
Towards a Foundation Model for Communication Systems

Davide Buffelli^{*1} Sowmen Das^{*1} Yu-Wei Lin^{*2} Sattar Vakili^{*1} Chien-Yi Wang⁵² Masoud Attarifar³
Pritthijit Nath⁴ Da-shan Shiu⁵

Abstract

Artificial Intelligence (AI) has demonstrated unprecedented performance across various domains, and its application to communication systems is an active area of research. While current methods focus on task-specific solutions, the broader trend in AI is shifting toward large *general* models capable of supporting multiple applications. In this work, we take a step toward a *foundation model for communication data*—a transformer-based, multi-modal model designed to operate directly on communication data. We propose methodologies to address key challenges, including tokenization, positional embedding, multimodality, variable feature sizes, and normalization. Furthermore, we empirically demonstrate that such a model can successfully estimate multiple features, including transmission rank, selected precoder, Doppler spread, and delay profile.

1. Introduction

In modern communication systems, numerous estimation and optimization tasks are required to ensure effective signal transmission and reception. Some examples include channel estimation, Doppler spectrum and delay spread estimation, and channel state information (CSI) feedback, which includes precoding matrix and its rank.

In the past few years, an increasingly large amount of research has focused on replacing traditional solutions to the above tasks with machine learning, and deep learning in particular (Žeger & Šišul, 2021; Thakkar et al., 2020). These methods are based on the idea of learning from data rather than relying on handcrafted heuristics, paving the way for

more intelligent and efficient networks. Indeed, AI is already expected to become a key component of 6G networks (Wang et al., 2023) and will play an even more central role in future communication systems (Saad et al., 2024).

Current research in this domain has focused on designing task-specific machine learning solutions. These approaches involve collecting large amounts of supervised data, which is costly and time-consuming, and require a dedicated model for each task. In contrast, the broader field of AI is increasingly adopting a paradigm in which a single large model, referred to as a *foundation model*, is trained on vast amounts of unsupervised (and thus easier to obtain) data. Such a foundation model learns intricate relationships in the data, and is then leveraged to perform downstream tasks, by either applying it directly, specializing it with some additional fine-tuning, or by using its learned representations.

In this work, we move the first steps towards a *Foundation Model for Communication Systems*. This is a large deep learning model trained on vast quantities of unsupervised data coming from network communication. While foundation models have been proposed for domains like natural language, code, images, physical data, and time series, communication systems present unique challenges that are not addressed by the current literature. In more detail, dealing with multiple features at each slot in the input sequence, features from different domains (categorical, integer, real and complex valued, matrices), and with varying sizes and magnitudes, requires carefully designed ad-hoc procedures and components. We summarize our contributions as follows:

- We design and train a foundation model for communication systems addressing key challenges including tokenization, positional embedding, multimodality, features of varying size, and normalization.
- We design a simulation system to generate unsupervised data for the training of our foundation model.
- We test our foundation model in several scenarios, showing it can successfully perform multiple estimations.
- We report experimental results on how model size and dataset size affect estimation performance, providing

^{*}Equal contribution ¹MediaTek Research, London, UK ²MediaTek, Hsinchu, R.O.C. ³MediaTek, Cambourne, UK ⁴University of Cambridge, Cambridge, UK ⁵MediaTek Research, Taipei, R.O.C.. Correspondence to: Davide Buffelli <davide.buffelli@mtkresearch.com>, Sattar Vakili <sattar.vakili@mtkresearch.com>.

initial insights into the scaling behavior of foundation models for communication systems (see Figures 2 and 3).

2. Related Work

We focus our discussion on foundation models, and large language models (LLMs) applied to communication applications. The literature on machine learning for communication systems is rapidly growing, and a complete presentation would be out of the scope of this paper. We refer the interested reader to recent surveys on machine learning (Sun et al., 2024; Wang & Li, 2024; Hamdan et al., 2023), and deep learning in particular (Yu et al., 2022; Liao et al., 2020), in this domain.

2.1. Foundation Models

The term *foundation model* was introduced in the context of LLMs (Bommasani et al., 2021), i.e., models trained on large quantities of textual data. Multimodal foundation models that combine data from different modalities have also been introduced. The most popular such models combine text with images (Alayrac et al., 2022; Chen et al., 2025; Dai et al., 2024; Team et al., 2024; OpenAI, 2024).

Recently, significant effort has been dedicated to creating foundation models outside of the traditional textual, visual, or audio domains. Examples of this are foundation models for signals from wearable devices (Narayanswamy et al., 2025; Abbaspourazad et al., 2024), earth-related signals (Bodnar et al., 2024), and time-series (Goswami et al., 2024; Rasul et al., 2024; Shi et al., 2025; Das et al., 2024; Liu et al., 2024; Ekambaram et al., 2024; Liu et al., 2025).

While these models share commonalities to our work, communication data poses unique challenges that are not addressed in current literature. In particular, prior works focus on settings with few features and low heterogeneity—unlike communication systems, as discussed in Section 4.

2.2. LLMs for Communication Systems

With the rise in popularity of LLMs, several works have focused on introducing them into the context of communication systems. (Shao et al., 2024) presents the challenges that need to be addressed to make existing pre-trained LLMs able to understand data from communication system. (Nikbakht et al., 2024) introduce a dataset containing numerous technical documents that can be used to introduce knowledge related to communication systems to LLMs. Finally, (Piovesan et al., 2024) shows that “small” language models enhanced with Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) approaches can already understand communications data, and (Bornea et al., 2024) further improve the RAG procedure for specialized tech-

nical documents. These works aim to enrich LLMs with communication-related knowledge but do not operate on low-level data. In contrast, our model directly processes raw, heterogeneously structured communication data to perform effective feature estimation, without relying on external knowledge sources such as natural language.

