

# Integrating Pre-Trained Language Model With Physical Layer Communications

Ju-Hyung Lee<sup>1</sup>, Dong-Ho Lee, Joohan Lee<sup>2</sup>, and Jay Pujara

**Abstract**—The burgeoning field of on-device AI communication, where devices exchange information directly through embedded foundation models, such as language models (LMs), requires robust, efficient, and generalizable communication frameworks. However, integrating these frameworks with existing wireless systems and effectively managing noise and bit errors pose significant challenges. In this work, we introduce a practical on-device AI communication framework, integrated with physical layer (PHY) communication functions, demonstrated through its performance on a link-level simulator. Our framework incorporates end-to-end training with channel noise to enhance resilience, incorporates vector quantized variational autoencoders (VQ-VAE) for efficient and robust communication, and utilizes pre-trained encoder-decoder transformers for improved generalization capabilities. Simulations, across various communication scenarios, reveal that our framework achieves a 50% reduction in transmission size while demonstrating substantial generalization ability and noise robustness under standardized 3GPP channel models.

**Index Terms**—Physical layer communications, language model, VQ-VAE, natural language processing (NLP), link-level simulation.

## I. INTRODUCTION

THE increasing capabilities of mobile devices and the advancements in large language model (LLM), particularly foundation models, have paved the way for a new era of on-device artificial intelligence (AI).<sup>1</sup> This paradigm shift allows devices to possess their own AI capabilities, enabling tasks such as real-time translation, personalized recommendations, and even autonomous driving, all while ensuring privacy and efficiency [1], [2], [3], [4].

Manuscript received 14 February 2024; revised 9 June 2024; accepted 17 August 2024. Date of publication 9 September 2024; date of current version 13 November 2024. This work was supported in part by the National Science Foundation (NSF) under Grant 2008443. The associate editor coordinating the review of this article and approving it for publication was B. Guler. (Corresponding author: Ju-Hyung Lee.)

Ju-Hyung Lee was with the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California (USC), Los Angeles, CA 90007 USA. He is now with Nokia, Sunnyvale, CA 94085 USA (e-mail: juhyung.lee@outlook.com).

Dong-Ho Lee and Jay Pujara are with the Information Science Institute, University of Southern California (USC), Los Angeles, CA 90007 USA (e-mail: dongho.lee@usc.edu; jpujara@usc.edu).

Joohan Lee is with the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California (USC), Los Angeles, CA 90007 USA (e-mail: joohanl@usc.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TWC.2024.3452481>.

Digital Object Identifier 10.1109/TWC.2024.3452481

<sup>1</sup>Throughout this paper, we will use the term “on-device AI” and “on-device AI/language model (LM)” interchangeably to refer to a foundation model or a language model deployed at both transmitter (TX) side and receiver (RX) side in a distributed manner.

On-device AI communication enables devices equipped with advanced foundation models like LMs to directly exchange information, forming a network of distributed intelligence. This distributed AI-to-AI communication requires efficient and accurate transmission of information between devices. The “encoded representation” generated by the AI/LLM at the TX side (called *AI-Src-Enc*) needs to be reliably transmitted and accurately interpreted as “decoding input” by the AI/LLM at the RX side (called *AI-Src-Dec*) [5].

However, seamlessly integrating this novel technology with existing communication infrastructure presents a significant challenge. Maintaining backward compatibility is a crucial requirement in wireless communication systems, allowing devices designed for older networks (e.g., 4G and 5G) to operate seamlessly with newer technologies (e.g., 6G). This interoperability ensures a smooth transition between networks, such as handover and roaming, enhancing user experience. For on-device AI communication, ensuring backward compatibility is still essential, enabling existing communication infrastructure to support this new paradigm without requiring major modifications. Thus, it is required to integrate the on-device AI model application into wireless communication systems, ensuring they adhere to existing legacy protocols for seamless interoperability [6].

Furthermore, effective on-device AI communication demands several key qualities: (1) *Robustness*: The system must be resilient to communication errors (e.g., bit error) and noise (e.g., channel fading) inherent in dynamic wireless environments; (2) *Efficiency*: The system must efficiently compress and transmit data while maintaining accuracy and fidelity; and (3) *Generalization Capability*: The system should be able to handle diverse inputs and adapt to various communication scenarios. Meeting these requirements presents another challenge, requiring novel approaches in system design and optimization.

## A. Related Works

1) *Semantic Communication*: Semantic communication, which focuses on effectively transmitting semantic (meaningful) information rather than merely receiving individual symbols or bits [7], is the most relevant research direction for distributed AI-to-AI communications. In order to accomplish this, semantic communication systems leverage neural network (NN) at both the TX and RX to extract and decode the semantic information. Semantic communication research can

be classified based on the types of data they utilize, including image [8], [9], [10], video [11], [12], [13], and speech/text data [14], [15], [16]. These works pioneered the inclusion of AI in source coding within transceivers, optimizing joint source-and-channel coding (JSCC) under specific noisy channels and data distributions.

However, such semantic communication framework poses significant challenges for practical application in on-device AI communication. Firstly, it demands major modifications to existing legacy communication protocols (e.g., 5G-NR protocol), distinguishing it from practical distributed AI-to-AI communication systems. JSCC necessitates that both channel coding and communication modules be modified and trainable, requiring major modifications in existing legacy communication protocols that are impractical in conventional wireless communication systems.

Secondly, much of the existing research assumes ideal channel conditions (e.g., additive white Gaussian noise (AWGN)) and bases analysis on information theory and symbol-level transmission. This overlooks the complexities of non-ideal (realistic) channel conditions and the necessity for link (or system)-level analysis and practical bit-level transmission.

Advancing the pioneering efforts in semantic communication towards practical application in on-device AI communication necessitates two critical components: (1) *Reliable* communication systems need to seamlessly operate with existing infrastructure. This typically involves interacting with various physical-layer communication function modules like mappers, forward error correction (FEC), and orthogonal frequency-division multiplexing (OFDM), each serving a distinct purpose. Therefore, it becomes necessary to integrate AI-Src-Enc and AI-Src-Dec within the transceiver system, with the operational parameters of modules compatible with, or analogous to, those in 5G-NR physical-layer modules. (2) *Efficient* communication systems should be able to compress the output of AI-Src-Enc and accurately decode the input of AI-Src-Dec. Developing an effective compression approach is crucial for reducing the size of semantic messages while retaining their essential information.

Fig. 1 shows the difference between semantic communication and on-device AI communication systems.

2) *Pre-Trained Language Model*: Recent studies on text semantic communication exploit transformer architectures [17] to extract semantics at the TX and recover the original information at the RX [14], [18]. However, such frameworks may have the following challenges: (i) training an end-to-end (E2E) semantic communication pipeline requires a huge computational effort due to the many randomly initialized parameters of semantic encoder/decoder to be trained on; (ii) difficulty in handling out-of-vocabulary (OOV) since they only use a set of whitespace-separated tokens in the training data. *General-purpose* communication systems should be able to effectively handle any data, including out-of-domain (OOD) data that lies outside the training data used to develop the framework. This necessitates using pre-trained language models with superior generalizability, allowing them to perform well even when encountering unforeseen data variations.

Pre-trained language models, which have been extensively trained on vast web datasets, are revolutionizing the field of natural language processing (NLP) [19], [20]. These models, often referred to as “foundation models”, possess the ability to effectively generalize even to data where they have not specifically been trained for [2]. Their capabilities are a result of leveraging transfer learning and scaling techniques. Transfer learning involves utilizing knowledge acquired from one task and applying it to another [21]. In the context of deep learning, pre-training is the prevailing approach to transfer learning, where a model is initially trained on a surrogate task and subsequently fine-tuned to adapt to the downstream task of interest.

## B. Contributions

To design and evaluate a practical and widely applicable on-device AI communication framework, we focus on three key questions:

- Q1:** *How can we integrate a pre-trained language model in practical physical layer (PHY) communication systems and evaluate its reliability in realistic network scenario?*
- Q2:** *How can we efficiently compress the output of the AI-Src-Enc and accurately decode the noisy input of the RX AI-Src-Dec?*
- Q3:** *How can we leverage the knowledge gained from a large dataset to build a general-purpose encoder/decoder for system-wide application?*

Our main contributions, which address these challenges, are summarized as follows:

- We integrate a language model with the link-level simulator, NVIDIA Sienna [22], incorporating 5G-NR PHY communication functions (e.g., Polar channel coding and QAM mapper). We evaluate its performance on a channel that contains not only noise but also delay dispersion (e.g., 3GPP CDL-family channel), seeing the efficacy of on-device AI communication. Furthermore, we propose a *noise-tuning* method to optimize its reliability. This addresses **Q1**.
- Converting on-device AI information (e.g., output of AI-Src-Enc and input of AI-Src-Dec) into bit-level information often increases data load. Additionally, even minor bit errors, induced by channel noise, can greatly hinder communication. Addressing this, we propose a novel approach that leverages vector quantization techniques [23] for efficient compression and decoding of on-device AI information. This approach transforms high-dimensional vectors into discrete data via a codebook, significantly reducing transmission overhead even while mitigating the impact of bit errors (i.e., robustness against channel noise). This addresses **Q2**.
- We incorporate encoder-decoder transformers into E2E on-device AI communication systems, initializing them with pre-trained weights. This approach reduces the reliance on channel-specific and data distribution-specific optimization, thus enhancing the system’s generalizability; This addresses **Q3**. In particular, we employ a pre-trained encoder-decoder transformer (BART [24]) to

