

Reverse Ordering Techniques for Attention-Based Channel Prediction

VALENTINA RIZZELLO ^{id} (Graduate Student Member, IEEE),
BENEDIKT BÖCK ^{id} (Graduate Student Member, IEEE), MICHAEL JOHAM ^{id} (Member, IEEE),
AND WOLFGANG UTSCHICK ^{id} (Fellow, IEEE)

School of Computation, Information and Technology, Technical University of Munich, 80333 München, Germany

CORRESPONDING AUTHOR: VALENTINA RIZZELLO (e-mail: valentina.rizzello@tum.de).

This work was supported by Futurewei Technologies, Inc., USA.

ABSTRACT Channel state information (CSI) is crucial for enhancing the performance of wireless systems by allowing to adjust the transmission strategies based on the current channel conditions. However, obtaining precise CSI is difficult because of the fast-changing channel conditions caused by multi-path fading. An inaccurate CSI hinders the performance of various adaptive wireless systems, highlighting the need for channel prediction techniques to effectively mitigate the drawbacks of outdated CSI. Conventional methods typically depend on assumptions regarding user velocity or require knowledge of the Doppler frequency. In contrast to existing approaches, we aim for a more robust and practical solution by training neural networks without making any assumptions about user velocity, relying solely on noisy channel observations during training. Specifically, we adapt both the sequence-to-sequence with attention (Seq2Seq-attn) and transformer models for channel prediction. Additionally, a new technique called reverse positional encoding is introduced in the transformer model to improve the robustness of the model against varying sequence lengths. Similarly, the encoder outputs of the Seq2Seq-attn model are reversed prior to the application of attention mechanisms. By means of simulations, we show that these proposed techniques enable the models to effectively capture relationships within sequences of channel snapshots without increasing the complexity. Importantly, this capability remains robust across varying sequence lengths, representing a substantial improvement over existing methodologies.

INDEX TERMS Channel prediction, positional encoding, Seq2Seq, transformer.

I. INTRODUCTION

In 5G, and beyond 5G, wireless communication systems, the channel state information (CSI) is essential for the base station (BS) to optimize its transmission strategy to communicate to the receiving mobile terminal (MT) [1]. The CSI, or channel, is a complex-valued matrix whose dimensions correspond to the number of transmit and receive antennas. It describes the link between each transmit and receive antenna pair, that can be affected by factors such as fading, multipath propagation, and interference from other signals. In a typical frequency division duplex (FDD) system, the BS sends a predefined sequence of symbols, called pilots, to the MT, which estimates the CSI and feeds the CSI coefficients back to the BS. Hence, there is an inevitable delay between the instant of when the

MT estimates the CSI and the one in which the BS receives the CSI coefficient. Therefore, since the channels change over time, it is crucial for the BS to predict the channel [2]. The problem of channel prediction is quite straightforward when the channel dynamics are known. In particular, when the Doppler frequency is known, linear predictors such as autoregressive (AR) models or Kalman filters (KFs) can effectively be used for tracking the CSI, see [3], [4], [5], [6]. However, in a typical wireless communication system, the MTs move with unknown channel statistics and different velocities. Therefore, a finite number of AR predictors need to be pre-trained for different Doppler frequencies and the channel parameters must be estimated from the available data. In order for this approach to work well, both, *i*) the Doppler frequency must

be correctly estimated, and *ii*) a potentially large number of linear predictors need to be stored. Additionally, a wrong or a coarse approximation of the Doppler frequency can cause a non-negligible loss. Given these challenges, it is crucial for both academia and industry to prioritize the development of reliable predictive models. Notably, predictive models play a crucial role in reducing latency by anticipating channel variations, which is especially vital for high-speed users facing rapid changes.

In recent years, neural networks (NNs) have become a promising solution in various research fields including wireless communications. In [7], convolutional neural networks (CNNs) are used in combination with AR models for CSI forecasting. In particular, CNNs are used to correctly identify the channel dynamics, and to load the corresponding pre-trained AR predictor to forecast the CSI. The authors of [8] also propose a hybrid approach called Hypernetwork Kalman Filter. There, a single-antenna setup is considered and only Kalman equations are utilized for prediction, whereas a hypernetwork continuously updates the Kalman parameters based on past observations. In [9] it was shown that the machine learning-based predictor requires smaller operational complexity compared to a vector Kalman filter-based predictor. In [10] a spatio-temporal neural network was proposed by considering channel prediction of MIMO-OFDM as a spatio-temporal series problem. The work in [11] instead considers frequency-selective channels and optimizes linear predictors by leveraging a long-short-term channel decomposition, in which the channel is decomposed into long-term space-time signatures and fading amplitudes. Additionally, meta-learning is used for training the model and reducing the required number of pilots. In [12] meta-learning is used to improve the prediction accuracy with only a few fine-tuning samples. In [7], [13], [14], [15] recurrent neural networks (RNNs) are used for CSI prediction. In particular, due to the ability of RNNs to incorporate the typical dynamics of time series data, they represent a valid alternative to AR models for time series forecasting. However, notably RNNs are difficult to train due to vanishing or exploding gradient issues, see [16]. In [17] it was shown that deep learning brings a notable performance gain compared with the conventional predictors built on shallow recurrent neural networks. Among the recent advances, we find the work in [18] where the CSI prediction is incorporated in a reinforcement learning-based setup with goal to maximize the multi-user sum rate over time. In this setting, the so-called Actor Network is responsible for CSI prediction and it is realized via a multi-layer perceptron (MLP). The most recent study in [19] presents a novel approach to predict future channels in parallel using a transformer-based parallel channel prediction scheme. The main objective is to prevent error propagation that often occurs in sequential prediction. To this end, the one-step ahead sequential prediction in the transformer decoder, also called dynamic decoding in the literature, is eliminated completely. Instead, the transformer decoder takes as input a certain number of the past channel realizations, along with a specific number of all-zero vectors

equal to the number of unknown channels, to predict all the future channels in parallel.