3. Communication System Model and Simulator

In this section, we describe the communication system model and the simulated data used to train the foundation model.

3.1. Communication System Model

We consider a multiple-input multiple-output (MIMO) orthogonal frequency-division multiplexing (OFDM) system of N_T transmit (Tx) antennas, N_R receive (Rx) antennas, K subcarriers per OFDM symbol with subcarrier spacing f_{sc} Hz, and L OFDM symbols per slot. We restrict our attention to channel state information (CSI) acquisition, which is an important subsystem of wireless communications. For CSI acquisition, the transmitter sends pilot signals to the receiver. The receiver estimates the CSI and then reports to the transmitter.

The input–output in the frequency domain are related as

$$\mathbf{y}[k, l, n] = \mathbf{H}[k, l, n]\mathbf{x}[k, l, n] + \mathbf{z}[k, l, n],$$

where $\mathbf{y}[k, l, n] \in \mathbb{C}^{N_R}$ is the Rx signal vector at subcarrier k and symbol l in slot n , $\mathbf{x}[k, l, n] \in \mathbb{C}^{N_T}$ is the Tx signal vector, $\mathbf{H}[k, l, n] \in \mathbb{C}^{N_R \times N_T}$ is the channel matrix drawn from a wide-sense stationary random process of mean zero and unit power per entry, and $\mathbf{z}[k, l, n] \in \mathbb{C}^{N_R}$ is the additive noise independent of \mathbf{H} . We assume $\{\mathbf{z}[k, l, n]\}$ are independent and identically distributed (i.i.d.) circularly-symmetric complex Gaussian $\mathcal{CN}(\mathbf{0}, \mathbf{C}_n)$, where $\mathbf{0}$ denotes an all-zero vector and \mathbf{C}_n denotes the noise covariance matrix. For data transmission, each modulated symbol vector $\mathbf{s} \in \mathcal{A}$ is precoded by $\mathbf{W}_{(R)} \in \mathcal{W}_R \subset \mathbb{C}^{N_T \times R}$, where \mathcal{A} is the set of constellation points, R denotes the rank (number of layers), and \mathcal{W}_R is the codebook of rank R known by the transmitter and the receiver. We assume that $\mathbb{E}[\|\mathbf{s}\|^2] = P$ and $\|\mathbf{W}_{(R)}\|_F^2 = 1$ for all R , where $\|\cdot\|_F$ denotes the Frobenius norm. We assume that the same precoder is applied to all subcarriers and symbols within a slot. Then, we have $\mathbf{x}[k, l, n] = \mathbf{W}_{(R)}[n]\mathbf{s}[k, l, n]$ and thus the input–output relation can be expanded as

$$\mathbf{y}[k, l, n] = \mathbf{H}[k, l, n]\mathbf{W}_{(R)}[n]\mathbf{s}[k, l, n] + \mathbf{z}[k, l, n].$$

We divide the K subcarriers into B subcarrier groups of size M , so that $K = BM$. Denote $\mathcal{P} = \{(Mm, l) : m \in$

$\{1, \dots, B\}, l \in \mathcal{S}\}$, where \mathcal{S} denotes the set of symbols with pilots. The performance of CSI acquisition can be assessed by spectral efficiency defined as

$$G(R, \mathbf{W}_{(R)}) = \frac{1}{|\mathcal{P}|} \sum_{(k,l) \in \mathcal{P}} \log \det \left(\mathbf{C}_n + P \mathbf{H}[k, l] \mathbf{W}_{(R)} \mathbf{W}_{(R)}^H \mathbf{H}^H[k, l] \right),$$

which follows the capacity of MIMO channel with CSI at receiver (see e.g., Section 8.2.1 of (Tse & Viswanath, 2005)).

To measure the channel matrices $\{\mathbf{H}[k, l]\}$, we transmit pilots following a comb-type structure: For each Tx antenna j , we set $\mathbf{x}[Mm + j, l, n] = \sqrt{P} \mathbf{e}_j$, where $l \in \mathcal{S}$ and \mathbf{e}_j denotes the standard unit vector with one on the j -th position and zero otherwise. Besides, the noise covariance \mathbf{C}_n needs to be estimated, so we set $\mathbf{x}[Mm + j, l, n] = \mathbf{0}$ for $j = N_T + 1, \dots, N_T + N_Z$ and assume $N_Z = M - N_T$.

The receiver needs to determine $(R, \mathbf{W}_{(R)})$ from the received signals at pilot subcarriers, also known as channel frequency response (CFR): For $j \in \{1, \dots, N_T\}$,

$$\begin{aligned} \tilde{\mathbf{h}}_j[m, l, n] &\triangleq \mathbf{y}[Mm + j, l, n] \\ &= \sqrt{P} \mathbf{h}_j[Mm + j, l, n] + \mathbf{z}[Mm + j, l, n], \end{aligned}$$

where \mathbf{h}_j is the column j of \mathbf{H} , and for $k \in \{1, \dots, BN_Z\}$,

$$\tilde{\mathbf{z}}[k, l, n] \triangleq \mathbf{z}[M \lfloor (k-1)/N_Z \rfloor + ((k-1) \bmod N_Z) + 1, l, n].$$

To simplify notation, we assume per-slot operation and hence omit the slot index n in the remainder of this section.

3.1.1. NOISE COVARIANCE ESTIMATION

The noise covariance can be estimated by

$$\hat{\mathbf{C}}_n = \frac{1}{BN_Z |\mathcal{S}|} \sum_{k=1}^{BN_Z} \sum_{l \in \mathcal{S}} \tilde{\mathbf{z}}[k, l] \tilde{\mathbf{z}}^H[k, l].$$

The average noise power at Rx antenna i is denoted as $\hat{\sigma}_i^2 \triangleq [\hat{\mathbf{C}}_n]_{ii}$.