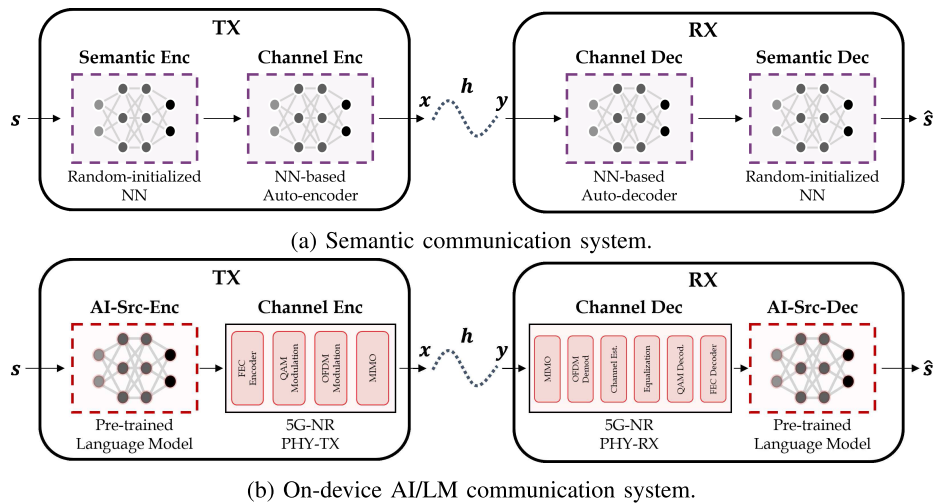


Fig. 1. Overview of semantic communication systems and on-device AI/LM communication system.

initialize the parameters of AI-Src-Enc/Dec so that the pipeline itself requires little or no computational effort and use a pre-trained tokenizer to effectively handle OOV so that our pipeline can be generalized to any other text.

- We release source code for the experiments to promote reproducible ML research in wireless communication.<sup>2</sup>

### C. Paper Organization and Notation

This paper is organized as follows: Sec. II provides the background information on two key concepts: *on-device AI communication* and *vector quantised variational auto encoder (VQ-VAE)*. After presenting a detailed system model for the on-device AI communication systems in Sec. III, Sec. IV provides comparison results and ablation studies to assess the contribution of each of our proposed approaches, with a specific emphasis on the accuracy and efficiency of text transmission. Then, we also present an in-depth analysis of our proposed *noise-tuning* methods, followed by concluding remarks in Sec. VI.

*Notation:* Random variables are denoted by capital italic font, e.g.,  $X, Y$ , with realizations  $x, y$ , respectively.  $I(X; Y)$ ,  $p(y|x)$  and  $p(x, y)$  represent the mutual information, conditional probability, and joint probability distribution of the two random variables  $X$  and  $Y$ . Multivariate random variables are represented with capital bold font, e.g.,  $\mathbf{Y} = [Y_0, Y_1]^T$ . Vectors are represented using a lowercase bold font, e.g.,  $\mathbf{y}$ . We use  $\mathbb{R}^{D \times 1}$  to represent the  $D$ -dimensional space of real-valued vectors. We also use  $\|\cdot\|$  to denote the  $L^2$ -norm, which is an Euclidean norm.

## II. BACKGROUND

### A. On-Device AI (Distributed AI-to-AI) Communication

On-device AI (distributed AI-to-AI) communication refers to a paradigm where AI models, e.g., LM, are embedded locally within devices and communicate with each other over a network. This decentralized approach allows devices

to exchange complex, context-rich information, leveraging the capabilities of AI models to facilitate personalized and intelligent interactions without relying on centralized servers.

In the process of AI-to-AI communication using via LM, the TX side initially encodes the input sequence. This encoded data is then sent across through the physical layer to the RX side, where it is decoded back into its original form. To enable this, several key components are required. Both AI-Src-Enc at the TX and AI-Src-Dec at the RX must employ a shared tokenizer and embedder. These tools play a pivotal role in accurately encoding and decoding the data. Initially, the tokenizer divides the input sequence into discrete tokens. Subsequently, each token is transformed into an embedding representation through the embedder. This series of mapped embeddings is then input into the encoder to generate a transmittable representation. At the RX end, to reconstruct the original sequence from this representation, the same embedder is used to deduce the token index, which is then converted back to the corresponding token using the tokenizer.

The transmittable representation itself is a high-dimensional vector made up of floating point numbers. For its transmission via the physical layer, this vector needs to be converted into a form of bit-level information, suitable for digital transmission. However, these floating-point numbers are particularly susceptible to bit-flips, which can occur due to noise within the physical layer as mentioned in Q2. For instance, a 32-bit floating-point number representing 1.0 could change to infinity with just one bit-flip (bit error) in the second position, leading to a significant decline in transmission accuracy [25].

### B. Vector Quantized Variational Autoencoders (VQ-VAE)

VQ-VAE are a version of variational autoencoders [26], distinguished by their use of vector quantization techniques to effectively encode information into discrete latent representations [23]. At the core, an autoencoder is a neural network that identifies latent spaces, which are complex and non-linear functions derived from the data. This neural network architecture is divided into encoder and decoder. The encoder's role is to process and transform the input data into a latent vector

<sup>2</sup><https://github.com/abman23/on-device-ai-comm>

representation. Following this, the decoder is tasked with the accurate reconstruction of the original data, using only the latent vector representation provided by the encoder. Unlike traditional approaches where this latent representation is continuous [26], the VQ-VAE creates discrete latent representation by incorporating a discrete codebook. This codebook, forming an array of vector entries each assigned a specific index, is utilized to discretize the latent space within the autoencoder. By implementing this strategy, the VQ-VAE facilitates the learning of a discrete latent space that effectively captures key features of the data. The resultant latent representation in this model is a sequence of integer indices corresponding to the codebook entries, which not only allows for more efficient data compression but also maintains a level of reconstruction quality that is comparable to traditional methods.

### III. ON-DEVICE AI COMMUNICATION SYSTEMS

#### A. Problem Description

Consider a sentence  $s$  that maps to a symbol stream  $x$  using the AI-Src-Enc  $S_E(\cdot)$ , vector quantization  $Q_E(\cdot; Z)$ , and channel encoder  $C_E(\cdot)$  as follows:

$$x = C_E(Q_E(S_E(s); Z)), \quad (1)$$

where  $Z$  is the shared codebook. The symbol stream  $x$  passes through a physical channel  $h$  with noise  $n \sim \mathcal{CN}(0, \sigma_n^2)$ , resulting in the received signal  $y$ :

$$y = hx + n. \quad (2)$$

At the receiver (RX) end, the received signal  $y$  is processed by the channel decoder  $C_D(\cdot)$ , vector dequantizer  $Q_D(\cdot; Z)$ , and AI-Src-Dec  $S_D(\cdot)$  to estimate the original sentence  $\hat{s}$ :

$$\hat{s} = S_D(Q_D(C_D(y); Z)). \quad (3)$$

The architecture of the communication system is illustrated in Fig. 2, where the transmitter (TX) consists of the AI-Src-Enc  $S_E(\cdot)$ , vector quantizer  $Q_E(\cdot; Z)$ , and channel encoder  $C_E(\cdot)$ , and the receiver (RX) includes the channel decoder  $C_D(\cdot)$ , vector dequantizer  $Q_D(\cdot; Z)$ , and AI-Src-Dec  $S_D(\cdot)$ . Both TX and RX share the same codebook  $Z$ .

At the TX end, the sentence  $s$  is encoded into a vector representation  $r$  by  $S_E(\cdot)$ , mapped to discrete indices  $t$  via  $Q_E(\cdot; Z)$ , and encoded into the symbol stream  $x$  by  $C_E(\cdot)$ . This process adds redundancy to ensure reliable detection and correction of bit errors.

Conversely, at the RX end, the channel decoder  $C_D(\cdot)$  decodes the received signal  $y$ , the vector dequantizer  $Q_D(\cdot; Z)$  recovers the discrete indices  $\hat{t}$  into vector representation  $\hat{r}$ , and the AI-Src-Dec  $S_D(\cdot)$  decodes this to recover the estimated sentence  $\hat{s}$ .

*On-device AI/LM = (New) Lossy Source Coding:* While conventional (lossy) source coding focuses on compressing the source input by ensuring statistical similarity between the original and reconstructed signals ( $s$  and  $\hat{s}$ ), AI-Src-Enc takes a different approach. The system aims to minimize lexical errors (*i.e.*, lexical similarity) and semantic errors (*i.e.*, semantic similarity) while also reducing the number of bits/symbols retrieved from AI-Src-Enc, thereby achieving compression.

In particular, it prioritizes compressing the embedding size (*i.e.*, the dimensionality of  $r$ ) while maintaining, or even enhancing, both lexical and semantic similarities between the original and reconstructed data.

While leveraging advanced traditional PHY communication system's techniques (*e.g.*, multiple-input multiple-output (MIMO), OFDM, channel coding, etc.) which prioritize achieving low bit error rate (BER) (or symbol error rate (SER)), on-device AI/LM communication system focuses more on preserving the meaning between the original and reconstructed data to ensure successful distributed AI-to-AI (inter-AI) transmission. To this end, we design and evaluate our E2E on-device AI communication system within the context of bit-level PHY communication transmission to see its E2E performance in realistic scenarios.<sup>3</sup> Our rationale for this system is to preserve the meaning and context of the data, even in the presence of realistic channel noise and bit errors.

