

Camera Based mmWave Beam Prediction: Towards Multi-Candidate Real-World Scenarios

Gouranga Charan , Muhammad Alrabeiah , Tawfik Osman, and Ahmed Alkhateeb 

Abstract—Leveraging sensory information to aid the millimeter-wave (mmWave) and sub-terahertz (sub-THz) beam selection process is attracting increasing interest. This sensory data, captured for example by cameras at the basestations, has the potential of significantly reducing the beam sweeping overhead and enabling highly-mobile applications. The solutions developed so far, however, have mainly considered single-candidate scenarios, i.e., scenarios with a single candidate user in the visual scene, and were evaluated using synthetic datasets. To address these limitations, this paper extensively investigates the sensing-aided beam prediction problem in a real-world multi-object vehicle-to-infrastructure (V2I) scenario and presents a comprehensive machine learning based framework. In particular, this paper proposes to utilize visual and positional data to predict the optimal beam indices as an alternative to the conventional beam sweeping approaches. For this, a novel user (transmitter) identification solution has been developed, a key step in realizing sensing-aided multi-candidate and multi-user beam prediction solutions. The proposed solutions are evaluated on the large-scale real-world DeepSense 6G dataset. Experimental results in realistic V2I communication scenarios indicate that the proposed solutions achieve between 67 – 84% top-1 and close to 100% top-5 beam prediction accuracy for the scenarios with single-user, and between 65 – 80% top-1 and close to 95% top-5 beam prediction accuracy for multi-candidate scenarios. Furthermore, the proposed approach can identify the probable transmitting candidate with more than 93% accuracy across the different scenarios. This highlights a promising approach for significantly reducing the beam training overhead in mmWave/THz communication systems.

Index Terms—Beam prediction, computer vision, deep learning, mmWave communication, multi-user.

I. INTRODUCTION

THE promise that 5G and beyond hold for supporting revolutionary applications (such as autonomous vehicles, intelligent factories, and Internet of Things (IoT)) is contingent on those systems meeting unprecedented performance requirements in terms of achievable rates, latency, and reliability [1],

Received 13 August 2023; revised 1 July 2024; accepted 19 November 2024. Date of publication 3 December 2024; date of current version 18 April 2025. This work was supported by National Science Foundation (NSF) under Grant No. 2048021. The review of this article was coordinated by Dr. Haijun Zhang. (Corresponding author: Ahmed Alkhateeb.)

Gouranga Charan, Tawfik Osman, and Ahmed Alkhateeb are with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: gcharan@asu.edu; tosm@asu.edu; alkhateeb@asu.edu).

Muhammad Alrabeiah is with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287 USA, and also with Electrical Engineering Department, King Saud University, Riyadh 11421, Saudi Arabia (e-mail: malrabei@asu.edu).

Digital Object Identifier 10.1109/TVT.2024.3506948

[2], [3]. Communication systems in high-frequency ranges, e.g., millimeter-wave (mmWave) and sub-terahertz (sub THz), present a way to meet the first of those demands. This is primarily due to their abundance of bandwidth that helps achieve data rates in excess of tens of Gbps [2], [4]. However, high-frequency systems face challenges on many levels, and one of the most critical challenges is the relatively large beam-training overhead. Signal propagation in the high-frequency domain is characterized by poor penetration ability and suffers from high power loss due to scattering [5]. Therefore these systems need to periodically update the choice of beamforming vectors at both transmitters and receivers to maintain satisfactory Signal-to-Noise Ratios (SNRs). Such need is a common source of strain for those systems, and it has driven much research in the wireless community for innovative solutions to reduce the training overhead.

Some recently emerging approaches to deal with many high-frequency wireless communication challenges revolve around machine learning, and computer vision [6], [7], [8], [9], [10], [11], [12]. Those approaches collectively define the *Vision-Aided Wireless Communications* (ViWiComm) framework. Within that framework, a wireless system utilizes computer vision, multimodal machine learning, and deep learning to develop an understanding of the wireless environment and its elements. The system, then, taps into that understanding to address some of the adversities it faces, like the beam-training overhead. A key advantage of the ViWiComm framework is the information-rich sensors it introduces to the wireless system, RGB cameras [12], [13], radars [14], and LiDAR sensors [15], to name a few. These sensors provide much-needed information about the wireless environment that is commonly under-utilized or ignored altogether.