In this study, we draw inspiration from the achievements in natural language processing by so-called attention-based models. We adapt both, the transformer and the sequence-to-sequence with attention (Seq2Seq-attn) models, to the channel prediction task. The decision to utilize Seq2Seq-attn and transformer models for predicting dynamic channels is driven by their inherent compatibility with sequential data. These models are specifically designed to handle tasks where the input and output exhibit varying lengths, such as machine translation or time series prediction. Instead of settling for the vanilla architecture [20], [21], we introduce a novel reversed positional encoding (RPE) technique in the transformer model to improve the model's robustness against variable sequence lengths during testing, that may differ from the lengths assumed for training. With the same goal in mind, we also reverse the encoder outputs of the Seq2Seq-attn model before applying attention. Unlike the state of the art, we investigate the challenging setup where only noisy channels are available for training and where the users are moving inside a cell within a wide range of velocities, i.e., between 0 km/h and 120 km/h. We evaluate our models for varying noise levels and sequence lengths, including lengths that differ from those used during training. Additionally, we analyze the complexity of the proposed neural networks with respect to existing benchmarks both in terms of number of parameters and in terms of floating point operations. The contributions of this paper are summarized as follows:

- We adapt both the transformer and the Seq2Seq-attn model to the channel prediction task.
- We introduce novel ordering techniques in those models to make them robust in adapting to CSI sequences of any length,
- We show by means of simulations that the proposed methods overcome the limitations of both classical methods and existing learning based methods without introducing undue complexity.

For the source code, refer to our GitHub repository: <https://github.com/vrizz/reverse-ordering>. The rest of the paper is organized as follows. In Section II, the system model is described; in Section III, the proposed Transformer-RPE model is presented; in Section IV, the proposed Seq2Seq-attn-R model is presented; in Section VI, the dataset used and the training setup are presented, and the simulation results are discussed; in Section VII, we draw our conclusions.

II. SYSTEM MODEL

We consider a BS serving multiple MTs in a typical 5G cell. The BS is equipped with M antennas, whereas the single-antenna users are moving with different velocities and, therefore, experience different fading conditions. In accordance to the 5G standard [22], we assume that the channel remains constant for the duration of a slot which we denote as T_{slot} , and that a frame contains N_{slot} slots. Additionally, we assume that the velocities of the users are constant within

the duration of a frame. From now on, we denote as $\mathbf{h}_i \in \mathbb{C}^M$ the CSI vector corresponding to the i th generic slot of the CSI time series. The sequence of N_{slot} subsequent CSI vectors $\{\mathbf{h}_i\}_{i=1}^{N_{\text{slot}}}$ is assumed to be strongly correlated. The goal of multi-step CSI prediction is to find the best estimator $f: \mathbb{C}^{M \times \ell} \rightarrow \mathbb{C}^{M \times \delta}$ which predicts the CSI vectors $\{\mathbf{h}_i\}_{i=\ell+1}^{\ell+\delta}$ based on preceding ℓ observations $\{\mathbf{h}_i\}_{i=1}^{\ell}$ with $\ell + \delta \leq N_{\text{slot}}$. Note that, we assume that the channels $\{\mathbf{h}_i\}_{i=1}^{\ell}$ are not perfectly known, and that only the corresponding noisy observations $\{\check{\mathbf{h}}_i\}_{i=1}^{\ell}$ are available, i.e.,

$$\check{\mathbf{h}}_i \leftarrow \mathbf{h}_i + \mathbf{n}_i \quad i = 1, \dots, \ell, \quad (1)$$

where \mathbf{n}_i denotes the complex-valued noise vector with independent elements distributed as $\mathcal{N}_{\mathbb{C}}(0, \sigma_n^2)$, and such that $\mathbb{E}[\mathbf{n}_i \mathbf{n}_j^H] = \mathbf{0}$ for all $i \neq j$. Our assumption of continuous synthetic channel data availability provides a controlled training and evaluation environment. Though different from real-world situations, it enables systematic assessment of our model's predictive capabilities. In practice, we acknowledge the need of digital twins [23] or actual measurements to capture real-world dynamics. In addition, throughout this study, we consider real-valued neural networks. Therefore, we transform the complex vector \mathbf{h}_i into a real vector where the real and imaginary parts of the original vector are concatenated as

$$\mathbb{R}^{2M} \ni \check{\mathbf{h}}_i = \text{concat}(\Re\{\check{\mathbf{h}}_i\}, \Im\{\check{\mathbf{h}}_i\}). \quad (2)$$

III. TRANSFORMER-RPE MODEL

In this section, we describe the proposed Transformer with RPE model for CSI prediction, which has the Transformer model [20] as baseline. The Transformer-RPE model is illustrated in Fig. 1. In the following, we provide a brief description of the Transformer model. However, the reader can refer to [20] for a more detailed explanation. A Transformer consists of an *encoder* and a *decoder* NN. The encoder aims to extract the important information from its input sequence, which can help the decoder to predict the next slots one by one in a subsequent step. This input sequence is represented by the ℓ known noisy channels $\{\check{\mathbf{h}}_i\}_{i=1}^{\ell}$. The encoder has L_{enc} layers and each layer contains two consecutive residual networks. The first residual network has the ‘‘multi-head attention’’ as the main layer, whereas the second residual network contains an MLP as the main module. Moreover, the layer normalization (LN) [24] precedes each of these modules. The MLP contains two fully-connected layers and a GeLU [25] activation function after the first layer. A concise pseudo-code of the multi-head attention layer can be found in Algorithm 1. Multi-head attention in transformer models processes input sequences by using multiple sets of attention weights, or ‘‘heads,’’ instead of a single set. Each head attends to different parts of the input sequence, allowing the model to adapt to various data patterns. The outputs from all heads are combined, providing a comprehensive representation of the input. This approach allows the model to focus on diverse input aspects and learn more expressive features. In the transformer encoder, we compute a multi-head self-attention,

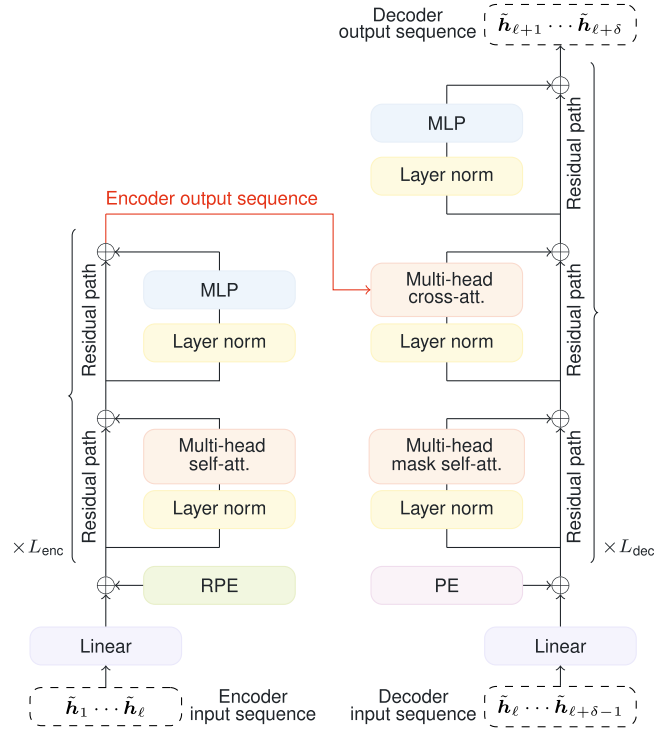


FIGURE 1. Transformer-RPE for CSI prediction.