#### B. System Model

1) *AI-Source-Encoder and AI-Source-Decoder:* Both the AI-Src-Enc  $S_E(\cdot)$  and AI-Src-Dec  $S_D(\cdot)$  are constructed from a series of 6 transformer layers [17]. Each of these layers integrates a self-attention mechanism, positional encoding, and a densely connected layer, fortified by residual connections. Notably, while they share a similar foundational architecture, the specific manner in which the self-attention operates diverges between the encoder and decoder components. This distinction in the self-attention mechanism ensures specialized processing tailored to the unique demands of both encoding and decoding tasks. The encoder employs a fully-visible self-attention strategy, granting the model the capability to focus on any token of the input while the decoder leverages an auto-regressive self-attention mechanism, limiting the model's attention solely to previous outputs, which is more appropriate for the streaming paradigm in which the decoder operates. These architectures can be pre-trained on a large scale corpus by corrupting documents and computing the cross entropy loss between the decoder's output and the original document to learn the model generalizable knowledge [24].

Here, we employ such pre-trained checkpoints (BART-base [24]) and use the encoder and decoder weights to initialize the weights of  $S_E(\cdot)$  and  $S_D(\cdot)$  respectively. For a more detailed understanding of the en/decoder's operation, both  $S_E(\cdot)$  and  $S_D(\cdot)$  share the pre-trained embedding  $\mathcal{E}$  and the pre-trained tokenizer  $\mathcal{T}$ . Once the sentence  $s$  is given to  $S_E(\cdot)$ ,  $\mathcal{T}$  tokenizes  $s$  into tokens  $s_t = [s_{t_1}, s_{t_2}, \dots, s_{t_n}]$  and maps each token to embedding  $s_e = [s_{e_1}, s_{e_2}, \dots, s_{e_n}]$  by  $\mathcal{E}$ . Then,  $S_E(\cdot)$  encodes  $s_e$  into hidden states  $r = [r_1, r_2, \dots, r_n] \in \mathbb{R}^{n \times d_r}$  through multiple transformer layers, and passes it to channel encoder

<sup>3</sup>In this study, our primary focus is on PHY layer functions, and we do not address MAC (or higher) layer functions, like automatic repeat request (ARQ) and hybrid automatic repeat request (HARQ), which aim to ensure error-free transmission. It is important to note that in certain network scenarios—where simpler transmission is required, where transmitted data are latency-sensitive and cannot afford the delay caused by HARQ, or in broadcast situations where HARQ is not applicable, such as in Sidelink for device-to-device (D2D), URLLC, or LEO satellite networks—the retransmission process may be bypassed.

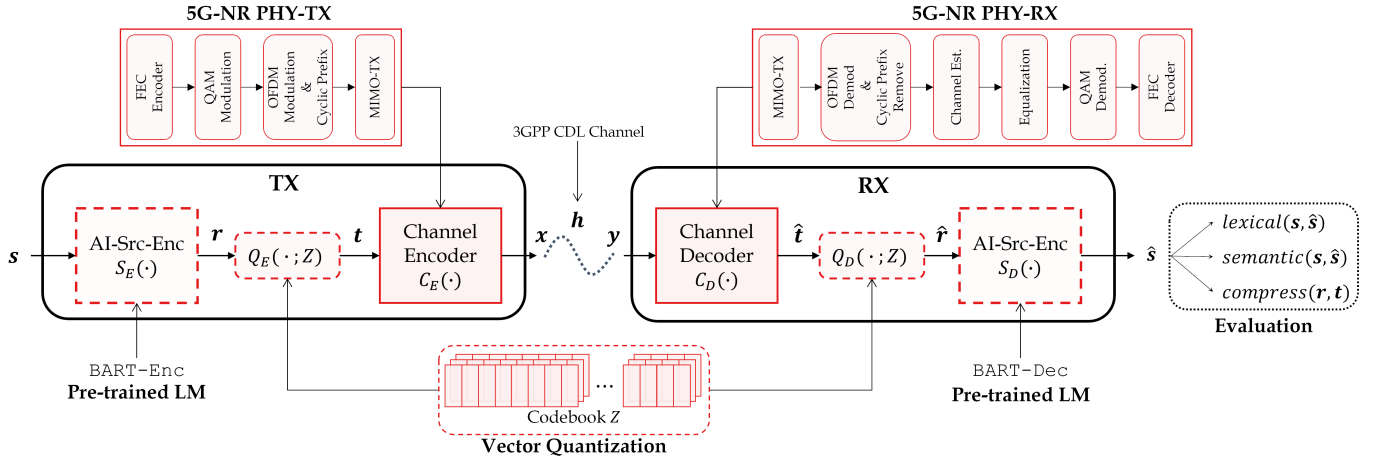


Fig. 2. **Framework overview of on-device AI/LM PHY communication systems** integrated with pre-trained language model. This framework incorporates a link-level simulator to realistically emulate bit-level transmission within PHY communication systems. At the transmission (TX) end, the symbol stream  $s$  undergoes AI-Src-Enc  $S_E(\cdot)$ , vector quantization  $Q_E(\cdot; Z)$ , and channel encoder  $C_E(\cdot)$  to produce  $x$ . Conversely, at the receiver (RX) end, the received signal  $y$  is channel-decoded, vector-dequantized, and semantically decoded to recover the symbol  $\hat{s}$ . We evaluate the system in three different criteria: lexical similarity, semantic similarity, and compression rate.

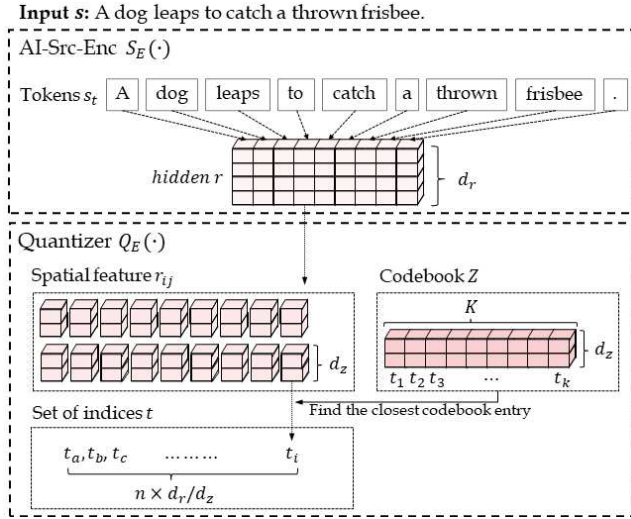


Fig. 3. Architecture of AI-Source-Encoder (AI-Src-Enc) and Quantizer.

$C_E(\cdot)$ . In this context,  $n$  denotes the total number of tokens, while  $d_r$  specifies the dimension of the feature representation of each token.

After transmission is complete and the hidden states  $\hat{r}$  are retrieved from the channel decoder  $C_D(\cdot)$  and provided to  $S_D(\cdot)$ , the semantic decoder  $S_D(\hat{r})$  establishes the conditional distribution  $p_{\theta_{S_D}}(\hat{s}_i | \hat{s}_{0:i-1}, \hat{r})$  and auto-regressively samples words from this distribution for each index.

2) *Vector Quantization (VQ-VAE)*: Our framework employs vector quantization to transform continuous feature representations into discrete indices, enabling efficient data transmission. The core of this process is a discrete codebook

$$Z = \{z_k\}_{k=1}^K \in \mathbb{R}^{K \times d_z}, \quad (4)$$

where  $K$  denotes the codebook size and  $d_z$  the dimension of each code vector.

At the TX end, the continuous feature vector  $r \in \mathbb{R}^{n \times d_r}$ , generated by the AI-Src-Enc  $S_E(\cdot)$ , is quantized into discrete indices  $t \in \mathbb{R}^{n \times 1}$  using  $Z$ . This quantization aims for  $d_z < d_r$  to ensure transmission efficiency. Specifically, each feature vector  $r_i \in r$  is divided into segments  $r_{ij}$ , each with dimension  $d_z$ . For instance, a feature vector  $r_i$  with dimension 768 and a codebook dimension of 256 can be segmented into three parts ( $[r_{i1}, r_{i2}, r_{i3}]$ ), each of dimension 256. The quantizer  $Q_E(\cdot; Z)$  maps each segment  $r_{ij}$  to the closest codebook entry  $z_k \in Z$ , forming the discrete representation  $t_i \in t$ :

$$t_i = Q_E(r_i; Z) = \left( \underset{z_k \in Z}{\operatorname{argmin}} \|r_{ij} - z_k\|^2 \right) \in \mathbb{R}^{(d_r/d_z) \times 1}, \quad (5)$$

where  $(d_r/d_z)$  is the length of the encoded sequence. In the given example,  $(d_r/d_z)$  equals 3.

Conversely, at the RX end, the dequantizer  $Q_D(\cdot; Z)$  converts the received discrete indices  $\hat{t}_i$  back into continuous vectors  $\hat{r}_i$  by directly referencing  $Z$ , reconstructing the feature representation  $\hat{r}_i$  from  $\hat{t}_i$ .

3) *Channel En/Decoder*: The channel encoder  $C_E(\cdot)$  and decoder  $C_D(\cdot)$  incorporate various PHY communication functions, such as Polar code, QAM mapper, OFDM, MIMO, selected for their alignment with 5G-NR standards. This approach ensures that our framework mirrors real-world communication scenarios while acknowledging its partial compliance with 5G-NR protocols.

Firstly, the input bit for  $C_E(\cdot)$  is grouped with other bits to form a codeword. Polar codes, a form of linear block error correction codes, are considered one of the channel coding schemes in 5G-NR, where low complexity error correction is available [27]. It adds redundancy to the codeword with a certain coderate  $\rho$ . This helps detect and correct bit errors introduced during transmission. The codeword is then further processed and segmented for transmission.

