

LiDAR Aided Future Beam Prediction in Real-World Millimeter Wave V2I Communications

Shuaifeng Jiang, Gouranga Charan, and Ahmed Alkhateeb

Abstract—This paper presents the first large-scale real-world evaluation for using LiDAR data to guide the mmWave beam prediction task. A machine learning (ML) model that leverages LiDAR sensory data to predict the current and future beams was developed. Based on the large-scale real-world dataset, DeepSense 6G, this model was evaluated in a vehicle-to-infrastructure communication scenario with highly-mobile vehicles. The experimental results show that the developed LiDAR-aided beam prediction and tracking model can predict the optimal beam in 95% of the cases and with around 90% reduction in the beam training overhead. The LiDAR-aided beam tracking achieves comparable accuracy performance to a baseline solution that has perfect knowledge of the previous optimal beams, without requiring any knowledge about the previous optimal beam information and without any need for beam calibration. This highlights a promising solution for the critical beam alignment challenges in mmWave and terahertz communication systems.

Index Terms—beam tracking, LiDAR, machine learning, DeepSense 6G, real-world data.

I. INTRODUCTION

Harvesting the millimeter wave (mmWave) and terahertz (THz) data rate gains requires deploying large antenna arrays and using narrow beams [1]. Adjusting the optimal narrow beams, however, often requires large beam training overhead, which makes it hard for these systems to support highly mobile applications. This motivates developing novel approaches that can reduce the beam training overhead. An important observation here is that due to the highly directional nature of the mmWave/THz communication, the beam management problem in these systems implicitly relies on the positions of the base station (BS) and the user equipment (UE) and the geometry of the surrounding environment. Based on that, the sensing information about the user location and the environment geometry could potentially be utilized to guide the beam management and reduce beam training overhead [3]–[6].

Prior work has studied improving the mmWave/THz beam selection and blockage prediction based on a variety of sensing modalities [3]–[6]. One conventional direction is using wireless sensing information. In [3], the authors proposed to leverage wireless received signals as a unique signature for the UE position and its interaction with the surrounding environment. [4] exploits the sub-6 GHz channel, which is easier to obtain compared to the mmWave channel, to predict the mmWave beam and blockage status. Another direction gaining increasing interest is exploiting the vision/camera

information. [5] leverages the camera signals to proactively predict dynamic link blockages for mmWave systems. In [6], the authors proposed to use the radar signals to sense the UE to predict the 6G beam.

However, each sensing modality mentioned above has its own drawbacks. The wireless received signal and the radar sensing signals occupy wireless resources. The performance of the camera/vision-based sensing degrades under poor lighting conditions, and it leads to privacy concerns. Extra signaling overhead is needed for the BS to obtain position information of the UEs, and the accurate position information can also cause privacy problems. This motivates us to investigate the other alternative: the light detection and ranging (LiDAR) sensors. In this paper, we propose to achieve mmWave/THz beam prediction and beam tracking relying on the LiDAR sensing data. Our contribution is summarized as follows.

- We propose to leverage LiDAR sensory data to aid the mmWave beam prediction and tracking tasks. While using LiDAR to predict current beams has been previously investigated in [7] and [8], this is the first paper that leverages LiDAR data to also predict *future* beams.
- We propose an efficient machine learning (ML) model for the LiDAR-aided beam prediction tasks.
- To evaluate the performance of the proposed LiDAR beam prediction and tracking approach, this paper presents the first large-scale real-world evaluation results based on a vehicle-to-infrastructure communication scenario of the DeepSense 6G dataset.

Evaluation results show that the proposed LiDAR beam prediction and tracking approach can achieve high accuracy performance: it can predict the optimal beam in 95% of the cases with beam training overhead reduced from 64 to 5. This demonstrates the potential of LiDAR sensing information in real-world mmWave/THz communication systems.

II. SYSTEM MODEL AND PROBLEM FORMULATION

In this section, we first introduce the adopted system model. Then, we define the two considered beam management tasks, namely LiDAR-aided current, and future beam prediction.

A. System Model

As shown in Fig. 1, we consider a communication system model where a base station (BS) serves a single mobile user equipment (UE) over a mmWave frequency band. Extensions to multi-user scenarios is an important future research direction. The BS has N antennas, and is also equipped with a LiDAR sensor to aid the communication. For simplicity, the




Fig. 1. This figure illustrates the considered system model: the BS senses the environment and the moving UE with a LiDAR. The obtained sensing information is then utilized for the BS beam prediction and beam tracking. The figure also shows the block diagram of the proposed ML models. The ML model consists of the embedding block, the gated recurrent unit (GRU) feature extractor block, and the classifier block.