Algorithm 1: Multi-Head (Masked) Self- or Cross- Attention [20].

Input: $\mathbf{X} \in \mathbb{R}^{d_x \times l_x}$, $\mathbf{Z} \in \mathbb{R}^{d_z \times l_z}$, $\text{Mask} \in \{0, 1\}^{l_z \times l_x}$, primary sequence, context sequence, and an optional mask
Output: $\mathbf{Y} \in \mathbb{R}^{d_{\text{out}} \times l_x}$, updated representation of \mathbf{X}
Hyperparameters: H , number of attention-heads
Learnable parameters: $\mathbf{W}_q \in \mathbb{R}^{H d_{\text{attn}} \times d_x}$,
 $\mathbf{W}_k \in \mathbb{R}^{H d_{\text{attn}} \times d_z}$, $\mathbf{W}_v \in \mathbb{R}^{H d_{\text{mid}} \times d_z}$, $\mathbf{W}_o \in \mathbb{R}^{d_{\text{out}} \times H d_{\text{mid}}}$
 $\mathbf{Q} \leftarrow \mathbf{W}_q \mathbf{X}$ \triangleright queries $\in \mathbb{R}^{H d_{\text{attn}} \times l_x}$
 $\mathbf{K} \leftarrow \mathbf{W}_k \mathbf{Z}$ \triangleright keys $\in \mathbb{R}^{H d_{\text{attn}} \times l_z}$
 $\mathbf{V} \leftarrow \mathbf{W}_v \mathbf{Z}$ \triangleright values $\in \mathbb{R}^{H d_{\text{mid}} \times l_z}$
for $h = 1$ **to** H **do**
 $\mathbf{S}^{(h)} \leftarrow \mathbf{K}^{(h), T} \mathbf{Q}^{(h)}$ \triangleright scores $\in \mathbb{R}^{l_z \times l_x}$
if Mask **then**
 $\mathbf{S}^{(h)}[\neg \text{Mask}] \leftarrow -\infty$
end if
 $\tilde{\mathbf{V}}^{(h)} \leftarrow \mathbf{V}^{(h)} \cdot \text{softmax}(\mathbf{S}^{(h)} / \sqrt{d_{\text{attn}}})$ $\triangleright \in \mathbb{R}^{d_{\text{mid}} \times l_x}$
end for
 $\tilde{\mathbf{V}} \leftarrow [\tilde{\mathbf{V}}^{(1)}; \dots; \tilde{\mathbf{V}}^{(H)}]$ $\triangleright \in \mathbb{R}^{H d_{\text{mid}} \times l_x}$
 $\mathbf{Y} \leftarrow \mathbf{W}_o \tilde{\mathbf{V}}$ $\triangleright \in \mathbb{R}^{d_{\text{out}} \times l_x}$

which means that in Algorithm 1, the context sequence \mathbf{Z} is equal to the primary sequence \mathbf{X} .

The decoder has L_{dec} layers and each layer contains three consecutive residual networks. The first and the second residual networks have the multi-head attention layer as the main module of the residual block, whereas the third residual network contains an MLP as the main module. As for the encoder

layer, the LN precedes each of these modules. In the first residual network of the decoder, Algorithm 1 takes as input, in addition to the primary sequence \mathbf{X} , a mask that ensures that the prediction of $\mathbf{h}_{\ell+k}$ only depends on $\{\mathbf{h}_{\ell+j}\}_{j=0}^{k-1}$. In the Algorithm 1, this is achieved by setting all the values of the input of the softmax activation function [26] to $-\infty$ which correspond to “illegal” connections. The primary and context sequence coincide. On the contrary, in the second residual network of the decoder, Algorithm 1 takes the encoder output sequence as context sequence \mathbf{Z} . In this way, the information contained in the known CSI slots can be leveraged to predict the next slot.

For the proposed model, the first element of the decoder input sequence is the last known CSI snapshot $\tilde{\mathbf{h}}_{\ell}$. This is in contrast to what happens in natural language processing since there, due to the absence of previous output at the beginning of a translation, a predefined start-of-sentence token is given as first decoder input. Additionally, during training, teacher-forcing [27] is used in the decoder, which means that during training, we use the true noisy CSI observation $\tilde{\mathbf{h}}_{\ell+1}, \dots, \tilde{\mathbf{h}}_{\ell+\delta-1}$ as further decoder inputs, instead of the predicted ones obtained at the decoder output. This helps to speed up the training process since the decoder outputs $\tilde{\mathbf{h}}_{\ell+1}, \dots, \tilde{\mathbf{h}}_{\ell+\delta}$ can be obtained in parallel. However, during testing, we have to stick to a sequential one-by-one prediction.

Similarly to the original implementation in [20] also in this case, we transform the input sequences of both, the encoder and the decoder, first by a linear layer and then by adding a constant bias term, called positional encoding (PE). The PE consists of constant, non-learnable vectors that are added after the first linear layer in both, the encoder and the decoder. Since there is no recurrence in the model, the PE is the only way to inject information about the order of the sequence. Therefore, it is crucial for the transformer architecture. In [20], and in our work, the PE is constructed with sine and cosine functions as

$$\begin{aligned} \text{PE}(j, 2i) &= \sin(j/(10000^{2i/d_{\text{model}}})) \\ \text{PE}(j, 2i+1) &= \cos(j/(10000^{2i/d_{\text{model}}}), \end{aligned} \quad (3)$$

where $j \in \{0, \ell-1\}$ is the index corresponding to the position within the sequence, $i \in \{0, \lfloor d_{\text{model}}/2 \rfloor\}$ is the index of the dimension, and d_{model} corresponds to the dimension of each CSI slot after the first linear layer. However, differently from the vanilla architecture, we introduce a novel RPE in the transformer encoder, while keeping the standard PE in the transformer decoder. Intuitively, this means that we start counting the CSI snapshots in the encoder from the last known snapshot. In other words, this enhances the robustness of the transformer to sequences of variable lengths, as the PE linked to the latest known slots remains consistent for shorter or longer sequences. The RPE can be obtained by first computing the standard PE in (3) and then by reversing the order with respect to the position index j .