Each segmented codeword is mapped to a specific point on a constellation diagram, such as our considered QAM. Higher

QAM schemes offer more bits per symbol (better efficiency) but are more susceptible to noise. The chosen QAM symbol is then represented by a combination of amplitude and phase variations in a carrier signal. The number of bits per symbol used is denoted by  $m$  and the mapper takes as input a message  $u \in \{0, \dots, 2^m - 1\}$

The modulated symbols are divided into subcarriers across a wide spectrum. This distributes the signal energy, making it more resilient to frequency-selective fading channels. Each user (or subframe) can be modulated independently, enabling adaptive bit rate and channel equalization by using the received modulation and coding scheme (MCS) index. The MCS index is dynamically selected and can be adjusted on a per-subframe basis, based on the latest channel quality information (CQI).

Data is transmitted simultaneously leveraging multiple antennas, *i.e.*, MIMO, to exploit spatial diversity or multiplexing. This creates multiple independent channels, increasing reliability, data rate, or both. The processed signal is transmitted over the air. At the receiver, the signal is received by multiple antennas. The received signal is then demodulated (QAM decoding) and demultiplexed (OFDM decoding). Channel coding algorithms (Polar decoding) are applied to correct any errors introduced during transmission. Finally, the decoded bits are reassembled into the original data.

4) *Channel*: In the on-device AI communication, the wireless propagation channel, henceforth simply referred to as *channel*, introduces various types of noise and impairments, which can significantly affect the transmitted signal. These impacts include, but are not limited to, path loss, fading, Doppler shift, and AWGN.

3GPP has defined standardized models that are used for the simulation and testing of various wireless system concepts. These channel models describe time variation, delay dispersion, angular dispersion (at both link ends) and amplitude characteristics. We particularly leverage the standardized 3GPP clustered delay line (CDL)-family channel models, such as CDL-A, CDL-B, and CDL-C [28]. It is worth noting that while these CDL models are useful for testing, particularly for link-level simulation, they often diverge significantly from real-world channels.

### C. Train Objectives

Our framework is jointly trained using two loss functions:  $\mathcal{L}_{\text{CE}}$  and  $\mathcal{L}_{\text{VQVAE}}$ . The  $\mathcal{L}_{\text{CE}}$  minimizes the discrepancy between input sentence  $\mathbf{s}$  at the TX end and its predicted sentence  $\hat{\mathbf{s}}$  at the RX end, employing the cross entropy loss:

$$\mathcal{L}_{\text{CE}}(\hat{\mathbf{s}}, \mathbf{s}) = - \sum_{i=1}^L p(\mathbf{s}_{t_i}) \log(p(\hat{\mathbf{s}}_{t_i})) + (1 - p(\mathbf{s}_{t_i})) \log(1 - p(\hat{\mathbf{s}}_{t_i})), \quad (6)$$

where  $p(\mathbf{s}_{t_i})$  is the true distribution in which the correct token at the  $i$ -th index has a probability of 1, and all other tokens have a probability of 0. On the other hand,  $p(\hat{\mathbf{s}}_{t_i})$  denotes the predicted probability distribution over all the possible tokens for the  $i$ -th index.  $\mathcal{L}_{\text{VQVAE}}$  aims to jointly train codebook  $Z$  and the framework in end-to-end to bring the selected codebook vector ( $\mathbf{t}$ ) as close as possible to the encoder output

( $\mathbf{r}$ ) as follows:

$$\mathcal{L}_{\text{VQVAE}}(\mathbf{r}, \mathbf{t}; Z) = \|\text{sg}[\mathbf{r}] - \mathbf{t}\|_2^2 + \beta \|\text{sg}[\mathbf{t}] - \mathbf{r}\|_2^2. \quad (7)$$

where  $\text{sg}[\cdot]$  represents the stop-gradient operation, ensuring no gradient is passed through and treating it as a constant that doesn't update. It is applied to  $\mathbf{r}$  to update only the codebook in the first term, while it is applied to  $\mathbf{t}$  to update only the encoder output, thereby bringing these two closer together. Our framework is trained by combining the two aforementioned loss functions as follows:

$$\mathcal{L} = \mathcal{L}_{\text{CE}}(\hat{\mathbf{s}}, \mathbf{s}) + \mathcal{L}_{\text{VQVAE}}(\mathbf{r}, \mathbf{t}; Z), \quad (8)$$

In this training (fine-tuning) process, the trainable parameters include AI-Src-Enc  $S_{\text{E}}(\cdot)$ , AI-Src-Dec  $S_{\text{D}}(\cdot)$ , the embedding  $\mathcal{E}$ , quantizer  $Q_{\text{E}}(\cdot; Z)$ , dequantizer  $Q_{\text{D}}(\cdot; Z)$ , and the codebook  $Z$ .<sup>4</sup>

### D. Optimization Techniques

Our system employs three key optimization techniques to enhance performance:

1) *Noise-Tuning*: We optimize robustness against communication noise by adaptively tuning AI-Src-Enc and AI-Src-Dec to bit (or symbol) errors or channel noise. During fine-tuning, we expose them to simulated channel impairments (e.g., under the 3GPP CDL-A channel) alongside the channel encoder/decoder. This simple fine-tuning, without modifying any PHY module, significantly improves reliability across diverse channel/communication conditions.

2) *Codebook*: To achieve efficient data compression, we utilize a VQ-VAE. It learns a codebook of discrete vectors capturing the essence of the AI-Src-Enc's continuous representations. These codebook vectors, significantly smaller than the originals, are transmitted, reducing bandwidth demands. Interestingly, it also mitigates the impact of bit errors during transmission as its representation learns the bit (or symbol) error pattern - potentially due to it learning the bit (or symbol) error pattern during pre-training. Upon reception, the VQ-VAE decodes these indices back into high-dimensional vectors. This decoding utilizes the shared latent space (*i.e.*,  $Z$ ) learned during pre-training, ensuring effective recovery of the semantic content within the compressed representation. This method balances accuracy, compression, and inference complexity through the adjustable compression rate (and embedding size).

3) *Pre-Training*: To reduce the need for extensive application-specific training data, we leverage pre-trained models like BART-base. This pre-training equips the system with a broad base of generalizable knowledge, enhancing generalization performance.

## IV. EXPERIMENTS

### A. Dataset

1) *Train Dataset*: Theoretically, for training purposes, any text can be utilized. We can set both the input  $\mathbf{s}$  and the output

<sup>4</sup>Wireless PHY components, including code rate  $\rho$ , were not directly incorporated in the NN training. We acknowledge the potential for adaptive adjustment of specific standardized functions to improve robustness, which will be explored in future research.

TABLE I  
TYPE (OR PARAMETER) FOR CHANNEL EN/DECODER

Type (or Parameter)	Value
<b>Channel En/Decoder</b>	
Channel coding	Polar code
Coderate ( $\rho$ )	0.5
# of coded bits	960
# of information bits	480
Mapper	QAM
# of bits per symbol ( $m$ )	4
# of OFDM symbols	14
# of TX×RX antennas	2×2 (1×1) <sup>a</sup>
Direction	Downlink
Carrier frequency	2.6 [GHz]
Subcarrier spacing	15 [kHz]
FFT size (# of subcarrier)	72
Channel (train)	CDL-A, Rayleigh
Channel (test)	CDL- $\{A\sim D\}$ , Rayleigh
<b>AI-Source-En/Decoder</b>	
Pre-trained LM ( $\mathcal{T}$ , $\mathcal{E}$ , $S_{\mathbf{E}}(\cdot)$ , $S_{\mathbf{D}}(\cdot)$ )	BART-base
embedding dimension ( $d_r$ )	768
<b>Codebook</b>	
Dimension of each code ( $d_z$ )	2
Codebook size ( $K$ )	1024

<sup>a</sup>Other than the results in Sec. V-B, which uses a 1x1 TX/RX configuration for a fair comparison, a MIMO (2x2) configuration is considered for the rest of the simulations.

$\hat{s}$  to be identical for each sentence. This approach aims to train the framework to ensure the faithful transmission of sentences without any modifications, aligning with the primary objective of the system. However, following the precedent set by earlier studies in text semantic communication [14], [18], we utilize the European Parliament dataset [29] (called *EuroParl*) which is extracted from the proceedings of the European Parliament.

2) *Test Dataset*: For the evaluation, it is important to evaluate the generalization efficacy of framework on out-of-distribution data, particularly on sentences or even tokens not encountered during training. To evaluate the generalizability, we randomly sample 1K sentences from the image-caption dataset *Flickr* [30]. The token distribution in this dataset differs from our training data, a result of reporting bias. This bias emerges because people tend to report what interests them (e.g., parliamentary discussions) rather than typical and general facts (e.g., describing an image).

### B. Implementation Details

Our entire framework is built using the Keras framework [31]. For modeling various channel structures, we utilize the Sionna library [22]. We employ the Huggingface library [32] to load the pre-trained weights of the BART-base model, which serves as the foundation for initializing the embedding  $\mathcal{E}$ , tokenizer  $\mathcal{T}$ , AI-Src-Enc  $S_{\mathbf{E}}(\cdot)$  and the AI-Src-Dec  $S_{\mathbf{D}}(\cdot)$ , respectively.  $\mathcal{T}$ , pre-trained embedding  $\mathcal{E}$ . In the fine-tuning phase, the Adam optimizer [33] is employed, with the sequence length capped at 64 tokens, and the training has been done for 3 epochs, using training loss in (8).<sup>5</sup> We conduct

<sup>5</sup>The pre-trained model may or may not use the same training loss function during its pre-training phase.

a grid search to find the best learning rate from the set [3e-4, 1e-4, 5e-5, 2e-5, 1e-5] and the batch size from [2, 4, 8], using the development dataset to determine the optimal parameters. All of our experiments are run on an RTX 2080-Ti using 32-bit floating-point precision.