UE is assumed to be single-antenna. Adopting a narrowband block-fading channel model, the channel vector between the BS and the UE at time step t , $\mathbf{h}[t] \in \mathbb{C}^{N \times 1}$, can be written as

$$\mathbf{h}[t] = \sum_{l=1}^L \alpha_l \mathbf{a}(\theta_l, \phi_l), \quad (1)$$

where L is the total number of paths. α_l , θ_l , and ϕ_l represent the complex gain, azimuth angle, and elevation angle of the l -th path. $\mathbf{a}(\theta_l, \phi_l)$ is the array response vector of the two angles. The L , α_l , θ_l , and ϕ_l are all functions of t due to the UE and environment dynamics. The channel coherence time T_c is assumed to be larger than the duration of each time step.

In the downlink transmission, if the BS sends a complex symbol $s \in \mathbb{C}$, the received signal at the UE can be written as

$$y[t] = \mathbf{h}^H[t] \mathbf{f}[t] s + n[t], \quad (2)$$

where $\mathbf{f}[t] \in \mathbb{C}^{N \times 1}$ is the transmit beamforming vector the BS used at time step t , and n is the receive additive white Gaussian noise, satisfying $\mathbb{E}[n[t]n^H[t]] = \sigma_n^2$. The downlink signal s satisfies $\mathbb{E}[ss^H] = P$ with P denoting the transmit power. The BS is assumed to adopt a pre-defined beamforming codebook $\mathcal{F} = \{\mathbf{f}_1, \dots, \mathbf{f}_M\}$ of size M , i.e., $\mathbf{f}[t] \in \mathcal{F}$. Since analog beamsteering is considered, the beamforming vectors in $\mathcal{F}$ satisfy the constant modulus constraint and $\|\mathbf{f}\|_2^2 = 1$.

Under the pre-defined codebook constraint, the beamforming vector $\mathbf{f}[t] \in \mathcal{F}$ is uniquely indicated by the beam index $p[t] \in \{1, \dots, M\}$. The optimal beam index $p^*[t]$ can be obtained by maximizing the receive SNR:

$$p^*[t] = \arg \max_{p[t] \in \{1, \dots, M\}} |\mathbf{h}^H[t] \mathbf{f}_{p[t]}|^2. \quad (3)$$

B. Problem Formulation and ML Task Definition

In this paper, we investigate how to leverage the LiDAR sensory information to achieve beam prediction (predicting the current beams) and beam tracking (predicting the future beams) with reduced beam training overhead. For that, we define $\mathbf{x}[t] \in \mathbb{R}^{D \times 1}$ as the LiDAR sensing signal captured by the BS at time step t , where D is the number of quantized angles in the LiDAR field-of-view and each entry of $\mathbf{x}[t]$ represents the LiDAR-estimated distance to the obstacle at the corresponding angular direction. Furthermore, the sequence of observed LiDAR information can be written as $\mathcal{X}_{t,W} = \{\mathbf{x}[t-W+1], \dots, \mathbf{x}[t]\}$, where $W \in \mathbb{Z}^+$ is an observation

window. The objective is then to find a mapping function that predicts the current/future optimal beam index from the $\mathcal{X}_{t,W}$. In this paper, we focus on the use of ML models to learn this mapping function. Next, we formally define the beam prediction and tracking problems.

LiDAR Beam Prediction: The beam prediction task is defined as predicting the *current* optimal beam based on the LiDAR data captured until the *current* time step. If $f_{pr}(\cdot)$ denotes the model that takes the current LiDAR sensory information as the input, then the beam prediction optimization problem can be written as

$$f_{pr}^*(\mathcal{X}_{t,W}; \Theta_{pr}^*) = \arg \max_{f_{pr}(\cdot), \Theta_{pr}} \mathbb{P}\{f_{pr}(\mathcal{X}_{t,W}; \Theta_{pr}) = p^*[t]\},$$

where $f_{pr}(\cdot)$ denotes the ML model mapping function for the beam prediction task, and Θ_{pr} is its training parameters.