To better understand the motivation behind this procedure, we can consider a simple example in which the transformer [20] with standard PE in the encoder is trained with an

encoder input sequence of length ℓ , whereas it is tested with an encoder input sequence of length ν , with $\nu \neq \ell$. During training, the transformer implicitly exploits the information contained in the most recent snapshots more than the information contained in the initial snapshots to make a good prediction. In the transformer with standard PE those are the snapshots associated with, e.g., $\text{PE}(\ell-1, :), \dots, \text{PE}(\ell-\zeta, :)$ with $\zeta < \ell$, and the colon that denotes all the elements of the corresponding row. However, when $\nu < \ell$, such model fails to make a good prediction because in this case the most important (or recent) snapshots are associated with the PEs that were linked with the initial snapshots during training, thus their importance is underestimated. On the other hand, when $\nu > \ell$, the usual ordering leads to the situation in which those snapshots, which are now outdated, are interpreted as the most recent ones by the model, and their importance is overestimated for making prediction. This problem is solved with the proposed RPE, which introduces consistency when mapping the PEs to the corresponding snapshots, and makes the transformer robust, allowing to capture the relationships between different snapshots in the sequence, regardless of the sequence length.

IV. SEQ2SEQ-ATTN-R MODEL

Another relevant framework to solve the problem at hand is the sequence-to-sequence (Seq2Seq) architecture, see [28]. Like the transformer, also the Seq2Seq model comprises an *encoder* and a *decoder* neural network, which are both RNNs in the simplest case. Specifically, the encoder RNN encodes the input sequence to produce a final state which in turn is used as initial state for the decoder RNN. The hope is that the final state of the encoder encodes all the important information about the source or input sequence such that the decoder can generate the target sequence based on this vector. However, in such setting, the decoder has to extract meaningful information from a single representation (the final state of the encoder), which can be a daunting task, especially when taking into account long sequences, or sentences. In [29] an attention mechanism has been introduced in the decoder neural network to address this problem. In particular, instead of passing only the final state of the encoder RNN, this approach involves passing all the encoder RNN states to the decoder. Hence, at each decoder step, the attention mechanism decides which parts of the source sequence are more relevant.

In the following, we propose an adapted model of [29], called Seq2Seq-attn-R model, to tackle the channel prediction task. To avoid vanishing or exploding gradient problems,¹ we opt for a GRU as RNN for the encoder. The main steps of the GRU are

$$\begin{aligned} \mathbf{z}_t &= \sigma(\mathbf{W}_z \text{concat}(\tilde{\mathbf{h}}_t, \mathbf{u}_{t-1}) + \mathbf{b}_z) \\ \mathbf{r}_t &= \sigma(\mathbf{W}_r \text{concat}(\tilde{\mathbf{h}}_t, \mathbf{u}_{t-1}) + \mathbf{b}_r) \end{aligned}$$

¹Because RNNs allow for information to be fed back to the same node multiple times, the gradients can become too small or too large, leading to unstable training. The gating mechanism of both gated recurrent unit (GRU) and long-short-term memory (LSTM) models addresses this issue.

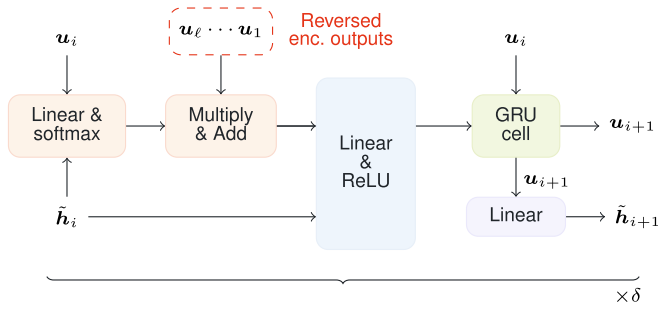


FIGURE 2. Decoder of Seq2Seq-attn-R model for CSI prediction.

$$\begin{aligned}\tilde{\mathbf{u}}_t &= \tanh(\mathbf{W}_{\tilde{u}} \text{concat}(\tilde{\mathbf{h}}_t, \mathbf{r}_t \odot \mathbf{u}_{t-1}) + \mathbf{b}_{\tilde{u}}) \\ \mathbf{u}_t &= (1 - z_t) \odot \tilde{\mathbf{u}}_t + z_t \odot \mathbf{u}_{t-1},\end{aligned}\quad (4)$$

where z_t and r_t represent the update and the reset gate, respectively, and σ denotes the sigmoid activation function [26]. In particular, when z_t is close to $\mathbf{1}$, we ignore completely the current input $\tilde{\mathbf{h}}_t$ for the update of the current hidden state \mathbf{u}_t . On the other hand, when both r_t and z_t are equal to zero, the hidden state only depends on the current input. The decoder also comprises a GRU. However, at each step, and in order to encourage the decoder to leverage the important parts of the encoder outputs before making the prediction, an attention mechanism with respect to the encoder outputs precedes the GRU. The model used for the decoder is shown in Fig. 2. In particular, the current hidden state, and the current input are concatenated, and then projected via a single layer onto a dimension equal to the maximum number of encoder outputs ℓ_{\max} . At this point, the first ℓ out of the ℓ_{\max} units are selected and the obtained vector is normalized with the softmax activation function [26] to obtain the weights (probabilities), which are multiplied with the reversed encoder outputs $\mathbf{u}_\ell, \dots, \mathbf{u}_1$. The rationale behind reversing the encoder outputs is similar to the idea of using the reverse positional encoding in the transformer encoder. Essentially, by reversing the encoder outputs, we ensure that the weights associated with the initial units out of the ℓ_{\max} units correspond to the most recent known slots. This enables the network to generalize to sequences of varying lengths. Alternatively, instead of reversing the encoder outputs, we could achieve the same goal by selecting the last (instead of the first) ℓ units out of ℓ_{\max} before applying the softmax.

Next, the result of the weighted sum of the encoder outputs, or “attention” with respect to the encoder outputs, is combined with the current decoder input and fed to a linear layer followed by a rectified linear unit (ReLU) [30] activation function to produce the second input vector for the decoder GRU. The current hidden state of the decoder GRU serves as the first input vector. Therefore, before entering the GRU, the current decoder input is preprocessed to take into account the contribution of the known slots.

Finally, and analogously with typical RNNs, the output of the GRU is fed to a linear layer to output the prediction of the next CSI vector. Like the Transformer-RPE model, the

first input of the decoder is represented by the last known CSI snapshot, and teacher-forcing [27] is deployed during training. However, differently from the Transformer-RPE model, in the Seq2Seq-attn-R model, the training happens sequentially.