Detailed system parameters are summarized in Table I.

### C. Evaluation

In traditional communication systems, the performance of information transmission is evaluated by measuring the accuracy of transmitting individual bits (0s and 1s), as reflected by the Bit Error Rate (BER), or by assessing the fidelity of conveying symbols, which are collections of bits, denoted by the Symbol Error Rate (SER).

In contrast, on-device AI communication prioritizes the transmission of meaningful content,<sup>6</sup> focusing on the efficient utilization of bandwidth for more effective data transfer. To accurately evaluate the performance of on-device AI communication systems, we utilize three specific metrics: lexical similarity (e.g., BLEU [34]), semantic similarity (e.g., SBERT [35]) between  $s$  and  $\hat{s}$ , and the compression rate of  $\rho$ .

1) *BLEU*: The BLEU score [34], initially developed for evaluating machine translation, quantifies the correspondence between the  $n$ -grams of a generated sentence  $\hat{s}$  and those in a reference sentence  $s$  [14], [18]. Here,  $n$ -grams mean a collection of  $n$  successive words in a sentence. BLEU score involves two key factors: (1)  $n$ -gram-based precision of the generated sentence  $\hat{s}$  and the reference sentence  $s$  as follow:

$$p_n = \frac{\sum_{n\text{-gram} \in \hat{s}} \text{Count}_{\text{clip}}(n\text{-gram})}{\sum_{n\text{-gram} \in \hat{s}} \text{Count}(n\text{-gram})}. \quad (9)$$

Here,  $\text{Count}_{\text{clip}}(n\text{-gram})$  represents the number of  $n$ -grams from the generated sentence  $\hat{s}$  that are found in a reference sentence  $s$  while  $\text{Count}(n\text{-gram})$  represents the total number of  $n$ -grams in the generated sentence  $\hat{s}$ .

However, simply counting identical  $n$ -grams can lead to an overestimation. Consider the case where the reference sentence is “I am a boy” and the generated sentence is “a a a a”. In this scenario, the unigram precision  $p_1$  would erroneously be computed as 1, as each occurrence of the unigram “a” is found in the reference sentence.

To address this issue,  $\text{Count}_{\text{clip}}(n\text{-gram})$  is employed to cap the count at the highest frequency observed in the reference sentence  $s$ , thereby preventing overcounting. In the given example,  $\text{Count}_{\text{clip}}(\text{“a”})$  is adjusted to 1 instead of 4, reflecting the maximum occurrence of the unigram “a” in the reference sentence  $s$ ; (2) A brevity penalty is used to mitigate the influence of sentence length, preventing the overfitting to the sentence length. This penalty comes into play when the length of generated sentence  $\hat{s}$  is shorter than the sentence length of the reference sentence  $s$  as follow:

$$\text{BP} = \begin{cases} 1, & \text{if } |\hat{s}| > |s| \\ e^{(1-|s|/|\hat{s}|)}, & \text{if } |\hat{s}| \leq |s| \end{cases} \quad (10)$$

<sup>6</sup>This objective aligns with those of other research in the field of semantic communication.

TABLE II

PERFORMANCE COMPARISON STUDY. EACH BASELINE IS TRAINED WITH EUROPARL TRAINING SET AND EVALUATED USING BOTH THE EUROPARL AND FLICKR TEST SETS (CHANNEL = RAYLEIGH, EbN0 = 7 [dB]). THE EUROPARL TEST SET ASSESSES THE IN-DISTRIBUTION PERFORMANCE, WHEREAS THE FLICKR TEST SET EVALUATES THE OUT-OF-DISTRIBUTION PERFORMANCE

Framework	PHY Comm.	Transmission	Pre-Training	Noise-Tuning	EuroParl (In-dist.)			Flickr (Out-dist.)		
					BLEU-3	BLEU-4	SBERT	BLEU-3	BLEU-4	SBERT
DeepSC [14]	✗ (symbol-level)	✗	✗	✗	0.6329	0.5802	0.7669	0.1878	0.1111	0.3737
Seq2Seq-SC [25]	✓ (bit-level)	Tanh	✓	✗	0.6428	0.5732	0.8142	0.5522	0.4795	0.7910
Ours (1 × 1)	✓ (bit-level)	VQ-VAE	✓	✓	<b>0.9959</b>	<b>0.9946</b>	<b>0.9959</b>	<b>0.9714</b>	<b>0.9682</b>	<b>0.9998</b>

The BLEU score for  $n$ -grams, represented as BLEU- $n$ , is calculated as the product of this brevity penalty and the exponential of the precision score for  $n$ -grams  $\exp(\log p_n)$ . The overall BLEU score is the brevity penalty multiplied by the exponential of the weighted sum of the log precision scores for different  $n$ -grams, represented as  $\exp\left(\sum_{n=1}^N w_n \log p_n\right)$ , where  $w_n$  denotes the weight of the  $n$ -gram. Here, BLEU- $n$  refers to the BLEU score considering only  $n$ -grams, and the general BLEU score is a weighted average of BLEU-1, BLEU-2, BLEU-3, and BLEU-4. Higher-order  $n$ -grams (*i.e.*, longer  $n$ -grams) are indicative of the fluency and grammatical accuracy of the generated sentence  $\hat{s}$ .

2) *SBERT*: Despite a low lexical overlap between the sentences  $s$  and  $\hat{s}$ , their semantic content may be closely aligned. For example, the words “child” and “children” bear a close semantic relationship, yet a BLEU score based on lexical matching would not recognize this and would rate the similarity as zero. To compute such semantic similarity, sentences can be represented into vector embeddings through an embedding model denoted as  $M$ . The degree of semantic similarity can then be quantified by computing the cosine similarity between these vector representations as follow:

$$\text{match}(\hat{s}, s) = \frac{M(s) \cdot M(\hat{s})^T}{\|M(s)\| \|M(\hat{s})\|}. \quad (11)$$

Existing research in semantic communication often employs BERT [36] as the embedding model  $M$  to encode sentences. To derive sentence embedding, they typically use methods like average pooling or extract the embedding from the first token (*i.e.*, CLS pooling). These embeddings are then used to calculate cosine similarity [14], [18]. Nonetheless, such pre-trained models, when not fine-tuned for the specific task of semantic textual similarity, may not effectively capture the true semantic nuances of sentences. This is attributed to the anisotropic nature of the embedding space, which can lead to embeddings that do not properly represent the semantic variance across sentences [37]. Here, we use SBERT [35], which is fine-tuned on semantic textual similarity tasks, to encode the sentence embedding.

3) *Compression Efficiency*: General (lossy) source coding in conventional systems emphasizes compact representation of information to enhance transmission/spectral efficiency, often at the expense of minor data loss. This contrasts with on-device AI communication, where the exchange of efficient representations prioritizes semantic fidelity over mere data volume reduction. The aim here is to transmit the essence and context of messages with minimal semantic distortion.

The compression rate in on-device AI communication, therefore, assesses how effectively AI-Src-Enc output is compressed. It is quantified by the ratio of the sizes of AI-Src-Enc output ( $\mathbf{r}$ ) and channel encoder input ( $\mathbf{t}$ ), described as:

$$\text{compress}(\mathbf{r}, \mathbf{t}) = \frac{\text{size}(\mathbf{t})}{\text{size}(\mathbf{r})}, \quad (12)$$

where  $\text{size}(\mathbf{x})$  denotes the size of vector  $\mathbf{x}$ .

To evaluate how compression impacts the conveyance of a message’s meaning, we define compression efficiency ( $\mathcal{E}_{\text{compress}}$ ) as a metric that examines the balance between the compression rate and semantic similarity:

$$\mathcal{E}_{\text{compress}} \triangleq \frac{\text{match}(\hat{s}, s; \rho)}{\log_2(\text{compress}(\mathbf{r}, \mathbf{t}))}, \quad (13)$$

where  $\text{match}(\cdot, \cdot; \rho)$  assesses the SBERT for a given  $\rho$  compress rate. A large value of  $\mathcal{E}_{\text{compress}}$  indicates high compression efficiency, meaning that the model effectively balances compression and semantic fidelity. Conversely, a small value signifies lower compression efficiency.

## V. EXPERIMENTAL RESULTS

### A. Setup/Baseline

We primarily evaluate our proposed on-device AI/LLM communication frameworks (“Ours”) from four distinct viewpoints:

1) *PHY Communication*: PHY communication conduct bit-level transmission by leveraging specific PHY functions, *e.g.*, Polar coding, QAM mapper, OFDM. On the other hand, existing semantic communication works, exemplified by systems like DeepSC [14]), conduct symbol-level transmission, not considering the practical PHY communication functions.

2) *Pre-Training Utilization*: (1) w/o pre-training utilizes a transformer architecture for the AI-Src-Enc and AI-Src-Dec, starting with random initialization; while (2) w/ pre-training begins with pre-trained model weights to initialize the AI-Src-Enc and AI-Src-Dec.

3) *Noise-Tuning Utilization*: (1) w/o noise-tuning fine-tunes the framework in the absence of any simulated channel noise; while (2) w/ noise-tuning incorporates simulated channel noise, inducing bit errors, into the fine-tuning process. In particular, the systems is noise-tuned (fine-tuned) under the CDL family of channel models and bit-level PHY communication setup (see details in Table I) and energy per-bit to noise ratio ( $\text{EbN0}$ ) = 5 ~ 15 [dB].