3.1.2. SIGNAL POWER ESTIMATION

The signal power P can be estimated as

$$\hat{P} = \frac{1}{N_T B |\mathcal{S}| N_R} \sum_{j=1}^{N_T} \sum_{m=1}^B \sum_{l \in \mathcal{S}} \left\| \tilde{\mathbf{h}}_j[m, l] \right\|^2 - \frac{1}{N_R} \sum_{i=1}^{N_R} \hat{\sigma}_i^2.$$

It is desirable to perform channel denoising before determining the rank and precoder. We perform robust channel estimation by assuming the power delay profile and the Doppler spectrum are both of rectangular shape (Li et al., 1998).

3.1.3. DELAY PROFILE ESTIMATION

For robust channel estimation in frequency domain, we assume that any two channel gains separated by Δ pilots in the same OFDM symbol, say $h_{(f),m}$ and $h_{(f),m+\Delta}$, satisfy that

$$\begin{aligned} &\frac{\mathbb{E}[h_{(f),m+\Delta} h_{(f),m}^*]}{\sqrt{\mathbb{E}[|h_{(f),m+\Delta}|^2] \mathbb{E}[|h_{(f),m}|^2]}} \\ &= e^{-j2\pi\mu\Delta M f_{sc}} \frac{\sin(\pi\ell\Delta M f_{sc})}{\pi\ell\Delta M f_{sc}}, \end{aligned}$$

where μ and ℓ are the center and the length of the delay profile, respectively. The center and the length are estimated by the following procedure: For the tuple of Tx antenna j , Rx antenna i , and symbol l , we collect its associated CFR into a vector $\tilde{\mathbf{h}}_{(f)}[i, j, l]$ and transform the CFR into the channel impulse response (CIR) $\tilde{\mathbf{h}}_{(t)}[i, j, l]$ through inverse discrete Fourier transform (IDFT) after zero padding to size N_{FFT} , which is the smallest power of 2 larger than K . The noisy delay profile $\tilde{\mathbf{p}}_{(t)}[i]$ at Rx antenna i is then calculated as

$$\tilde{\mathbf{p}}_{(t)}[i] = \frac{1}{N_T |\mathcal{S}|} \sum_{j=1}^{N_T} \sum_{l \in \mathcal{S}} \left| \tilde{\mathbf{h}}_{(t)}[i, j, l] \right|^2.$$

Only the tap(s) with power larger than $3\hat{\sigma}_i^2$ are considered as including the desired signal. Denote $\mathcal{D}[i] = \{n \mid [\tilde{\mathbf{p}}_{(t)}[i]]_n > 3\hat{\sigma}_i^2, 0 \leq n < N_{\text{FFT}}\}$. Then, the starting and ending positions of the delay profile ($\hat{n}_{\text{start}}[i], \hat{n}_{\text{end}}[i]$) are found as the pair that covers $\mathcal{D}[i]$ with the minimum length, considering invariance under circular shift. Finally, we have

$$\begin{aligned} \hat{\mu}[i] &= \frac{1}{M f_{sc} N_{\text{FFT}}} \left(\frac{\hat{n}_{\text{start}}[i] + \hat{n}_{\text{end}}[i]}{2} \right), \\ \hat{\ell}[i] &= \frac{1}{M f_{sc} N_{\text{FFT}}} (\hat{n}_{\text{start}}[i] - \hat{n}_{\text{end}}[i] + 1), \end{aligned}$$

and the estimated correlation matrix $\hat{\mathbf{R}}_{f,\text{robust}}$ with

$$\begin{aligned} &[\hat{\mathbf{R}}_{f,\text{robust}}]_{m_1, m_2} \\ &= e^{-j2\pi\hat{\mu}[i](m_1-m_2)M f_{sc}} \frac{\sin(\pi\hat{\ell}[i](m_1-m_2)M f_{sc})}{\pi\hat{\ell}[i](m_1-m_2)M f_{sc}}. \end{aligned}$$

In case that $\mathcal{D}[i]$ is empty, we set $\hat{\mu}[i] = \frac{\arg \max_n [\tilde{\mathbf{p}}_{(t)}[i]]_n}{M f_{sc} N_{\text{FFT}}}$ and $\hat{\ell}[i] = \frac{1}{M f_{sc} N_{\text{FFT}}}$. Following the above procedure, for each Rx antenna i , the genie center $\mu[i]$ and the genie length $\ell[i]$ for the noiseless delay profile $\mathbf{p}_{(t)}[i]$ can be calculated. With an abuse of notation, we denote by $\mathbf{R}_{f,\text{robust}}[i]$ the genie robust frequency correlation matrix.

3.1.4. DOPPLER SPECTRUM ESTIMATION

For robust channel estimation in time domain, we assume that any two channel gains $h_{(t),l_1}$ and $h_{(t),l_2}$ of the same

subcarrier in symbol l_1 and l_2 , respectively, satisfy that

$$\frac{\mathbb{E}[h_{(t),l_1}h_{(t),l_2}^*]}{\sqrt{\mathbb{E}[|h_{(t),l_1}|^2]\mathbb{E}[|h_{(t),l_2}|^2]}} = \frac{\sin(\pi w(l_1 - l_2)T)}{\pi w(l_1 - l_2)T},$$

where w is the width of the Doppler spectrum and T is the OFDM symbol duration. For the tuple of Tx antenna j , Rx antenna i , and subcarrier group m , we collect its associated received signals into a vector $\tilde{\mathbf{h}}_{(t)}[i, j, m]$. Then, the time covariance matrix of Rx antenna i can be estimated as

$$\hat{\mathbf{C}}_{\text{time}}[i] = \frac{1}{N_TB} \sum_{j=1}^N \sum_{m=1}^B \tilde{\mathbf{h}}_{(t)}[i, j, m] \tilde{\mathbf{h}}_{(t)}^H[i, j, m] - \hat{\sigma}_i^2 \mathbf{I}_{|S|}.$$