LiDAR Beam Tracking: The beam tracking targets predicting the optimal *future* beams using the LiDAR data captured until the current time step. If $f_{tr}(\cdot)$ denotes the model that takes the previous LiDAR sensory information as the input, then the beam tracking optimization problem predicting the beam for the v -th future time step can be written as

$$f_{tr}^*(\mathcal{X}_{t,W}, v; \Theta_{tr}^*) = \arg \max_{f_{tr}(\cdot), \Theta_{tr}} \mathbb{P}\{f_{tr}(\mathcal{X}_{t,W}, v; \Theta_{tr}) = p^*[t+v]\}, \quad (4)$$

where Θ_{tr}^* is the associated trainable parameters, and $v \in \{1, \dots, V\}$ denotes the lead-time (future instance) of the beam tracking process. Next, we present a baseline for the beam tracking problem which relies only on wireless data.

Baseline Beam Tracking: The sequence of the previously adopted beams contains information about the UE moving pattern, which can be exploited to predict the future optimal beams. In this paper, we employ this baseline ML model which predicts the future beams using the previously used beams. Similar to (4), this optimization problem is defined as

$$f_b^*(\mathcal{F}_{t,W}, v; \Theta_b^*) = \arg \max_{f_b(\cdot), \Theta_b} \mathbb{P}\{f_b(\mathcal{F}_{t,W}, v; \Theta_b) = p^*[t+v]\},$$

where $\mathcal{F}_{t,W} = \{\mathbf{f}[t-W+1], \dots, \mathbf{f}[t]\}$ denotes the beam sequence used at time steps $(t-W+1), \dots, t$. Similar to the two LiDAR beam management tasks, $f_b^*(\cdot)$ and Θ_b^* are the optimal ML model and trainable parameters associated to this baseline beam tracking task.

In Section III, we explain the proposed ML models for the mmWave beam prediction and beam tracking tasks.

III. PROPOSED LiDAR AIDED BEAM MANAGEMENT SOLUTION

First, we present why LiDAR sensor is promising for aiding the mmWave beam management tasks. Then, we explain the system operation for LiDAR-aided communications. Last, we describe the proposed ML model for the LiDAR beam prediction/tracking tasks and the baseline (beam history based) solution.

A. Key Idea: Why LiDAR?

LiDARs are often seen installed in vehicles to enable autonomous driving. Although the use of LiDAR for wireless communications has not been commonly investigated, LiDAR can be a promising sensing option for mmWave BSs. First, the BS is more stationary than vehicles, and the height of the BS is also usually higher. Therefore, installing LiDARs at the BSs can have better observing angles and produce more accurate measurements. Second, since the BSs are less power-constrained compared to vehicles, the power consumption of LiDARs could be more affordable for BSs. The main limitation of LiDARs is the inaccuracy of measuring distance under some weather conditions, such as heavy rain, snow, and fog. In these cases, complementary sensors can be adopted to enhance the sensing capability at the BSs.

B. Proposed LiDAR-Aided System Operation

We propose a LiDAR-aided mmWave communication system that operates as follows: In each time step (*e.g.*, coherence time [9]), the BS captures a LiDAR image using its sensor as shown in Fig. 1. A sequence of L captured LiDAR sensing images is then presented to an ML model, which predicts the top- k promising beams that should be used to serve the user. Given the top- k predicted beams, the BS can either (i) directly adopt the top-1 beam and completely eliminate the beam training overhead or (ii) perform an over-the-air beam refinement/training using only the predicted subset of beams (the top- k beams). Next, we describe the proposed ML model.

C. Deep Learning Model

To solve the beam prediction/tracking tasks, we utilize the recurrent neural network (RNN) based encoder-decoder architecture in [10]. The adopted RNN predicts a sequence of current and future optimal beams given an input sequence of previous LiDAR or wireless beam measurements. We adopt the encoder-decoder RNN architecture for the following reasons. (i) Since the optimal beam evolves together with the UE and environment dynamics, useful information for predicting the optimal beams can be extracted from the previously observed sensory data using *sequential modeling*. RNNs have been extensively studied and empirically proven for sequential modeling tasks, *e.g.*, natural language processing and speech recognition. (ii) The beam prediction/tracking task is latency-sensitive due to the short coherence time of the wireless

channels. Compared to other NN architectures for sequential modeling, *e.g.* Transformers [11], RNNs have advantages in computational complexity and inference time.

Fig. 1 shows the block diagram of the proposed beam prediction and tracking ML model. As shown in Fig. 1, the ML model consists of (i) $W - 1$ repeated blocks, each has an embedding and a gated recurrent unit (GRU) [10] feature extractor components, in addition to (ii) $V + 1$ repeated blocks, each has an embedding, a GRU feature extractor, and classifier components. Next, we describe the different components.

