
- [19] H. K. Karl Chahine, Nanyang Ye, “DeepIC: Coding for interference channels via deep learning,” *GlobeCom*, 2021.