Much of the work on addressing beam training [7], [16], [17] with ViWiComm either assumes simplified wireless settings or is based on synthetic datasets like the ViWi dataset [6]. One may question the practicality of the developed framework, given that real wireless communication environments are characterized by two inherent properties: dynamism and visual diversity. Dynamism in wireless environments refers to the continuous change in the locations of radio transmitters and receivers, which naturally leads to time-varying wireless channels. This property is a serious and well-acknowledged challenge in wireless communications [1], [18], for it is the main reason that beam training needs to be performed frequently. Furthermore, dynamism partially contributes to the second property, visual diversity. The wireless environment is fairly visually complex; it is composed of visually diverse objects (e.g., trees, buildings,

people, cars, buses, etc.), some of which are continuously on the move, causing the visual scene to vary. From a vision perspective, visual diversity is a serious challenge to a ViWiComm as it leads to the multi-candidate dilemma; it is the case when the composition of the wireless environment includes multiple objects that could constitute a possible wireless transmitter, see [11]. By considering these two properties, the practicality issue of ViWiComm could be summarized in the form of two questions:

Q.1: *Could the encouraging results obtained with synthetic datasets (e.g., [7], [16]) be extended to datasets collected from real wireless environments?*

Q.2: *Could the vision-aided wireless communication framework tackle the beam-training challenge in wireless environments with multiple candidate transmitters?*

Answering these two questions constitutes the cornerstone of this study. More specifically, the study builds on top of the work in [11], [12], [19] and provides a detailed evaluation of the beam training challenge with the ViWiComm framework.

A. Prior Work

The beam-training challenge in high-frequency wireless systems has been investigated in several studies like [20], [21], to name a few. In recent years, a considerable amount of research has explored machine-learning-based approaches to tackle that challenge. For this study, that research is surveyed from the perspective of computer vision. This means the proposed approaches in the literature will be categorized based on whether they utilize visual data or not.

Beam prediction without visual data: Many studies have considered developing machine learning algorithms for beam prediction using wireless and/or position sensory data. Good examples could be found in [1], [22], [23]. The work in [1] introduces a novel coordinated beamforming approach based on mmWave omni- or quasi-omni-channels. A set of coordinating mmWave basestations estimates their local mmWave channels using omni or quasi-omni beam patterns and then uses those channels to train a Deep Neural Network (DNN) to predict the beamforming vectors at each basestation. The study in [22] takes a different look at what could be used as sensory data. It proposes a new approach based on sub-6GHz channels and DNNs. It poses the beam prediction problem as a classification problem to which a DNN is trained to observe the sub-6GHz channels and predict the optimal beamforming vector. The use of position sensory data to tackle the beam-training challenge has been explored in [23]. It proposes to use receiver position and types of neighboring vehicles to handcraft a per receiver feature vector, and it trains an ensemble classifier based on a random forest model to predict the optimal beamforming vectors using that feature vector.

The existing work has shown promise in addressing the beam-training challenge. However, there is potential for further improvement by considering better sensory data and learning methods. Currently, the chosen sensory data is insufficient to

capture the complexity and dynamics of the wireless environment. The wireless or position data only offer partial information about objects and their movements. Additionally, the previous research relies solely on unimodal machine learning. Future wireless systems, including basestations and user equipment, are expected to incorporate multiple sensors such as RGB cameras, sub-6GHz transceivers, LiDARs, radars, and GPS [24]. Utilizing all available sensory data instead of relying on a single modality would be more practical. Therefore, machine learning research should focus on developing algorithms that can effectively extract meaningful features by combining different modalities.

Beam prediction with visual data: This category includes beam prediction approaches that address the two shortcomings mentioned above by utilizing visual sensory data and multimodal machine learning. Using computer vision to address the beam-training challenge in high-frequency wireless systems is rooted in the early work in [6], [7]. The ViWiComm framework, in which wireless systems are equipped with visual data sources, has first been introduced in [6] along with the first Vision-Wireless (ViWi) dataset. Using ViWi, [7] presents the first case study for beam prediction using the ViWiComm framework. It proposes a beam prediction approach based on Convolutional Neural Networks (CNNs) for wireless settings with a single-candidate user. The work in [16] extends that in [7] by utilizing object detectors. It takes a step closer to studying ViWiComm in realistic wireless settings by synthesizing images with multiple candidate users using a scenario from the ViWi dataset. Some common limitations of the early work include focusing on simple communication scenarios, such as settings with single-candidate users or artificially generated multi-candidate users, and not fully exploring the potential of multimodal machine learning. Wireless environments, as mentioned earlier, are characterized by dynamic changes and visual variations. This poses a significant challenge when predicting beams based solely on visual data, as visual information alone does not provide any insight into the identity of the object responsible for the radio signal. To address this issue, previous studies such as [11], [25] have proposed multimodal machine learning algorithms that combine both visual and wireless data to identify the radio transmitter in the environment. However, it is important to note that these studies do not specifically tackle the beam-training challenge in practical wireless settings.