Embedding Block: The embedding block transforms the high-dimensional input data into a low-dimensional representation and reduces the trainable parameters of the model. For the LiDAR data embedding, we adopt a fully connected layer that projects the $\mathbf{x}[t] \in \mathbb{R}^{D \times 1}$ into $\tilde{\mathbf{x}}[t] \in \mathbb{R}^{64 \times 1}$. To embed the previous beam indices, we apply a trainable look-up table of M embedding vectors $\{\tilde{\mathbf{p}}_1, \dots, \tilde{\mathbf{p}}_M\}$, where the embedding vector $\tilde{\mathbf{p}}_m \in \mathbb{R}^{64 \times 1}$ corresponds to the beam index m .

RNN Feature Extractor Block: We employ W single-layer GRUs to extract features from the input LiDAR or beam sequence. In our implementation, the hidden state size of the GRU is 64, and it is initialized with the all-zero vectors. For fair comparisons, we use the same RNN architecture in all three ML beam management tasks.

Classifier Block: Following the RNN feature extractor, a fully connected layer works as the classifier. Since the beam prediction and tracking tasks are classification tasks, we employ the softmax activation function at the output of this fully connected layer. The output of the classifier is a score vector $\hat{\mathbf{s}} = [s_1, \dots, s_M]^T$. The score $s_m \in (0, 1)$ corresponds to the m -th beam $\mathbf{f}_m$ in the codebook. The beam index with the highest score indicates the predicted beam.

Learning Process: The ML model is trained offline in a supervised way. Each data point consists of (i) the input sequence $\{\mathcal{X}_{t,W}, \mathcal{Z}\}$, where $\mathcal{Z}$ is a padding sequence of V all-zero vectors, each of dimension $D \times 1$, and (ii) the label which is the desired output sequence $\{\mathbf{s}^*[t], \dots, \mathbf{s}^*[t+V]\}$, where $\mathbf{s}^*[t+v]$ is the one-hot representation of $p^*[t+v]$.

For the baseline model, each data point consists of (i) the input sequence $\{\mathcal{F}_{t,W}, \mathcal{Z}\}$, and (ii) the desired output sequence $\{\mathbf{s}^*[t], \dots, \mathbf{s}^*[t+V]\}$. To train the ML models, we adopt a cross-entropy loss function applied to the last γ outputs among the output sequence, where γ is a design parameter that is set to 4 in our implementation.

IV. EVALUATION SETUP

In this section, we present the first evaluation of LiDAR aided beam prediction based on a large-scale real-world dataset. In particular, we adopt our DeepSense 6G dataset [12], which comprises co-existing multi-modal sensing and communications data. Next, we describe the adopted DeepSense scenario and the task-specific development dataset.

A. Data Acquisition

We adopt Scenario 8 of the DeepSense 6G dataset. As shown by Fig. 2, the system setup consists of (i) a moving UE working as a transmitter, (ii) a fixed BS working as a receiver,

TABLE I

TOP- k ACCURACY AND RELATIVE RECEIVE POWER OF ML MODELS WITH LiDAR AND BEAM INPUT TRAINED ON 80% AND EVALUATED ON 20% DATA.

Metric	Current beam		Future beam 1		Future beam 2		Future beam 3	
	LiDAR input	Beam input	LiDAR input	Beam input	LiDAR input	Beam input	LiDAR input	Beam input
Top-1 acc.	57.5%	-	55.3%	60.3%	51.9%	56.8%	46.0%	50.9%
Top-5 acc.	95.6%	-	95.0%	97.8%	94.5%	97.0%	93.3%	96.7%
Top-1 pow.	95.9%	-	95.0%	98.1%	93.8%	97.1%	91.8%	96.3%
Top-5 pow.	99.9%	-	99.8%	99.9%	99.7%	99.9%	99.4%	99.9%




Fig. 2. System setup of the Scenario 8 of the DeepSense 6G dataset. The mmWave BS antenna array receives the signal transmitted by the moving UE. A LiDAR is placed in front of the BS to obtain sensing information.

and (iii) a LiDAR sensor with which the BS obtains sensing information. The moving vehicle carries an omni-directional mmWave 60 GHz transmitter to communicate with the BS. The BS is equipped with a 60 GHz receiver that has a 16-element phased array and uses a pre-defined beam codebook of 64 beams. During the data collection process, and at each sampling time t , one LiDAR raw sensing data vector and one 64-element beam training power vector (carrying the receive power values with the 64 beams) are collected.