TABLE III  
RECEIVED OUTPUT EXAMPLES OF BASELINES AND OURS (CHANNEL = RAYLEIGH, EbN0 = 7 [dB])

Transmit $s$	A man pushes a bicycle along the beach as the sun sets behind him.
Baseline	Receive $\hat{s}$
DeepSC [14]	a man between a pity along the bureau as the danish sets behind him.
Seq2Seq-SC [25]	A man pushes a bicycle along the beach as sun sets behind him.
Ours (1 × 1)	A man pushes a bicycle along the beach as the sun sets behind him.

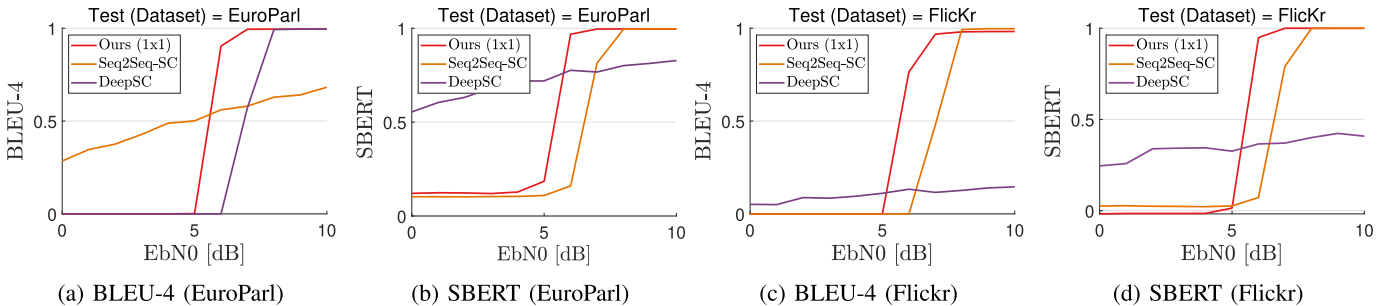


Fig. 4. Performance comparison with existing frameworks under different EbN0 [dB] (Channel = Rayleigh).

4) *Codebook Utilization*: In PHY communication, it is necessary to transform the vector output of the AI-Src-Enc into a discrete bit-level representation. One straightforward method is to directly map each floating-point number in the vector to its bit-level equivalent (e.g., 768-dimensional vector results in  $768 \times 32$  bits). However, this method is highly susceptible to channel noise, where even a single bit error can significantly distort the vector representation (e.g., *floating-point arithmetic error*). To mitigate this issue, we explore and compare various techniques that are more robust to noise. (1) *Direct* encodes the floating-point numbers of  $r$  straight into bits for transmission; (2) *Tanh* also transforms the floating-point numbers of  $r$  into bits for sending but applies a hyperbolic tangent activation function to adjust for the cases when bit errors from channel noise push the received representation  $\hat{r}$  beyond certain range (e.g.,  $[-1, 1]$ ). By doing so, it realigns the out-of-bound values back into the  $[-1, 1]$  range, and has been used in Seq2Seq-SC [25]; (3) *VQ-VAE* employs vector quantization along with a codebook  $Z$  to encode vector representation  $r$  into a set of discrete codebook indices  $t$  for transmission.

### B. Comparative Study

We begin by evaluating the overall efficacy of our proposed framework against earlier works in semantic communication. Each baseline is trained using the EuroParl training set and evaluated on both the EuroParl and Flickr test sets. There is a significant overlap in vocabulary between the EuroParl training and test sets, while the overlap is minimal between the EuroParl training set and the Flickr test set. This setup allows for an assessment of the framework's generalizability. For a fair comparison, all baseline models, including the baseline DeepSC which focuses on symbol-level transmission, are evaluated under identical conditions of  $1 \times 1$  SISO and Rayleigh channel throughout the comparative study.

Table II illustrates the performance comparison under a Rayleigh Channel with  $EbN0 = 7$  [dB]. Meanwhile, Table III offers qualitative examples demonstrating the transmission

accuracy of each methodology, and Fig. 4 highlights the robustness of each approach.

Tables II and III demonstrate that our framework significantly outperforms both DeepSC and Seq2Seq-SC. A key observation from Table II is the substantial decline in DeepSC's performance on the Flickr dataset, whereas the performance reduction for Seq2Seq-SC on Flickr is less severe in comparison. This difference can be mainly due to the use of *pre-training* which exhibits generalization capability for OOV - that will be elaborated on a later section.

Fig. 4 also demonstrates the robustness of our framework compared to Seq2Seq-SC across various EbN0 [dB]. An interesting point in Fig. 4 is the gradual increase in DeepSC's trend compared to the sharp rise in Seq2Seq-SC and our framework. This primarily stems from whether the transmission takes place through bit-level transmission via practical PHY communication. Transmitting through the PHY process information in bit-level, which requires converting the floating-point vector into bits, where bit errors caused by noise can lead to substantial changes in the information (called *floating-point arithmetic error*). Therefore, in conditions of substantial noise, such as when  $EbN0 < 5$  [dB], transmitting vector representations through the PHY becomes challenging. Conversely, DeepSC transmits representations based on symbol-level, where noise is directly added into the continuous vector representation, while not considering the practical issues in bit-level operation. This results in such a gradual increase in its performance trend as it does not experience any floating-point error.

### C. Ablation Study

In this part of the discussion, we explore the impact of each component within our optimized E2E on-device AI communication framework and examine their respective roles. Table IV shows how performance varies with and without the use of *pre-training* and *noise-tuning*. Additionally, Table VI shows the performance differences when using

TABLE IV  
PERFORMANCE VARIANCE w/ AND w/o PRE-TRAINING AND NOISE-TUNING (CHANNEL = CDL-A, EbN0 = 4 [dB])

Combination	Codebook	Pre-Training	Noise-Tuning	EuroParl (In-dist.)			Flickr (Out-dist.)		
				BLEU-3	BLEU-4	SBERT	BLEU-3	BLEU-4	SBERT
Ours	VQ-VAE	✓	✓	0.9918	0.9888	0.9944	0.9296	0.9178	0.9959
w/o pre-training (3 epochs)	VQ-VAE	✗	✓	0.0690	0.0393	0.3571	0.0002	0.0000	0.0080
w/o pre-training (10 epochs)	VQ-VAE	✗	✓	0.6136	0.5376	0.7277	0.0573	0.0246	0.0933
w/o noise-tuning	VQ-VAE	✓	✗	0.9515	0.9325	0.9804	0.9147	0.8907	0.9787

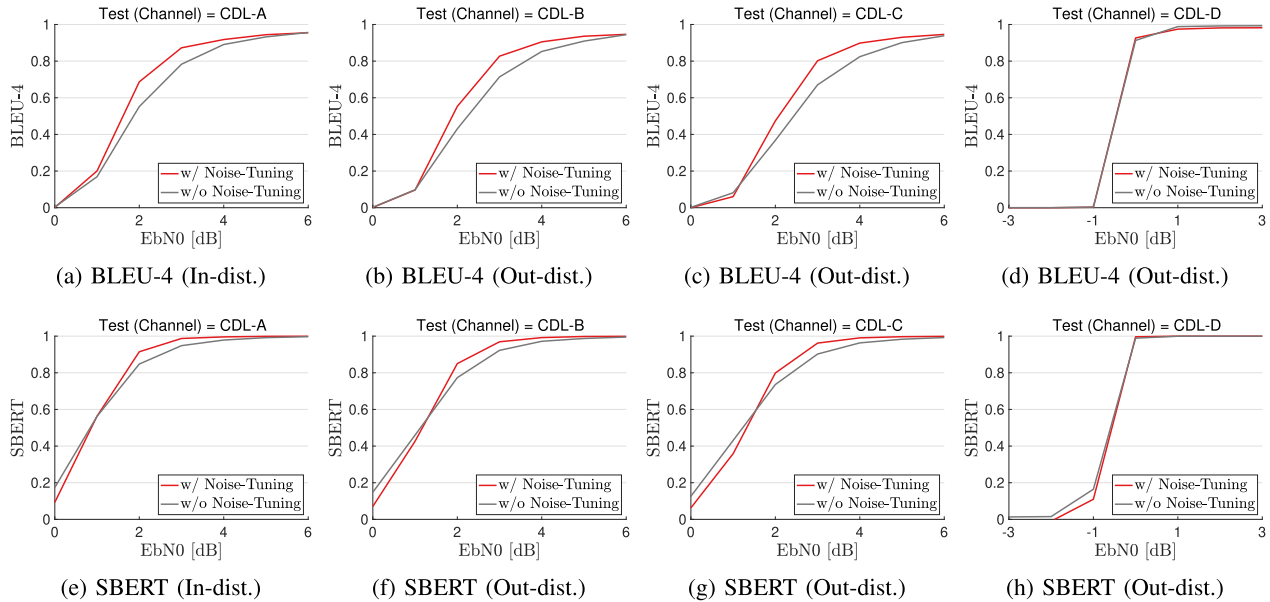


Fig. 5. Impact of noise-tuning (w/o Noise-Tuning vs. w/ Noise-Tuning). The CDL-A channel assesses the in-distribution performance, whereas the CDL-B, -C, and -D evaluate the out-of-distribution performance.

different transmission approaches including Naive, Tanh and VQ-VAE, showing the impact of the use of Codebook.

1) *Impact of Pre-Training*: As shown in Table IV, w/o pre-training leads to a considerable performance drop. After increasing the number of training epochs, the framework shows improvement on in-distribution data that shares a substantial vocabulary and distribution overlap with the training data. However, its performance on other datasets with minimal overlap with the training data remains close to zero. This indicates that pre-training plays a crucial role in the framework's ability to generalize - which allows this framework to be usable across a wide range of applications.