The corresponding estimated time correlation matrix is denoted by $\hat{\mathbf{R}}_{\text{time}}[i]$, which can be calculated from the entries of $\hat{\mathbf{C}}_{\text{time}}$ by $r_{ij} = c_{ij} / \sqrt{c_{ii}c_{jj}}$. Then, the width w can be found by

$$\hat{w}[i] = \arg \min_w \|\mathbf{R}_{\text{t,robust}}(w) - \hat{\mathbf{R}}_{\text{time}}[i]\|_F^2,$$

where

$$[\mathbf{R}_{\text{t,robust}}(w)]_{l_1, l_2} = \frac{\sin(\pi w(l_1 - l_2)T)}{\pi w(l_1 - l_2)T}.$$

Denote $\hat{\mathbf{R}}_{\text{t,robust}}[i] = \mathbf{R}_{\text{t,robust}}(\hat{w}[i])$ the estimated robust time correlation matrix of Rx antenna i . The genie time covariance matrix $\hat{\mathbf{C}}_{\text{time}}[i]$ is calculated by replacing $\tilde{\mathbf{h}}_{(t)}[i, j, m]$ and $\hat{\sigma}_i^2$ by $\mathbf{h}_{(t)}[i, j, m]$ and 0, respectively. Following the above procedure, the genie width $w[i]$ and the genie time correlation matrix $\mathbf{R}_{\text{time}}[i]$ can be calculated. We denote by $\mathbf{R}_{\text{t,robust}}[i] = \mathbf{R}_{\text{t,robust}}(w[i])$ the genie robust time correlation matrix.

Once we have $\hat{\mathbf{R}}_{\text{f,robust}}[i]$ and $\hat{\mathbf{R}}_{\text{t,robust}}[i]$, the robust channel estimation is performed as follows: For Rx antenna i and Tx antenna j ,

$$\begin{aligned} & [\hat{\mathbf{h}}_f[i, j, l]]_{l \in S} \\ &= \hat{\mathbf{R}}_{\text{robust}}[i] \left(\frac{\hat{\sigma}_i^2}{\hat{P}} \mathbf{I} + \hat{\mathbf{R}}_{\text{robust}}[i] \right)^{-1} [\tilde{\mathbf{h}}_f[i, j, l]]_{l \in S}, \end{aligned}$$

where $\mathbf{R}_{\text{robust}}[i] = \mathbf{R}_{\text{t,robust}}[i] \otimes \mathbf{R}_{\text{f,robust}}[i]$, \otimes is the Kronecker product, and $[\tilde{\mathbf{h}}_f(i, j, l)]_{l \in S}$ denotes the concatenation of CFRs $\tilde{\mathbf{h}}_f[i, j, l]$ for all l in S .

3.1.5. PRECODER SELECTION

We denote by $\hat{\mathbf{H}}[m, l]$ the estimated channel matrix of subcarrier group m and symbol l . The average whitened spatial covariance matrix is given by

$$\hat{\mathbf{C}}_S = \frac{1}{B|S|} \sum_{m=1}^B \sum_{l \in S} \hat{\mathbf{H}}^H[m, l] \hat{\mathbf{C}}_n^{-1} \hat{\mathbf{H}}[m, l].$$

For each rank R , the precoder $\mathbf{W}_{(R)}$ is selected as

$$\hat{\mathbf{W}}_{(R)} = \arg \max_{\mathbf{W} \in \mathcal{W}_R} \log \det \left(\mathbf{I}_R + \mathbf{W}^H \hat{\mathbf{C}}_S \mathbf{W} \right).$$

3.1.6. RANK SELECTION

The reported rank \hat{R} is selected as

$$\hat{R} = \arg \max_{R \in \{1, \dots, N_R\}} \log \det \left(\mathbf{I}_R + \hat{\mathbf{W}}_{(R)}^H \hat{\mathbf{C}}_S \hat{\mathbf{W}}_{(R)} \right).$$

3.2. Dataset

We use the open-source SIONNA package (Hoydis et al., 2022), together with the system model described in this section, to generate the dataset. At each time step n , corresponding to a slot, we record a selected set of diverse features to conduct proof-of-concept experiments. The features are listed in Table 1.

Table 1. Selected features

Feature	Description
Channel type (categorical)	The type of SIONNA channel (UMi, UMa, RMa)
$K \in \mathbb{Z}^+$	Number of subcarriers per OFDM symbol
$\hat{\mathbf{C}}_n \in \mathbb{C}^{N_R \times N_R}$	Estimated noise covariance matrix
$\hat{\mathbf{R}}_{\text{f,robust}} \in \mathbb{C}^{B \times B}$	Estimated frequency correlation matrix
$\hat{\mathbf{C}}_{\text{time}} \in \mathbb{C}^{ S \times S }$	Estimated time covariance matrix
$\hat{\mathbf{R}}_{\text{time}} \in \mathbb{C}^{ S \times S }$	Estimated time correlation matrix
$\hat{\mu}, \hat{\ell} \in \mathbb{R}$	Estimated center and length of delay profile
$\hat{w} \in \mathbb{R}^+$	Estimated width of Doppler spectrum
$\hat{\mathbf{W}}_{(R)} \in \mathbb{C}^{N_T \times R}$	Selected precoder of rank R
$\hat{R} \in \mathbb{Z}^+$	Selected transmission rank
$\hat{G} \in \mathbb{R}^+$	Estimated Spectral Efficiency

The symbol duration and cyclic prefix length follow the fifth generation (5G) New Radio (NR) numerology (3GPP, 2018). Specifically, we use a normal cyclic prefix. In NR, a slot consists of 14 OFDM symbols, hence $L = 14$. The configurations used to generate the dataset is listed in Table 2, wherein SNR is randomly drawn following uniform distribution. To reduce implementation complexity, the parameters related to $(\hat{w}, \hat{\mu}, \hat{\ell})$ are selected from predefined candidate sets.