One raw LiDAR sensing data vector captures a 360° point cloud consisting of 460 samples. Each sample contains an angle and a range. The maximum range and the range resolution of the LiDAR sensor are 30 m and 0.03 m, respectively. The raw LiDAR data vector is preprocessed to obtain $\mathbf{x}[t]$. After preprocessing [13], $\mathbf{x}[t]$ has $D = 216$ angle bins, and the angle resolution is approximately 0.97° . Since the pre-defined beamforming codebook has 64 beams, the angle difference between two neighboring beams is around 2.8° . Therefore, the LiDAR angle resolution is reasonable.

The LiDAR and wireless beam data are synchronized in time. The interval between time steps is 128 millisecond (ms). Due to the 60 GHz carrier frequency and our system setup in Fig. 2, the communication channels are dominated by LoS components. Therefore, the channel coherence time can be approximated using equation (46) in [9]. The derived channel coherence time is around 143 ms under the 0.9 correlation target, which is longer than the time step interval. Furthermore, the DeepSense data collection testbed uses a 10 MHz bandwidth at the 60 GHz band. Therefore, it is reasonable to consider the narrowband channel model in Section II-A.

B. Development Dataset Generation

The DeepSense 6G Scenario 8 data consists of multiple data sequences. Each data sequence contains the data corresponding

to only one pass through the road. We extract the optimal beam indices of each data sequence from the receive powers. After that, each data point in these data sequences is a tuple of LiDAR sensing data and optimal beam index $(\mathbf{x}[t], p^*[t])$. Then, 80% of the data sequences are used for the training dataset, and the remaining data sequences form the test dataset. Note that, with this data split, no overlapping data point (and the corresponding LiDAR information) exists in the training and test datasets so the split is data-leakage free. In the training process, we use an observation window $W = 8$, and train the models to predict the beams up to the third future beam ($V = 3$). Therefore, we further format the training dataset such that each training sample consists of (i) a LiDAR data sequence of eight time steps, (ii) the corresponding eight optimal beam indices, and (iii) three optimal beam indices for the three future time steps we intend to predict.

V. EVALUATION RESULTS

In this section, we evaluate the performance of the proposed LiDAR aided beam prediction/tracking solution compared to the baseline approach. We adopt the top- k accuracy and the top- k relative receive power as the performance metrics. The top- k accuracy is defined as the percentage of the test samples whose ground-truth beam index lies in the predicted beams of top- k scores. The top- k relative receive power captures the ratio between the highest receive power achieved by the top- k predicted beams and the receive power of the optimal beam.

Beam Prediction and Tracking Performance: Table I presents the top- k accuracy and relative power performance of the two proposed approaches for the beam prediction/tracking tasks. It can be seen that the performance of the LiDAR-based solution is comparable to the baseline. The top-5 accuracy of the LiDAR beam prediction is 95.6%. This means that, with the proposed LiDAR beam prediction model, the BS can find the optimal beam in 95.6% of the cases, reducing the beam training overhead from 64 (in the exhaustive search case) to 5.

Although the top-1 accuracy of the LiDAR-based solution is relatively low, its top-1 relative receive power performance is over 95%. This demonstrates the efficacy of the LiDAR-based solution: most of the mispredictions come from the sub-optimal beam with near-optimal receive power.

It is important to mention that the baseline model is assumed to have perfect knowledge of the previous 8 optimal beams, which requires high beam training overhead. As an alternative solution, the baseline model can use the latest predicted beam index as the optimal index. The error in the predicted beam

indices, however, accumulates and the future beam prediction is expected to gradually deviate from the ground-truth optimal beams. This deviation is captured in Fig. 3 which presents the average top- k accuracy for predicting the *first* future beam versus different *operation windows*. The operation window represents the window over which the predicted beams are used as inputs to the baseline model; *i.e.*, instead of the optimal beams, before another exhaustive beam search is done. For example, for an operation window of 5 and an observation window of 8, the approach predicts the first future beam 5 times (in 5 sequential time instances), and then updates its beams by performing exhaustive beam search 8 times (the length of the observation window). It can be seen from Fig. 3 that the LiDAR-aided approach can keep track of the beams because the BS captures the up-to-date LiDAR sensing information at every time step with no cost on the wireless resources. On the other side, the accuracy of the baseline model degrades as the prediction window becomes longer. Without very frequent exhaustive beam training, the proposed LiDAR-aided beam prediction approach consistently outperforms the baseline solution.