2) *Impact of Noise-Tuning*: As shown in Table IV, w/o noise-tuning leads relatively small performance drop. However, it provides an intriguing insight that LM and VQ-VAE (i.e., AI-Scr-Enc and AI-Scr-Dec) are somewhat able to learn how to handle bit (or symbol) errors during their communication with another end.

To assess the impact of noise-tuning in more detail, we present Table V and Fig. 5. These results compare the performance of AI-Scr-Enc and AI-Scr-Dec trained with and without noise-tuning (specifically, the CDL-A channel at EbN0 of 5 ~ 15 [dB]) in four different communication conditions by testing those in CDL-A, B, C, and D channels. The performance difference showcases the impact of incorporating

TABLE V  
IMPACT OF NOISE-TUNING (w/o NOISE-TUNING VS. w/ NOISE-TUNING)

EbNo [dB]	Noise-Tuning (CDL-A)	Test (Channel)	BER	BLEU	SBERT
2 [dB]	✗	CDL-A	$2.124 \times 10^{-1}$	0.5801	0.8477
	✓	CDL-A		0.7278	0.9147
	✗	CDL-B	$2.332 \times 10^{-1}$	0.4588	0.7734
	✓	CDL-B		0.6125	0.8500
3 [dB]	✗	CDL-C	$2.419 \times 10^{-1}$	0.4072	0.7359
	✓	CDL-C		0.5394	0.7988
	✗	CDL-A	$1.641 \times 10^{-1}$	0.7931	0.9481
	✓	CDL-A		0.8893	0.9872
	✗	CDL-B	$1.844 \times 10^{-1}$	0.7312	0.9224
	✓	CDL-B		0.8441	0.9691
✗	CDL-C	$1.934 \times 10^{-1}$	0.6858	0.9023	
✓	CDL-C		0.8260	0.9619	

realistic communication errors into the fine-tuning process, indicating that AI-Scr-Enc and AI-Scr-Dec can learn and partially mitigate bit error patterns induced by noisy channel encountered during link-level PHY communication.<sup>7</sup> However,

<sup>7</sup>We conjecture that the AI-Scr-Enc and AI-Scr-Dec (e.g., LM and VQ-VAE) learn bit error pattern of a pre-determined constellation in mapper/demapper. When symbol error occurs, neighbor symbols likely are detected, not far-distant symbols. During noise-tuning, such intelligent en/decoder adaptively modifies its  $\mathcal{E}$  and  $\mathcal{Z}$  to be more robust to bit errors.

TABLE VI  
PERFORMANCE VARIANCE WITH DIFFERENT TRANSMISSION APPROACHES (CHANNEL = CDL-A, EbN0 = 4 [dB])

Case	Codebook	Pre-train	Noise-tuning	EuroParl (In-dist.)			Flickr (Out-dist.)		
				BLEU-3	BLEU-4	SBERT	BLEU-3	BLEU-4	SBERT
Naive Transmission	✗ (Naive)	✓	✓	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Tanh Transmission (Seq2Seq-SC)	✗ (Tanh)	✓	✓	0.3660	0.2759	0.6691	0.1636	0.0983	0.5094
VQ-VAE Transmission (Ours)	✓ (VQ-VAE)	✓	✓	0.9918	0.9888	0.9944	0.9296	0.9178	0.9959




Fig. 6. Impact of codebook. Figs. 6(a) and 6(b) compares w/o Codebook and w/ Codebook. Figs. 6(c)-6(e) show the tradeoff between compression and accuracy.




Fig. 7. Compression efficiency over compression rate.

how the AI-Src-Enc and AI-Src-Dec learn and overcome communication errors is still an open question that we defer to our future work. Overall, the AI-Src-Enc/AI-Src-Dec archives EbN0 gain up to  $\sim 1$  [dB] by utilizing the noise-tuning.

An important challenge for on-device AI communication is their robustness to test-time distributional shifts that naturally occur when the test environment no longer matches the conditions in a training environment. This is present in wireless communication systems, where the propagation conditions may change whenever the user is moving. Thus, an important question is whether the on-device AI communication can retain its performance, both from the same communication condition in training (in-distribution) and other communication conditions (out-of-distribution). To this end, we also test its performance in CDL-B, -C, and -D channels, whose condition is not seen in training, in Table V and Figs. 5.<sup>8</sup> CDL-B and CDL-C results indicate that noise-tuning enhances the robustness even in channels not seen in training, while the CDL-D (Figs. 5(a) and 5(b)) case highlights that there is still room for improvement for the noise-tuning to be more channel-agnostic.

<sup>8</sup>CDL-A channels encompass both line-of-sight (LoS) and non-LoS (NLoS) conditions, offering a more diverse representation of wireless channel propagation. In contrast, CDL-B and CDL-C channels are predominantly NLoS. CDL-D channels, on the other hand, are characterized by LoS.

3) *Impact of Codebook*: Table VI shows that the VQ-VAE transmission significantly outperforms other methods (e.g., Seq2Seq-SC), achieving near-flawless results under the CDL-A channel at an EbN0 = 4 [dB]. This demonstrates the effectiveness of using a discrete codebook for vector representation quantization in enhancing both accuracy and transmission robustness.

Figs. 6(a) and 6(b) further illustrate the impact of the codebook (VQ-VAE) over different EbN0 [dB]. In the scenario where a codebook is not used, the shape for transmitting information is an  $n \times d_r$ , where each element is a floating point vector. Thus, the total size of the transmitted information  $x$  becomes  $n \times d_r \times 32$  bits, considering that each floating-point vector is 32 bits. Conversely, when a codebook is used when  $d_z = 2$ , the size of transmission is reduced to  $n \times (d_r/2)$ . In this case, each point in the transmission is an integer index. The total size of the transmitted information  $x$  becomes  $n \times (d_r/2) \times 32$ , with the 32 bits now representing the size of each integer. It is worth noting that using integer datatype is notably robust against rounding errors and precision loss; additionally, it allows for using smaller data types, such as 16-bit integers. Remarkably, even with a 50% compression of the output of AI-Src-Enc, with-codebook consistently outperforms without-codebook, achieving EbN0 gains up to  $\sim 6$  [dB]. This suggests the potential of the encoded message from the AI-Src-Enc to be effectively compressed without losing its accuracy.

4) *Tradeoff: Compression ( $d_z$  and  $K$ ) vs. Accuracy*: The dimension of each code vector ( $d_z$ ) and codebook size ( $K$ ) in AI-Src-Enc/Dec impact the performance of the on-device AI communication systems. Figs. 6(c) and 6(d) show how  $d_z$  affects performance, examining increments through  $\{2, 4, 8, 16, 32, 64\}$ . Increasing  $d_z$  enhances compression but may reduce transmission accuracy, as compressing a data vector of size  $d_z$  into a single index increases vulnerability to transmission errors—a single index error can impact the entire  $d_z$ .

Fig. 6(e) shows the relationship between  $K$  and the performance. Smaller  $K$  reduces computational demands, yet it might cause an information loss during quantization, which can adversely affect the fidelity of the reconstructed representation. On the other hand, increasing  $K$  provides a more fine-grained resolution and greater accuracy. However, this comes with increased computational cost due to the cost of training and inference overhead, and it may even hinder the convergence of training.

Additionally, Fig. 7 the relationship between compression efficiency ( $\mathcal{E}_{\text{compress}}$ ) and code vector dimension, identifying a Pareto optimal point between the two. Therefore, finding the right balance and designing an appropriate code vector dimension and codebook size is crucial for achieving a desirable balance between accuracy and efficiency. Thus, optimizing the compression rate (e.g.,  $d_z$  and  $K$ ) for AI-Src-Enc/Dec is crucial for striking an optimal balance between accuracy and computational efficiency.

## VI. CONCLUSION

In this study, we integrated the BART pre-trained LM with the NVIDIA SIONNA link-level simulator to simulate on-device AI communication within a 5G-NR framework. The performance of on-device communication was assessed in bit-level transmission across various PHY communication setups, including 3GPP CDL-family channels, MIMO, OFDM, and Polar code.

Our analysis yielded several key insights from three optimization techniques: (1) Integration of pre-trained language models significantly enhances the generalization capabilities of the on-device AI communication system; (2) Noise-tuning results in an approximate EbN0 gain of around 1 [dB], effectively improving performance in both familiar (CDL-A channel) and new (CDL-B~D channels) communication conditions; (3) Codebook utilization leads to an EbN0 gain of up to approximately 6 [dB], accompanied by a 50% reduction in AI-Src-Enc output size. Even with a compression rate of 8 (utilizing only 12.5% of the compressed information), the system maintains reasonable performance levels.

The implementation strategy involves two stages:

- *Training Stage (Offline)*: The AI-Src-Enc/Dec embedding, along with the codebook, are trained by minimizing the loss function described in Sec. III-C. This computationally intensive stage occurs offline, leveraging powerful resources. The training (fine-tuning) process can be optimized using efficient fine-tuning approaches, such as QLoRA [38].
- *Inference Stage (On-Device)*: The on-device system employs the pre-trained AI-Src-Enc/Dec and codebooks, requiring only inference-level computations. The inference process can be optimized using on-device ML optimization techniques, such as pruning, 8-bit quantization, or knowledge distillation.

*Limitations*: While our proposed framework showcases the performance of on-device AI communication under a bit-level transmission with a link-level PHY communication setup, there are notable limitations that should be addressed to achieve full compliance with 5G-NR communication systems.