Table 2. Dataset configuration

Parameter	Values
Channel type	UMi, UMa, RMa
(Center frequency, subcarrier spacing) (Hz)	(2.6G, 15k), (3.5G, 30k)
SNR (dB)	[0, 30]
Receiver speed (km/h)	3, 10, 30, 60, 90
(N_T, N_R)	(4, 1), (4, 2), (4, 4), (8, 1), (8, 2), (8, 4)
Number of subcarrier group (B)	25, 50, 75, 100
Size of subcarrier group (M)	12, 24, 48
Set of pilot symbols (S)	{2, 8}, {2, 6, 10}, {4, 8, 12}, {2, 5, 8, 11}

Considering all possible combinations for the values in Table 2, there is a total of 86403 unique simulation settings. As generating data for all configurations would be intractable, we randomly select a subset of 29000 configurations. For each selected configuration we run the simulation 8 times, each time with a different SNR (sampled randomly with uniform distribution from the range specified in the above

table) for 100 slots. Each simulation uses a unique random seed. We then divide the obtained data into 5-slot sequences, each one representing one datapoint to be used as input for our model.

4. Methodology

The setting of communication systems presents several challenges that hinder the development of a foundation model. These systems are inherently complex and involve a large number of interdependent variables. Factors such as the choice of modulation, channel conditions (e.g., signal-to-noise ratio, multipath propagation, Doppler shift, fading), and the internal states of components like transmitters and receivers all influence system behavior. Moreover, the data in communication systems is highly heterogeneous. Some variables vary across consecutive transmissions over the same channel, while others remain constant for a given channel but change across different channels. Variables can be categorical, scalar, vector, or matrix-valued; their dimensionality may vary with system configuration, and they may belong to either the real or complex domain.

These characteristics require a foundation model for communication systems to reason over and capture relationships among a large number of diverse and heterogeneous variables. In the remainder of this section, we present our methodology for training such a model, covering data pre-processing, model design, and pre-training strategy.

Pre-processing Two important aspects must be considered. The first is that for multi-dimensional features, the model must handle varying dimensionality depending on the current system settings. Since neural networks require fixed input sizes, we identify the maximum dimensionality for each feature and apply zero-padding to those with smaller dimensions. Padding masks are then introduced in the neural network computations to ensure that the model ignores the padded values.

The second aspect is that different features have different magnitudes, which can further vary significantly depending on the system configuration. We adopt a normalization strategy tailored to the properties of each feature. Scalar features are normalized using global statistics—specifically, we compute the mean and standard deviation over the training set and use them for normalization. Matrix and vector features are handled individually: we either normalize them by computing statistics across the input sequence or leave them unchanged (e.g., correlation matrices, which are already normalized). For the features that require padding, the normalization statistics are computed on the unpadded elements (and the normalization is applied only on those elements).

Tokenization Given the heterogeneity of the features in our data, we divide them into four categories: scalars, vectors, matrices, and categorical. We design a tokenization method for each category as follows:

- **Scalars:** we use Fourier encodings: $\text{Tok}(x) = \left[\cos \frac{2\pi x}{\lambda_i}, \sin \frac{2\pi x}{\lambda_i} \right]$ for $0 \leq i < D/2$, where the λ_i are logarithmically spaced values taken from an interval $[\lambda_{\min}, \lambda_{\max}]$, chosen based on the scale and variability of the feature.
- **Vectors:** given a vector $x \in \mathbb{R}^d$, we apply a learnable linear transformation Wx , where $W \in \mathbb{R}^{D \times d}$.
- **Matrices:** following the approach from ViT (Dosovitskiy et al., 2021), we divide a matrix $X \in \mathbb{R}^{d_1 \times d_2}$ into “patches”: $\{X_j^{(p)} \in \mathbb{R}^{p_1 \times p_2}, j = 1, \dots, z\}$, where p_1 and p_2 are patch sizes, and $z = (d_1/p_1) \cdot (d_2/p_2)$ is the number of patches. Each patch is then flattened into a vector of size $p = p_1 \cdot p_2$ and transformed using a linear map $W \in \mathbb{R}^{D \times p}$. A matrix is thus converted into z tokens, providing finer granularity to encode its structure. We define patch sizes using a simple heuristic: we choose the size such that the number of tokens per matrix is at most 64, with a minimum patch size of 8×8 . Alternatively, the patch size could be treated as a tunable hyperparameter, but we avoid this extra cost as our heuristic performs well in our experiments.
- **Categorical:** we assign a learnable vector $V \in \mathbb{R}^D$ to each category.

To handle complex-valued features, we first normalize and pad them as needed, and then convert them into two-channel vectors or matrices by stacking the real and imaginary parts.

Transformer Model For the base architecture of the transformer, we follow the Llama model (Grattafiori et al., 2024), as adopted by several state-of-the-art open-source models.

Typically, transformer-based foundation models receive a single quantity (and thus a single token) per each element of the input sequence. In our case, however, we have multiple features—and hence multiple tokens—at each slot in the input sequence. We assign the same positional embedding to the tokens for all features within the same slot. This setup, however, means that the model cannot, by default, distinguish tokens corresponding to the same feature across different slots, as it has no information about which token, within those for a given slot, is related to which feature. To address this, we introduce a *feature embedding*: a learnable vector associated with each input feature, which is concatenated to the corresponding token to let the model recognize instances of the same feature across slots.