Training Overhead: Fig. 3 compares the average beam training overhead required for one time step versus the operation window length when using the top-5 beams. With a prediction window of three time steps, the LiDAR-aided approach only needs around 10% training overhead compared to the baseline solution.

Robustness to Vehicle Dynamics: We evaluate the top-1 relative receive power performance for the future beams predicted by the LiDAR-aided solution with different vehicle speed. The vehicle speed in our dataset ranges from 3 to 18 km/h, and the speed is quantized into four bins {4, 8, 12, 16} km/h in the evaluation. Although the relative receive power slightly decreases as the vehicle velocity increases, over 94% top-1 relative receive power can be achieved for the first future beam across all the vehicle speeds in the dataset.

NN runtime: The runtime of the NN is important for deploying the LiDAR-aided beam prediction and tracking systems in the real world. We tested the runtime using the Intel Xeon Silver 4216 CPU. Since the NN size is relatively small (around 4,000 parameters), the runtime can be small: 1.38 ms on the CPU when the first future beam is predicted, which is reasonable for real-world deployment.

VI. CONCLUSION

This paper develops a machine learning based mmWave beam prediction/tracking approach utilizing LiDAR data and demonstrates its performance on a large-scale real-world dataset for the first time. Based on the DeepSense 6G dataset, the proposed LiDAR-aided beam prediction and tracking model achieves only slightly lower accuracy than a baseline model that has perfect knowledge of the previous optimal beams. The top-5 accuracy of the LiDAR-aided approach is 95.6% and 95.0% for the current beam (beam prediction) and the first future beam (beam tracking), respectively. However, the LiDAR-aided approach only needs 10.4% beam training overhead to match the performance of the baseline approach.




Fig. 3. This figure shows the accuracy performance of the two future beam prediction approaches with different operation windows. The LiDAR-aided solution can keep track of the beam with lower training overhead.

These results highlight the potential of leveraging LiDAR sensory data in real-world mmWave communication systems.

REFERENCES

- [1] R. W. Heath *et al.*, "An overview of signal processing techniques for millimeter wave mimo systems," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 3, pp. 436–453, 2016.
- [2] 3GPP TR 38.802 version 14.1.0 Release 14, "Study on New Radio Access Technology—Physical Layer Aspects," Tech. Rep., Jun 2017.
- [3] A. Alkhateeb *et al.*, "Deep learning coordinated beamforming for highly-mobile millimeter wave systems," *IEEE Access*, vol. 6, pp. 37 328–37 348, 2018.
- [4] M. Alrabeiah and A. Alkhateeb, "Deep learning for mmwave beam and blockage prediction using sub-6 GHz channels," *IEEE Trans. Commun.*, vol. 68, no. 9, pp. 5504–5518, 2020.
- [5] G. Charan, M. Alrabeiah, and A. Alkhateeb, "Vision-aided dynamic blockage prediction for 6G wireless communication networks," in *2021 IEEE ICC Workshops*, pp. 1–6, 2021.
- [6] U. Demirhan and A. Alkhateeb, "Radar aided 6G beam prediction: Deep learning algorithms and real-world demonstration," *WCNC*, pp. 2655–2660, 2022.
- [7] A. Klautau, N. González-Precic, and R. W. Heath, "LiDAR data for deep learning-based mmwave beam-selection," *IEEE Wireless Commun. Lett.*, vol. 8, no. 3, pp. 909–912, 2019.
- [8] M. Zecchin *et al.*, "LiDAR and position-aided mmwave beam selection with non-local CNNs and curriculum training," *IEEE Trans. Veh. Technol.*, vol. 7, no. 3, pp. 2979–2990, 2022.
- [9] V. Va, J. Choi, and R.W. Heath, "The impact of beamwidth on temporal channel variation in vehicular channels and its implications," *IEEE Trans. Veh. Technol.*, vol. 66, no. 6, pp. 5014–5029, 2016.
- [10] K. Cho *et al.*, "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation," in *Proc. of EMNLP*, pp. 1724–1734, 2014.
- [11] A. Vaswani *et al.*, "Attention Is All You Need," *Advances in neural information processing systems*, vol. 66, 2017.
- [12] A. Alkhateeb *et al.*, "DeepSense 6G: A large-scale real-world multi-modal sensing and communication dataset," *available on arXiv*, 2021. [Online]. Available: <https://www.DeepSense6G.net>
- [13] S. Wu *et al.*, "LiDAR-Aided Mobile Blockage Prediction in Real-World Millimeter Wave Systems," *WCNC*, pp. 2631–2636, 2022.