V. BENCHMARKS

In this work, we consider an LSTM [31] model as further benchmark. An LSTM cell employs three different gates, an input gate i_t , a forget gate f_t , and an output gate o_t to prevent exploding or vanishing gradients. The main idea behind the gating system is not to retain information about all the inputs and to capture long-term dependencies. In formulas, we have:

$$\begin{aligned}i_t &= \sigma(\mathbf{W}_i \text{concat}(\tilde{\mathbf{h}}_t, \mathbf{u}_{t-1}) + \mathbf{b}_i) \\ f_t &= \sigma(\mathbf{W}_f \text{concat}(\tilde{\mathbf{h}}_t, \mathbf{u}_{t-1}) + \mathbf{b}_f) \\ o_t &= \sigma(\mathbf{W}_o \text{concat}(\tilde{\mathbf{h}}_t, \mathbf{u}_{t-1}) + \mathbf{b}_o) \\ \tilde{c}_t &= \tanh(\mathbf{W}_c \text{concat}(\tilde{\mathbf{h}}_t, \mathbf{u}_{t-1}) + \mathbf{b}_c) \\ c_t &= f_t \odot c_{t-1} + i_t \odot \tilde{c}_t, \quad \mathbf{u}_t = o_t \odot \tanh(c_t),\end{aligned}\quad (5)$$

where \mathbf{u}_t and c_t denote the hidden and the cell states, respectively, and σ denotes the sigmoid activation function [26]. Note that differently from the classical RNN, where the next hidden state is directly represented by \tilde{c}_t , here, the hidden state is updated using the cell state c_t . Therefore, \tilde{c}_t is first modulated by the input gate and then by the output gate. In this work, we utilize an LSTM to encode the input sequence $\{\tilde{\mathbf{h}}_i\}_{i=1}^\ell$ into \mathbf{u}_ℓ . Then, in order to predict the next δ CSI slots, we employ a final linear layer that takes \mathbf{u}_ℓ as input and directly outputs $\mathbf{y} = [\tilde{\mathbf{h}}_{\ell+1}^T, \dots, \tilde{\mathbf{h}}_{\ell+\delta}^T]^T$.

Apart from the LSTM-based model, additional benchmarks are: *i*) a two-layer MLP with a ReLU [30] activation function in the hidden layer; *ii*) a multivariate autoregressive (MAR) model of order equal to ℓ , where the coefficients are found with the ordinary least squares solution, see [32, Section 3.4.3]; *iii*) a transformer which utilizes the standard PE in the encoder, as in [20]; *iv*) the Transformer-Parallel architecture, proposed in [19]. Note that, the MAR approach, upon our investigation, emerged as the only conventional method adaptable to the challenging setup we are currently exploring. Unlike conventional solutions like the Kalman filter, the MAR based prediction does not require estimating user velocity or autocorrelation function, making it more efficient in terms of time and performance. Indeed, this would be quite inefficient in terms of both time and performance, since it would require the acquisition of many CSI snapshots during the online operational phase, and additional computation for estimating the velocity of the user. Moreover, the estimate of the velocity based on noisy observations may be inaccurate and using it as the basis for building a predictive model could compromise the reliability of the predictions.

VI. SIMULATIONS

A. SIMULATION SETUP

For the simulations, we consider CSI sequences generated with QuaDRiGa v.2.6, see [33]. In particular, we generate $N_{\text{samples}} = 150,000$ CSI sequences corresponding to 1,500 different velocities. Therefore, we have 100 users for each velocity. Every CSI sequence corresponds to a frame which contains $N_{\text{slot}} = 20$ slots, each with duration $T_{\text{slot}} = 0.5$ ms. The carrier frequency is 2.6 GHz and each velocity v measured in m/s is Rayleigh distributed: $\frac{v}{\gamma^2} \exp(-\frac{v^2}{2\gamma^2})$, where $\gamma = 8$. The reason for this choice is to simulate a realistic urban scenario, where the majority of the MTs move within a range of 20 to 50 km/h. However, there are MTs with $v < 20$ km/h (e.g., pedestrians and cyclists), as well as a few MTs with velocities exceeding 100 km/h (e.g., fast moving cars). The scenario is the “BERLIN_UMa_NLOS” which generates non-line-of-sight channels with 25 paths. Specifically, the channel model we used in this study is based on measurements conducted in Berlin. QuaDRiGa places scatterers in the surroundings of the BS. Scatterer positions are determined from drawn small-scale parameters, meeting specific requirements (e.g., 3GPP standards) for large-scale wireless channel parameters like delay and angular spreads. Once scatterer positions are set, QuaDRiGa deterministically calculates path lengths, angles, delays, etc., enabling the generation of time-evolving channels consistent with its environment [33], [34]. The BS at a height of 25,m has a uniform rectangular array with $M = 32$ antennas (8 vertical, 4 horizontal). Users’ initial positions are randomly distributed over a 120 deg sector, with distances from the BS ranging from 50,m to 150,m at a height of 1.5,m. Generated CSI sequences are normalized by path-gain, then split into training, validation, and test sets (80%, 10%, 10%).

Simulations consider varying noise levels, corrupting channel \mathbf{h}_i (described in (1)) with noise variance σ_n^2 to achieve a specific average SNR. The formula relates the average SNR to σ_n^2 .

$$\text{SNR} = \frac{1}{N_{\text{samples}}N_{\text{slot}}} \sum_{j=1}^{N_{\text{samples}}} \sum_{i=1}^{N_{\text{slot}}} \|\mathbf{h}_i^{(j)}\|^2 / M\sigma_n^2 \quad (6)$$

where $\mathbf{h}_i^{(j)}$ denotes the CSI vector in the i 'th slot of the j 'th sample in the dataset. For our simulations, we assume that the first $\ell = 16$ noisy CSI realizations are known. Therefore, the goal is to predict the next $\delta = 4$ noisy CSI vectors.