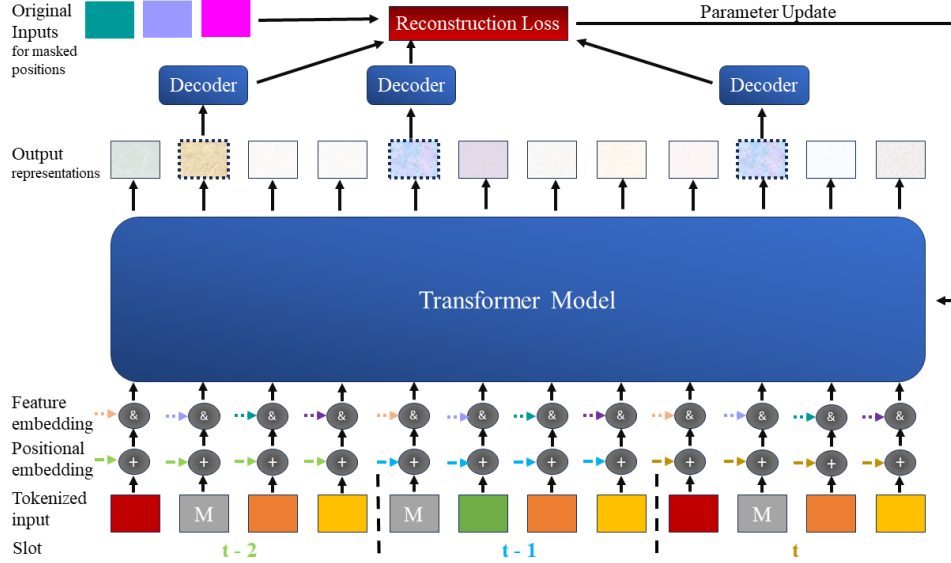


Figure 1. Scheme of our foundation model for communications data and the pre-training procedure. Tokens for each feature are first added to a positional embedding; features at the same input slot receive the same positional embedding (indicated by the color of the arrows). A feature embedding is then concatenated (&); the tokens for the same feature at different slots receive the same feature embedding. The transformer processes these tokens and outputs a representation for each input. During pre-training, a subset of input features is masked (denoted as “M”). The output representations corresponding to the masked inputs are decoded into the original feature space, and a reconstruction loss is computed. The model is trained to minimize this loss.

Once the tokens have been processed by the transformer, we obtain an embedding vector for each input token. To compute the loss for a given feature, we need to decode the embedding vector(s) back into the original space of that feature. We design the decoding mechanism to follow an “inverse” process to the one used for tokenization:

- **Scalars:** a learnable vector $U \in \mathbb{R}^{1 \times D}$ maps the hidden representation to a scalar
- **Vectors:** a learnable matrix $Q \in \mathbb{R}^{d \times D}$ maps the hidden representation to the original dimensionality
- **Matrices:** the token for each patch is mapped back to its original dimensionality with a learnable linear transformation $E \in \mathbb{R}^{p_1 \cdot p_2 \times D}$. The patches are then reshaped into a matrix in $\mathbb{R}^{p_1 \times p_2}$ and concatenated back into the original shape for the feature
- **Categorical:** a learnable matrix $R \in \mathbb{R}^{n_c \times D}$ maps the hidden representation to a 1-hot encoding of the category

The learnable components of the decoding mechanism are intentionally kept simple to encourage the model to capture complexity within the representations themselves, rather than in the decoding functions.

Pre-Training Procedure Our goal is to enable the foundation model to learn relationships between features. To this end, we adopt a self-supervised approach (requiring no annotations) based on masked token prediction, similar to the pretraining strategy used in language models such as BERT (Devlin et al., 2019). While training using *random masking* is common in language models, communication data requires more care. At each position in the input sequence, multiple features are present, and some features are not predictable from others. Naively masking such features can introduce instability during training.

To address this, we identify a subset of five *target* features from Table 1—namely, the transmission rank, selected precoder, Doppler spectrum, center and length of the delay profile. These are features that can be estimated from the remaining ones, as discussed in Section 3. Pretraining is then performed by randomly masking a subset of the target features at each slot in the input sequence and training the model to predict them. An overview of the model and pre-training procedure is provided in Figure 1.

5. Experiments

The evaluation of our foundation model focuses on demonstrating its effectiveness in understanding communication systems data and its ability to perform estimation tasks. Specifically, we consider two scenarios:

1. **Forecasting.** The model is tasked with estimating the value of a given feature in the next slot.
2. **Interpolation.** We randomly mask the values of a feature at certain slots in the input sequence, and the model is tasked with estimating the values of the masked feature. In more detail, we perform one evaluation run for each feature. During each run, for each batch, a slot of the input sequence is selected randomly (with uniform distribution), and the selected feature is masked for that slot.

5.1. Experimental Setting

Data. We generate 1 million datapoints using the simulator described in Section 3.2. The data is then randomly split into training, validation, and test sets, containing 80%, 10%, and 10% of the data, respectively.

Foundation Model Evaluation Procedure. As described above, we evaluate the model by measuring its ability to estimate missing values for features of interest through the tasks of forecasting and interpolation. The features used in these tasks are the target features defined in Section 4: transmission rank, selected precoder, Doppler spread, center and length of the delay profile.

Hyperparameters. We use the validation set to tune the hyperparameters of our model. Specifically, we perform a grid search over learning rate, learning rate scheduling, token dimension, number of layers, and number of attention heads. We provide the values used for the hyperparameter tuning procedure in Table 3.

Table 3. Values for hyperparameter tuning procedure.