Specifically, beyond the PHY, the medium access control (MAC) or network layers typically employs automatic repeat request (ARQ) or hybrid ARQ (HARQ) protocols for packet error correction. The absence of consideration for these higher-layer protocols in our study limits the assessment of the practical performance of our framework. Incorporating ARQ or HARQ mechanisms into our simulation setup would provide a more comprehensive evaluation of the on-device AI communication framework's efficacy in real-world communication scenarios.

*Open Questions*: Experimental results in Sec. IV-C highlight that on-device AI can improve communication performance through its predictive and contextual learning abilities. While integrating channel encoding/decoding with AI-Src-Enc/Dec, such as through JSCC or AI-native PHY/MAC, could unlock the full potential, this faces challenges like compatibility with 4G-LTE and 5G-NR, regulatory issues, and cost implications.

Therefore, employing AI as a complementary tool while keeping or carefully optimizing existing PHY or higher-layer functions presents a more feasible approach. Future research could focus on: (1) finding the best noise-tuning conditions without altering PHY functions; (2) exploring how language models can effectively counteract noisy conditions; (3) improving specific PHY functions to boost AI communication; (4) determining if AI/LM can also learn packet error patterns to ease the HARQ (or ARQ) process.

## ACKNOWLEDGMENT

This work was conducted while Ju-Hyung Lee was affiliated with the University of Southern California (USC). The authors express their deep gratitude to Prof. Andreas F. Molisch for his invaluable feedback and multifaceted support throughout this work. They also thank Thomas Choi for his insightful discussions during the initial stages of their research.

## REFERENCES

- [1] S. Dhar, J. Guo, J. J. Liu, S. Tripathi, U. Kurup, and M. Shah, "A survey of on-device machine learning: An algorithms and learning theory perspective," *ACM Trans. Internet Things*, vol. 2, no. 3, pp. 1–49, Jul. 2021.
- [2] R. Bommasani et al., "On the opportunities and risks of foundation models," 2021, *arXiv:2108.07258*.
- [3] S. Geng, S. Liu, Z. Fu, Y. Ge, and Y. Zhang, "Recommendation as language processing (RLP): A unified pretrain, personalized prompt & predict paradigm (P5)," in *Proc. 16th ACM Conf. Recommender Syst.*, New York, NY, USA, Sep. 2022, pp. 299–315.
- [4] C. Cui et al., "A survey on multimodal large language models for autonomous driving," 2023, *arXiv:2311.12320*.
- [5] G. Delétang et al., "Language modeling is compression," 2023, *arXiv:2309.10668*.
- [6] X. Lin, "An overview of 5G advanced evolution in 3GPP release 18," *IEEE Commun. Standards Mag.*, vol. 6, no. 3, pp. 77–83, Sep. 2022.
- [7] Z. Qin, X. Tao, J. Lu, W. Tong, and G. Ye Li, "Semantic communications: Principles and challenges," 2021, *arXiv:2201.01389*.
- [8] E. Boursoulatze, D. B. Kurka, and D. Gündüz, "Deep joint source-channel coding for wireless image transmission," *IEEE Trans. Cogn. Commun. Netw.*, vol. 5, no. 3, pp. 567–579, Sep. 2019.
- [9] J. Shao, Y. Mao, and J. Zhang, "Learning task-oriented communication for edge inference: An information bottleneck approach," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 1, pp. 197–211, Jan. 2022.
- [10] D. Huang, F. Gao, X. Tao, Q. Du, and J. Lu, "Toward semantic communications: Deep learning-based image semantic coding," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 55–71, Jan. 2023.
- [11] T.-Y. Tung and D. Gündüz, "DeepWiVe: Deep-learning-aided wireless video transmission," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 9, pp. 2570–2583, Sep. 2022.

- [12] S. Wang et al., "Wireless deep video semantic transmission," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 214–229, Jan. 2023.
- [13] P. Jiang, C.-K. Wen, S. Jin, and G. Y. Li, "Wireless semantic communications for video conferencing," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 230–244, Jan. 2023.
- [14] H. Xie, Z. Qin, G. Y. Li, and B.-H. Juang, "Deep learning enabled semantic communication systems," *IEEE Trans. Signal Process.*, vol. 69, pp. 2663–2675, 2021.
- [15] T. Han, Q. Yang, Z. Shi, S. He, and Z. Zhang, "Semantic-preserved communication system for highly efficient speech transmission," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 245–259, Jan. 2023.
- [16] Z. Weng, Z. Qin, X. Tao, C. Pan, G. Liu, and G. Y. Li, "Deep learning enabled semantic communications with speech recognition and synthesis," *IEEE Trans. Wireless Commun.*, vol. 22, no. 9, pp. 6227–6240, Sep. 2023.
- [17] A. Vaswani, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 30, 2017, pp. 6000–6010.
- [18] H. Hu, X. Zhu, F. Zhou, W. Wu, R. Q. Hu, and H. Zhu, "One-to-many semantic communication systems: Design, implementation, performance evaluation," *IEEE Commun. Lett.*, vol. 26, no. 12, pp. 2959–2963, Dec. 2022.
- [19] T. Brown et al., "Language models are few-shot learners," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 33, 2020, pp. 1877–1901.
- [20] A. Srivastava et al., "Beyond the imitation game: Quantifying and extrapolating the capabilities of language models," 2022, [arXiv:2206.04615](https://arxiv.org/abs/2206.04615).
- [21] S. Thrun, "Lifelong learning algorithms," in *Learning to Learn*. Boston, MA, USA: Springer, 1998, pp. 181–209.
- [22] J. Hoydis et al., "Sionna: An open-source library for next-generation physical layer research," 2022, [arXiv:2203.11854](https://arxiv.org/abs/2203.11854).
- [23] A. Van Den Oord et al., "Neural discrete representation learning," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 30, 2017, pp. 6309–6318.
- [24] M. Lewis, "BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, Jul. 2020, pp. 7871–7880.
- [25] J.-H. Lee, D.-H. Lee, E. Sheen, T. Choi, and J. Pujara, "Seq2Seq-SC: End-to-end semantic communication systems with pre-trained language model," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 2023, pp. 260–264.
- [26] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- [27] NR; Multiplexing and Channel Coding, document 3GPP TS 38.212, Version 17.3.0, Sep. 2022.
- [28] Study on Channel Model for Frequencies From 0.5 to 100 GHz, document 3GPP TS 38.901, Version 17.0.0, Mar. 2022.
- [29] P. Koehn, "Europarl: A parallel corpus for statistical machine translation," in *Proc. Mach. Transl. Summit X, Papers*, Phuket, Thailand, Sep. 2005, pp. 79–86.
- [30] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier, "From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions," *Trans. Assoc. Comput. Linguistics*, vol. 2, pp. 67–78, Dec. 2014.
- [31] F. Chollet et al. (2015). *Keras*. [Online]. Available: <https://keras.io>
- [32] T. Wolf et al., "Transformers: State-of-the-art natural language processing," in *Proc. Empirical Methods Natural Lang. Process.*, 2020, pp. 38–45.
- [33] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, San Diego, CA, USA, 2015, pp. 1–15.
- [34] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: A method for automatic evaluation of machine translation," in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, Philadelphia, PA, USA, 2002, pp. 311–318.
- [35] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Hong Kong, 2019, pp. 3982–3992.
- [36] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Assoc. Comput. Linguistics*, Minneapolis, MN, USA, 2019, pp. 4171–4186.
- [37] B. Li, H. Zhou, J. He, M. Wang, Y. Yang, and L. Li, "On the sentence embeddings from pre-trained language models," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Nov. 2020, pp. 9119–9130.
- [38] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, "QLoRA: Efficient finetuning of quantized LLMs," 2023, [arXiv:2305.14314](https://arxiv.org/abs/2305.14314).



**Ju-Hyung Lee** received the Ph.D. degree from Korea University, Seoul, South Korea. He was a Post-Doctoral Researcher of electrical and computer engineering with the University of Southern California (USC), Los Angeles, CA, USA, where he worked under the supervision of Prof. Andreas F. Molisch. Before that, he was a Research Professor with Korea University. He is currently a Principal Researcher with Nokia, Sunnyvale, CA, USA. His research interests include generative AI solutions for wireless communication systems and on-device AI/LLM. He has been recognized with several prestigious awards, including the First-Rank in a Machine Learning (ML) Competition, the Best Paper Awards at the Key IEEE Conferences, and the Grand Prize for Research Excellence from Korea University. For more information, visit <https://juhyung-lee.com/>.



**Dong-Ho Lee** is currently pursuing the Ph.D. degree in computer science specializing in natural language processing (NLP) and ML with the Information Sciences Institute (ISI), University of Southern California (USC), under the supervision of Prof. Jay Pujara. His research experience includes collaborations with Google DeepMind, Snap Inc., and Microsoft Research. He is contributing to the development of advanced AI agent systems with Softly AI, Seoul, South Korea. For further details, please visit <https://danny-lee.info>.



**Joohan Lee** received the master's degree in computer science from USC. He was a Research Assistant with USC, under the supervision of Dr. Ju-Hyung Lee and Prof. Andreas F. Molisch, focusing on generative AI for wireless communication systems. He also interned at Purdue University, conducting research in NLP, specifically software entity recognition, under the supervision of Prof. Tianyi Zhang. His research interests include tackling complex challenges in NLP, computer vision, and computer networking.



**Jay Pujara** is currently a Research Associate Professor of computer science with the University of Southern California (USC) and a Principal Scientist and the Director of the Center on Knowledge Graphs, Information Sciences Institute, USC. He is the author of over 100 peer-reviewed publications. His research interests include artificial intelligence, specifically knowledge graphs and statistical relational learning. He has received five best paper awards and has been featured in AI Magazine. For more information, visit <https://www.jaypujara.org>.