The performance metric that we consider is the normalized mean squared error (NMSE) with respect to the test set between the true noiseless CSI and the CSI predicted by the different models based on the noisy observations of the previous slots. In formulas, we have:

$$\text{NMSE} = \frac{1}{N_{\text{test}}} \sum_{j=1}^{N_{\text{test}}} \epsilon_j^2 \quad \epsilon_j = \frac{\|\mathbf{H}^{(j)} - \hat{\mathbf{H}}^{(j)}\|_F}{\|\mathbf{H}^{(j)}\|_F} \quad (7)$$

where \mathbf{H} is the matrix that consists of the clean CSI snapshots, and $\hat{\mathbf{H}} = [\hat{\mathbf{h}}_{\ell+1}, \dots, \hat{\mathbf{h}}_{\ell+\delta}]$ is the matrix that consists

of the corresponding predicted CSI snapshots. During training, we assume that only noisy data are available. Therefore, the loss function is given by the NMSE between $\hat{\mathbf{H}} = [\hat{\mathbf{h}}_{\ell+1}, \dots, \hat{\mathbf{h}}_{\ell+\delta}]$ and \mathbf{H} . All the models are trained separately for each SNR value. The number of epochs is set to 500 and the batch size is equal to 200. The Adam optimizer (see [35]) with learning rate equal to 10^{-3} is used. For each model, the parameters leading to the smallest NMSE concerning the validation set, between the noisy CSI and the predicted one, are saved and considered during the test phase. It is important to notice that the clean CSI is exclusively used during testing to evaluate performance. Throughout training and validation, only noisy CSI observations are employed. Additionally, all the NNs in this study were trained on a single GPU, specifically the Nvidia Tesla V100-PCIE-16GB.

B. MODELS PARAMETERS

For the Transformer-RPE described in Section III we consider $L_{\text{enc}} = L_{\text{dec}} = 2$. Furthermore, in Algorithm 1 we set $H = 4$, and $d_{\text{attn}} = d_{\text{mid}} = 16$, while the observed dimensions that correspond to the dimension of the real-valued CSI snapshots are $d_x = d_y = d_{\text{out}} = 64$. Consequently, $d_{\text{model}} = 64$. The MLP block in both, the encoder and the decoder, has a hidden dimension equal to 128.

For the Seq2Seq-attn-R model described in Section IV, we consider GRUs with two layers and hidden states with dimension equal to 128. In the decoder, the first linear layer followed by the softmax activation maps $(64 + 128)$ units to $\ell_{\text{max}} = 20$ units, where 64 is the input dimension (the dimension of $\hat{\mathbf{h}}_i$), 128 is the dimension of the hidden state \mathbf{u}_i , and the addition is due to the concatenation of the two. On the other hand, the second linear layer followed by a ReLU maps $(64 + 128)$ units to 64 units, where 128 is the resulting dimension after the “Multiply & Add” block. The decoder’s final linear layer, which produces the CSI prediction for the next step, maps the 128 units of the next hidden state to the 64 units of the next CSI. The dimensions of all the weights matrices and bias vectors of the GRU, that appear in (4) can be derived with the given information.

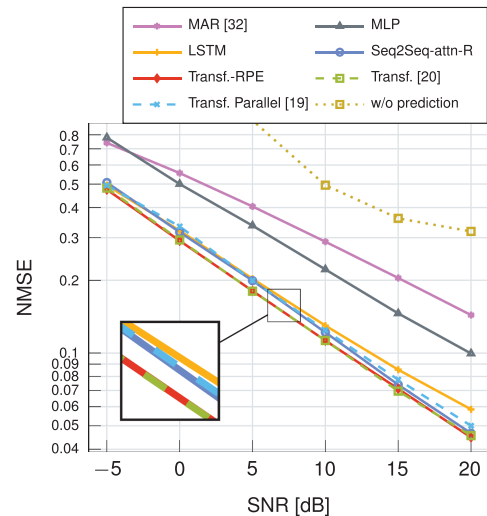
For the LSTM model, we have considered a two-layer LSTM, with hidden states with dimension equal to 128, and a final linear layer which maps the last hidden state to the output dimension which is equal to 256. The dimensions of all the parameters which appear in (5) can be inferred with the given information. In the MLP model, the observed dimension, the hidden dimension, and the output dimension are 1024, 512, and 256, respectively. The Transformer with standard PE has the same parameters as the Transformer-RPE. For the Transformer-Parallel, the decoder input is initialized with the past 8 CSI snapshots followed by δ snapshots initialized as all-zeros vectors, while the remaining parameters have the same values as for the Transformer-RPE. The parameters of our models are chosen heuristically, considering task complexity, dataset size, and the risk of overfitting.

C. NUMERICAL RESULTS

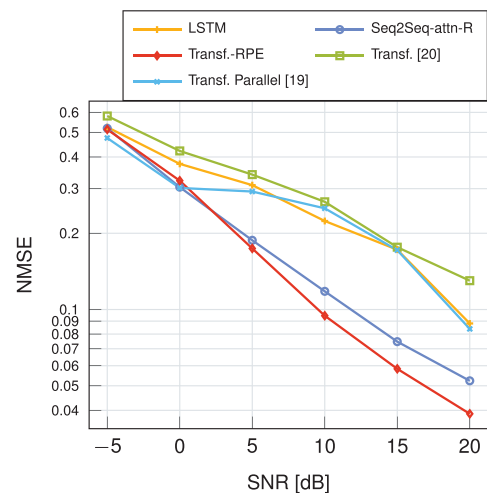
In Fig. 3, the NMSE vs. average SNR is displayed for all the models described within this work. In Fig. 3(a), we can observe that the models designed for sequential data, such as LSTM, Seq2Seq-attn-R, and all the Transformer-based models, outperform both the MAR and the MLP models. Moreover, among these, the models that include an attention mechanism outperform the LSTM model, despite the fact that the LSTM predict all the δ snapshots in one step. This means that in the LSTM, an imperfect prediction for the $\ell + 1$ snapshot has no influence regarding the prediction of the future $\delta - 1$ snapshots. Additionally, we observe that both the proposed Transformer-RPE and the one of [20] outperform the Transformer-Parallel [19] model in all the cases. This is because the former architectures can leverage the predicted CSI snapshots to make better prediction for the next one. And, at moderate or high SNR levels, this is advantageous for making a more accurate channel prediction. To highlight the importance of channel prediction, we compare it to a control group without prediction, which maintains the last known CSI snapshot for subsequent δ snapshots. Additionally, we assess the performance of the sequential models trained for $\ell = 16$ and $\delta = 4$ on different sequence lengths as a benchmark.

In Fig. 3(b), the NMSE of various models for $\ell = 8$ and $\delta = 2$ is displayed. Transformer-RPE and Seq2Seq-attn-R outperform other models, except at low SNR where Transformer-Parallel is slightly better. However, in low SNR for all models, NMSE is too high (>0.2), indicating poor performance in this region. Our approach outperforms Transformer-Parallel [19] mainly due to two factors. Firstly, we prioritize sequential prediction of each snapshot, leveraging the dynamic decoding process, unlike Transformer-Parallel, which predicts all snapshots in a single step. Secondly, our method uses reversed positional encoding, whereas Transformer-Parallel uses classical positional encoding in the encoder. By comparing Transformer-RPE and Transformer [20], improvements in generalization are evident due to the introduced RPE. It's essential to note that in Fig. 3(b), only half of the vector \mathbf{y} from Section V is considered for LSTM.