Feature	Values
Learning Rate	$10^{-4}, 10^{-5}, 10^{-6}$
Learning Rate Schedule	None, Step, Cosine Annealing
Token Dimension	128, 256, 512
Number of Layers	8, 12, 16
Number of Attention Heads	4, 8, 16, 32

The final hyperparameter values used for our experiments are shown in Table 4. During pre-training, we randomly (with uniform distribution) mask one target feature at each input slot. We use a batch size of 256, a learning rate of 10^{-6} , and no learning rate schedule for all model and dataset sizes.

Computational Resources. For data generation we found that it takes 0.5 seconds to generate a single slot. We parallelize the data generation over multiple Intel Xeon Platinum processing nodes. For training our largest model, we use $2 \times$ NVIDIA A6000 GPUs.

Table 4. Parameters and dimensions for different model sizes used in our experiments.

Model	#layers	#head	token dim	hidden dim
5M	6	6	387	1548
30M	32	18	387	1548
100M	75	24	429	1716

5.2. Results

Table 5 presents the results for the *forecasting* and *interpolation* tasks. The model achieves consistently low mean squared error (MSE) across all five target features, demonstrating its ability to capture dependencies in the communication data. Since all features are normalized prior to training, the reported MSE values are computed on standardized scales and can be interpreted as percentage errors. As expected, interpolation performs slightly better than forecasting, as it estimates missing features using other available features within the same slot, while forecasting requires predicting feature values at future slots.

Table 5. Estimation Error (MSE) for the **forecasting** and **interpolation** tasks on five key features: center of the delay profile ($\hat{\mu}$), length of the delay profile ($\hat{\ell}$), Doppler spectrum width (\hat{w}), transmission rank (\hat{R}), and selected precoder ($\hat{\mathbf{W}}$).

	$\hat{\mu}$	$\hat{\ell}$	\hat{w}	\hat{R}	$\hat{\mathbf{W}}$
Forecasting	0.019	0.021	0.077	0.129	0.101
Interpolation	0.019	0.020	0.054	0.128	0.100




Figure 2. Scaling behaviour vs. data generation compute. Each curve shows the test loss for a model of a given size (5M, 30M, or 100M parameters) as a function of compute used for data generation.

Among the features, Doppler spectrum width (\hat{w}) and delay profile parameters ($\hat{\mu}$, $\hat{\ell}$), which are all scalar-valued, are estimated with the highest accuracy. In contrast, transmission rank (\hat{R}) and especially the selected precoder ($\hat{\mathbf{W}}$)—a matrix-valued feature with variable dimensionality—exhibit higher errors. This reflects the increased difficulty of modeling high-dimensional, structured, and variable-size features. These observations are consistent with the challenges discussed in Section IV and highlight the model’s ability to

generalize across diverse modalities in communication systems data.

5.3. Scaling Behaviour.

We carried out experiments to study how model performance scales with compute, measured in floating point operations (FLOPs) allocated to data generation and training. Specifically, we trained foundation models of three sizes (5M, 30M, and 100M parameters) on datasets of varying sizes (10^6 , $5 \cdot 10^6$, and 10^7 training examples). These datasets are obtained by randomly subsampling the training set used for the estimation experiments. Each configuration was trained for up to 20 epochs.

Figure 2 shows how the test loss (sum of the losses for the target features) varies with compute allocated to data generation, which determines dataset size. Figure 3 shows test loss as a function of training compute, which controls the number of model updates. Performance improves with both model and dataset size, especially when scaled together. These trends are consistent with known scaling behaviour observed in other domains. Larger models benefit more from extended training. The results highlight the importance of co-scaling model size, data, and training budget to fully realize performance gains.




Figure 3. Scaling behaviour vs. training compute. Each curve shows the test loss for models of a given size (5M, 30M, 100M) as a function of compute used for training.

6. Conclusion

In this work, we laid the groundwork for a foundation model for communication systems. We discussed and addressed the unique challenges involved in designing such a model, developed a simulation system to generate training data, and demonstrated the model’s ability to perform forecasting and interpolation across several key features.

Future work will focus on scaling up both the dataset and model size, and incorporating additional features. We will release our data generation pipeline and model upon acceptance of this paper, with the hope that this work serves as a first step in this paradigm shift. We aim to support the

community’s progress toward building datasets, models, and benchmarks that advance the development of effective foundation models for communication systems, with potential applications in 6G.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- 3GPP. *NR; Physical Channels and Modulation*, 3GPP Standard TS 38.211, 2018.
- Abbaspourazad, S. et al. Large-scale training of foundation models for wearable biosignals. In *International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=pC3WJHf51j>.
- Alayrac, J.-B. et al. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=EbMuimAbPbs>.
- Bodnar, C. et al. A foundation model for the earth system. *arXiv*, 2024. URL <https://arxiv.org/abs/2405.13063>. *arxiv:2405.13063*.
- Bommasani, R. et al. On the opportunities and risks of foundation models. *arxiv*, 2021. URL <https://arxiv.org/abs/2108.07258>. *arxiv:2108.07258*.
- Bornea, A.-L., Ayed, F., Domenico, A. D., Piovesan, N., and Maatouk, A. Telco-rag: Navigating the challenges of retrieval-augmented language models for telecommunications. *arxiv*, 2024. URL <https://arxiv.org/abs/2404.15939>. *arxiv:2404.15939*.
- Chen, Z. et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arxiv*, 2025. URL <https://arxiv.org/abs/2412.05271>. *arxiv:2412.05271*.
- Dai, W. et al. Nvlm: Open frontier-class multimodal LLMs. *arxiv*, 2024. URL <https://arxiv.org/abs/2409.11402>. *arxiv:2409.11402*.
- Das, A., Kong, W., Sen, R., and Zhou, Y. A decoder-only foundation model for time-series forecasting. *arxiv*, 2024. URL <https://arxiv.org/abs/2310.10688>. *arxiv:2310.10688*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Conference of the*