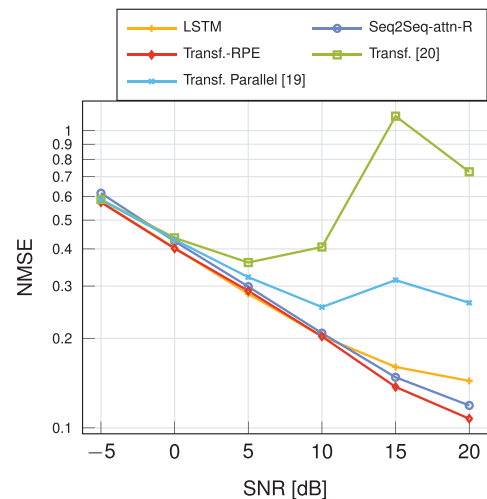
In Fig. 3(c), the NMSE of the different models for $\ell = 14$ and $\delta = 6$ is shown. In this case, the proposed models outperform both, the Transformer with PE and the Transformer-Parallel. At the same time, in Fig. 3(c), we can observe that the performance of the LSTM is very close to those of both, the Transformer-RPE and the Seq2Seq-attn-R model. Furthermore, to obtain the results in Fig. 3(c), we have to sweep twice through LSTM model: the first time to predict the next 4 snapshots and the second time, it uses the known snapshots together with the predicted ones to predict the remaining 2 snapshots. In both, Fig. 3(b) and (c), the presence of peaks in the lines of both, the Transformer with PE and the Transformer-Parallel, is due to the fact that these models trained for $\ell = 16$ and $\delta = 4$ cannot generalize to other sequence lengths. These results highlight the robustness of the



(a) $\ell = 16, \delta = 4$ (same as training).



(b) $\ell = 8, \delta = 2$.



(c) $\ell = 14, \delta = 6$.

FIGURE 3. NMSE vs. SNR.

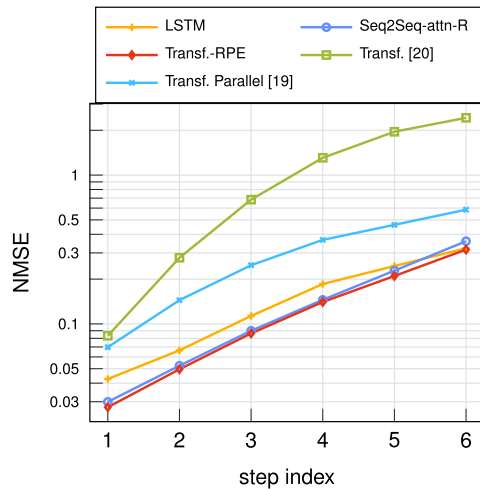


FIGURE 4. NMSE vs. prediction step for SNR=15 dB, $\ell = 14$, $\delta = 6$.

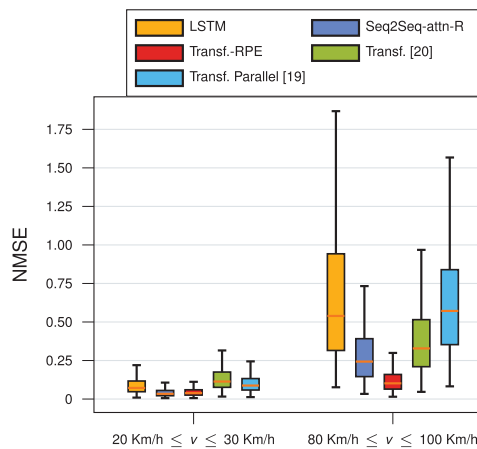


FIGURE 5. NMSE for different velocities and models for SNR=15 dB, $\ell = 8$, $\delta = 2$.

proposed models and show that they can generalize to any sequence length, as opposed to existing methods.

In Fig. 4, we depict NMSE versus prediction step for SNR = 15 dB, $\ell = 14$, and $\delta = 6$. Notably, prediction error increases with advancing time steps into the future. The proposed models, Transformer-RPE and Seq2Seq-attn-R outperform both, the Transformer with standard positional encoding [20] and the Transformer-Parallel [19]. The LSTM model shows a comparable prediction error, but it requires two passes, as discussed earlier. Additionally, in Fig. 5, we present the NMSE for two velocity ranges under SNR = 15 dB, $\ell = 8$, $\delta = 2$. The proposed models exhibit superior performance across various velocity ranges when compared to other models analyzed in this study. As expected, all models exhibit a smaller NMSE in the lower range of [20, 30] Km/h as opposed to the higher range of [80, 100] Km/h.

Table 1 shows model complexity in terms of parameters and floating point operations (FLOPs) for $\ell = 16$ and $\delta = 4$. Sequential data models have similar parameters and FLOPs.

TABLE 1. Number of Parameters and Complexity for $\ell = 16$, $\delta = 4$

Model	# parameters	# FLOPs
LSTM	264, 448	6.37×10^6
Seq2Seq-attn-R	370, 832	6.13×10^6
Transf.-RPE ²	178, 752	5.96×10^6
Transf. Parallel [19]	178, 752	4.68×10^6

Transformer-based models have the fewest parameters, while Transformer Parallel has the fewest FLOPs. This is due to Transformer Parallel predicting all δ CSI snapshots simultaneously, while Transformer-RPE iterates over the transformer decoder δ times, considering previous predictions.

In summary, the proposed models, particularly the Transformer-RPE, offer more accurate channel prediction compared to existing models while maintaining the complexity at the same order of magnitude. Additionally, the fact that the proposed models exhibit robust results for different sequence lengths highlights that in a practical scenario it is sufficient to train a single model instead of having to train a different model for each combination of ℓ and δ , which saves computational power, as well as storage requirements at the BS. The proposed approach, like many deep learning methods, relies on the assumption that the dataset accurately represents a specific cell's channel statistics. Consequently, the prediction model, such as Transformer-RPE, is tailored to a specific cell and requires training on data from that cell. If applied to a different cell, retraining the Transformer-RPE is necessary. However, transfer learning principles [36] can be used to extend the model's applicability to multiple cells simultaneously. Nevertheless, this extension is beyond the scope of the current work.

VII. CONCLUSION

In this study, we introduce two models for predicting time-varying channels: Transformer-RPE and Seq2Seq-attn-R. Both outperform existing methods in channel prediction accuracy across different noise levels and generalize to unseen sequence lengths. The proposed reverse ordering techniques are straightforward and effective, offering an attractive solution to avoid training multiple neural networks for varying lengths and thereby reducing training complexity and storage requirements at the BS. This is crucial in 5G and future wireless systems, emphasizing flexibility and resource efficiency. Future research could extend models to multiple-input-multiple-output (MIMO) channels, beyond the current focus on multiple-input-single-output (MISO) channels. Additionally, improving the models to predict precoding indices rather than channel values may offer a more practical solution for real-world systems.

²same as Transf. [20].

REFERENCES

- [1] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [2] K. T. Truong and R. W. Heath, "Effects of channel aging in massive MIMO systems," *J. Commun. Netw.*, vol. 15, no. 4, pp. 338–351, 2013.
- [3] K. E. Baddour and N. C. Beaulieu, "Autoregressive modeling for fading channel simulation," *IEEE Trans. Wireless Commun.*, vol. 4, no. 4, pp. 1650–1662, Jul. 2005.
- [4] A. Barbieri, A. Piemontese, and G. Colavolpe, "On the ARMA approximation for fading channels described by the Clarke model with applications to Kalman-based receivers," *IEEE Trans. Wireless Commun.*, vol. 8, no. 2, pp. 535–540, Feb. 2009.
- [5] A. H. E. Hussein, E. P. Simon, and L. Ros, "Optimization of the second order autoregressive model AR(2) for Rayleigh-Jakes flat fading channel estimation with Kalman filter," in *Proc. IEEE 22nd Int. Conf. Signal Process.*, 2017, pp. 1–5.
- [6] T. Zemen, C. F. Mecklenbrauker, F. Kaltenberger, and B. H. Fleury, "Minimum-energy band-limited predictor with dynamic subspace selection for time-variant flat-fading channels," *IEEE Trans. Signal Process.*, vol. 55, no. 9, pp. 4534–4548, Sep. 2007.
- [7] J. Yuan, H. Q. Ngo, and M. Matthaiou, "Machine learning-based channel prediction in massive MIMO with channel aging," *IEEE Trans. Wireless Commun.*, vol. 19, no. 5, pp. 2960–2973, May 2020.
- [8] K. Pratik, R. A. Amjad, A. Behboodi, J. B. Soriaga, and M. Welling, "Neural augmentation of Kalman filter with hypernetwork for channel tracking," in *Proc. IEEE Glob. Commun. Conf.*, 2021, pp. 1–6.
- [9] H. Kim, S. Kim, H. Lee, C. Jang, Y. Choi, and J. Choi, "Massive MIMO channel prediction: Kalman filtering vs machine learning," *IEEE Trans. Commun.*, vol. 69, no. 1, pp. 518–528, Jan. 2021.
- [10] G. Liu, Z. Hu, L. Wang, J. Xue, H. Yin, and D. Gesbert, "Spatio-temporal neural network for channel prediction in massive MIMO-OFDM systems," *IEEE Trans. Commun.*, vol. 70, no. 12, pp. 8003–8016, Dec. 2022.
- [11] S. Park and O. Simeone, "Predicting multi-antenna frequency-selective channels via meta-learned linear filters based on long-short term channel decomposition," 2022, *arXiv:2203.12715*.
- [12] H. Kim, J. Choi, and D. J. Love, "Massive MIMO channel prediction via meta-learning and deep denoising: Is a small dataset enough?," *IEEE Trans. Wireless Commun.*, vol. 22, no. 12, pp. 9278–9290, Dec. 2023.
- [13] W. Jiang, M. Strufe, and H. D. Schotten, "Long-range MIMO channel prediction using recurrent neural networks," in *Proc. IEEE 17th Annu. Consum. Commun. Netw. Conf.*, 2020, pp. 1–6.
- [14] W. Jiang and H. D. Schotten, "Neural network-based fading channel prediction: A comprehensive overview," *IEEE Access*, vol. 7, pp. 118112–118124, 2019.
- [15] M. K. Shehzad, L. Rose, S. Wesemann, M. Assaad, and S. A. Hassan, "Design of an efficient CSI feedback mechanism in massive MIMO systems: A machine learning approach using empirical data," 2022, *arXiv:2208.11951*.
- [16] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *Proc. 30th Annu. Int. Conf. Mach. Learn.*, 2013, pp. 1310–1318.
- [17] W. Jiang and H. D. Schotten, "Deep learning for fading channel prediction," *IEEE Open J. Commun. Soc.*, vol. 1, pp. 320–332, 2020.
- [18] M. Chu, A. Liu, C. Jiang, V. K. N. Lau, and T. Yang, "Wireless channel prediction for multi-user physical layer with deep reinforcement learning," in *Proc. IEEE 95th IEEE Veh. Technol. Conf.*, 2022, pp. 1–5.
- [19] H. Jiang, M. Cui, D. W. K. Ng, and L. Dai, "Accurate channel prediction based on transformer: Making mobility negligible," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 9, pp. 2717–2732, Sep. 2022.
- [20] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.
- [21] N. Wu, B. Green, X. Ben, and S. O'Banion, "Deep transformer models for time series forecasting: The influenza prevalence case," 2020, *arXiv:2001.08317*.
- [22] 3GPP "Radio resource control (RCC) protocol specification," Tech. Specification 38.331, 3GPP, Sophia Antipolis, France, 2019.
- [23] J. Hoydis et al., "Sionna: An open-source library for next-generation physical layer research," 2022, *arXiv: 2203.11854*.
- [24] L. J. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, *arXiv:1607.06450*.
- [25] D. Hendrycks and K. Gimpel, "Gaussian error linear units (GELUs)," 2016, *arXiv:1606.08415*.
- [26] T. Hastie, J. H. Friedman, and R. Tibshirani, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Berlin, Germany: Springer, 2001.
- [27] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural Comput.*, vol. 1, no. 2, pp. 270–280, 1989.
- [28] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 3104–3112.
- [29] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [30] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>
- [31] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [32] S. Weisberg, *Applied Linear Regression*, 4th ed. Hoboken, NJ, USA: Wiley, 2014.
- [33] S. Jaekel, L. Raschkowski, K. Boerner, and L. Thiele, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE Trans. Antennas Propag.*, vol. 62, no. 6, pp. 3242–3256, Jun. 2014.
- [34] 3GPP "Study on channel model for frequencies from 0.5 to 100 GHz," Tech. Specification 38.901, 3GPP, Sophia Antipolis, France, 2017.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [36] F. Zhuang et al., "A comprehensive survey on transfer learning," *Proc. IEEE*, vol. 109, no. 1, pp. 43–76, Jan. 2021.