- North American Chapter of the Association for Computational Linguistics: *Human Language Technologies*, pp. 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423/>.
- Dosovitskiy, A. et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Ekambaram, V. et al. Tiny time mixers (TTMs): Fast pre-trained models for enhanced zero/few-shot forecasting of multivariate time series. In *Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=3O5YCEWETq>.
- Goswami, M., Szafer, K., Choudhry, A., Cai, Y., Li, S., and Dubrawski, A. Moment: A family of open time-series foundation models. In *International Conference on Machine Learning*, 2024.
- Grattafiori, A. et al. The Llama 3 herd of models. *arxiv*, 2024. URL <https://arxiv.org/abs/2407.21783>. *arxiv:2407.21783*.
- Hamdan, M. Q. et al. Recent advances in machine learning for network automation in the o-ran. *Sensors*, 23 (21), 2023. ISSN 1424-8220. doi: 10.3390/s23218792. URL <https://www.mdpi.com/1424-8220/23/21/8792>.
- Hoydis, J. et al. Sionna: An open-source library for next-generation physical layer research. *arxiv*, 2022. URL <https://arxiv.org/abs/2203.11854>. *arxiv:2203.11854*.
- Lewis, P. et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS ’20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- Li, Y., Cimini, L., and Sollenberger, N. Robust channel estimation for ofdm systems with rapid dispersive fading channels. *IEEE Transactions on Communications*, 46(7): 902–915, 1998. doi: 10.1109/26.701317.
- Liao, F., Wei, S., and Zou, S. Deep learning methods in communication systems: A review. *Journal of Physics: Conference Series*, 1617(1):012024, aug 2020. doi: 10.1088/1742-6596/1617/1/012024. URL <https://dx.doi.org/10.1088/1742-6596/1617/1/012024>.
- Liu, X. et al. Moirai-moe: Empowering time series foundation models with sparse mixture of experts. *arxiv*, 2024. URL <https://arxiv.org/abs/2410.10469>. *arxiv:2410.10469*.
- Liu, Y. et al. Sundial: A family of highly capable time series foundation models. *arxiv*, 2025. URL <https://arxiv.org/abs/2502.00816>. *arxiv:2502.00816*.
- Narayanswamy, G. et al. Scaling wearable foundation models. In *International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=yb4QE6b22f>.
- Nikbakht, R., Benzaghta, M., and Geraci, G. Tspec-llm: An open-source dataset for llm understanding of 3gpp specifications. *arxiv*, 2024. URL <https://arxiv.org/abs/2406.01768>. *arxiv:2406.01768*.
- OpenAI. Openai o1 system card, 2024. <https://openai.com/index/openai-o1-system-card/>.
- Piovesan, N., Domenico, A. D., and Ayed, F. Telecom language models: Must they be large? *arxiv*, 2024. URL <https://arxiv.org/abs/2403.04666>. *arxiv:2403.04666*.
- Rasul, K. et al. Lag-Llama: Towards foundation models for probabilistic time series forecasting. *arxiv*, 2024. *arxiv:2310.08278*.
- Saad, W. et al. Artificial general intelligence (agi)-native wireless systems: A journey beyond 6g. *arxiv*, 2024. URL <https://arxiv.org/abs/2405.02336>. *arxiv:2405.02336*.
- Shao, J. et al. Wirelessllm: Empowering large language models towards wireless intelligence. *Journal of Communications and Information Networks*, 9(2):99–112, 2024. doi: 10.23919/JCIN.2024.10582827.
- Shi, X., Wang, S., Nie, Y., Li, D., Ye, Z., Wen, Q., and Jin, M. Time-moe: Billion-scale time series foundation models with mixture of experts. In *International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=elwDDFmlVu>.
- Sun, Y., Lee, H., and Simpson, O. Machine learning in communication systems and networks. *Sensors*, 24 (6), 2024. ISSN 1424-8220. doi: 10.3390/s24061925. URL <https://www.mdpi.com/1424-8220/24/6/1925>.
- Team, G. et al. Gemini: A family of highly capable multimodal models. *arxiv*, 2024. URL <https://arxiv.org/abs/2312.11805>. *arxiv:2312.11805*.
- Thakkar, K., Goyal, A., and Bhattacharyya, B. Emergence of deep learning as a potential solution for detection, recovery and de-noising of signals in communication systems. *International Journal of Intelligent Networks*, 1:119–127, 2020. ISSN 2666-6030. doi: <https://doi.org/10.1016/j.ijin.2020.12>.

001. URL <https://www.sciencedirect.com/science/article/pii/S2666603020300270>.

Tse, D. and Viswanath, P. *Fundamentals of wireless communication*. Cambridge University Press, USA, 2005. ISBN 0521845270.

Wang, S. and Li, G. Y. Machine learning in communications: A road to intelligent transmission and processing. *arxiv*, 2024. URL <https://arxiv.org/abs/2407.11595>. *arxiv:2407.11595*.

Wang, Y. et al. Transformer-empowered 6g intelligent networks: From massive mimo processing to semantic communication. *Wireless Communications*, 30(6):127–135, December 2023. ISSN 1536-1284. doi: 10.1109/MWC.008.2200157. URL <https://doi.org/10.1109/MWC.008.2200157>.

Yu, W., Sohrabi, F., and Jiang, T. Role of deep learning in wireless communications. *IEEE BITS the Information Theory Magazine*, 2(2):56–72, 2022. doi: 10.1109/MBITS.2022.3212978.

Žeger, I. and Šišul, G. Introduction to deep learning possibilities in communication systems. In *International Symposium ELMAR*, pp. 21–24, 2021. doi: 10.1109/ELMAR52657.2021.9550825.