B. Novelty and Contribution

This study marks a significant leap in ViWiComm, particularly in addressing beam-training challenges in practical, high-frequency wireless environments. Unlike previous research that often relies on simplified or synthetic scenarios, this work deals with real-world complexities characterized by dynamism and visual diversity. Our novel contribution lies in extending ViWiComm's promising results to authentic wireless settings and introducing an innovative multimodal deep learning framework that integrates visual and position data for beam prediction in multi-candidate environments. By utilizing DeepSense 6G scenarios, we bridge the gap between theoretical models and practical applications. This approach not only addresses the

crucial questions of ViWiComm’s practicality in real environments and its efficacy in multi-candidate settings but also paves the way for significant reductions in beam-training overhead in 5G and beyond systems. Consequently, our work represents a vital step towards making ViWiComm a practical reality in next-generation wireless networks, offering a solution to the challenges posed by frequent beam training in dynamic, visually complex environments. In particular, it addresses the two questions **Q.1** and **Q.2** by utilizing visual data and deep learning models. The main contributions of this paper could be summarized as follows:

- *Beam prediction in single-candidate settings:* As a stepping stone, this paper addresses **Q.1** using a dataset constructed using DeepSense 6G scenarios [19], which represent real wireless communication environments. It shows that those encouraging results in [7], [16] could indeed be achieved in real wireless communication settings with a single-candidate user.
- *Beam prediction in multi-candidate settings:* We address **Q.2** by introducing a pioneering deep neural network that fuses visual and position data for beam prediction in complex, multi-candidate wireless environments. This innovative approach not only identifies the correct radio transmitter among multiple candidates but also predicts its optimal beamforming vector. This solution represents the first demonstration of a multimodal ViWiComm DNN tackling beam prediction in practical, visually diverse settings with multiple potential transmitters.
- *Large multimodal dataset and comprehensive evaluation:* By utilizing the different scenarios in DeepSense, two datasets of co-existing multimodal data points (i.e., visual, position, and wireless) are constructed for development and testing purposes. Those datasets constitute the seeds for the comprehensive evaluation experiments conducted to answer the questions in **Q.1** and **Q.2**. We provide a detailed evaluation of our vision-position aided beam prediction approach using the DeepSense 6G dataset, which represents real-world wireless scenarios. This analysis offers concrete insights into the performance gains, overhead reduction, and efficiency improvements achievable in practical 5G systems.

II. VISION-AIDED BEAM PREDICTION

Overcoming the challenge of beam training in high-frequency wireless communication systems (mmWave and sub-THz) is a cornerstone in realizing the full potential of those systems. Recent efforts to deal with this challenge have seen increasing interest in leveraging artificial intelligence, and deep learning in particular [1], [7], [16], [22], [26]. Among those efforts, vision-aided beam prediction is a promising approach proposed to reduce the beam training overhead in high-frequency systems. As stated in Section I-B, this work provides a comprehensive and realistic evaluation of the potential of vision-aided beam prediction solutions. In the following three subsections, we i) motivate the need for vision, ii) introduce the key idea of vision-aided beam prediction, and iii) present the overall flow of this paper.

A. Motivation

Harvesting the large bandwidth available at the mmWave and sub-THz bands requires using narrow beams and overcoming their alignment challenge [2], [5]. The true challenge with beam alignment stems from the inherently dynamic nature of the wireless communication environment; transmitters, receivers, and scatterers could all be on the move. A direct implication of that dynamics is that transmitters and receivers need to periodically update their choice of beamforming vectors to maintain a satisfactory SNR level and, preferably, a LOS connection. This update comes in the form of beam-training, which is a recognized burden in high-frequency wireless communications [1], [22]. An interesting and novel approach for handling the beam-training burden could be found in embracing a striking resemblance between high-frequency communication and computer vision systems, which is their reliance on LOS [8], [27]. Since high-frequency signals struggle in penetrating objects in the wireless environment and lose a significant amount of power due to scattering [5], there is a quite large SNR margin between LOS and NLOS communication links that skews in favor of LOS. This makes LOS a preferable setting in high-frequency communications, and it draws a connection with computer vision, which is inherently LOS. The data usually captured and analyzed in a computer vision system depicts what is *visible* in the scene, starting with simple patterns (e.g., edges, colors, etc.) to abstract concepts (e.g., human, dog, tree, etc.). As such, the information contained in visual data could be as valuable to a high-frequency system as it is to a computer vision system, *begging the question of how computer vision could be used to mitigate the beam-training challenge.*

B. Key Idea

The reliance on LOS is the key feature that links high-frequency communications to computer vision and is the bedrock of vision-aided beam prediction. To understand the connection and how it is used to overcome the large beam training overhead challenge, consider the example depicted in Fig. 1, where a high-frequency basestation serves some users in its surrounding environment. Without loss of generality, the basestation is assumed to employ a well-calibrated high-frequency antenna array ¹. Some users experience LOS connections with the basestation in this environment, while others have NLOS links. In a classical system operation, the basestation performs beam training to identify suitable beams for each user. Two examples are illustrated in the “top-view” image in Fig. 1, namely beams \mathbf{f}_{LOS} and \mathbf{f}_{NLOS} . An important factor in the beam choice is the user’s relative position with respect to the basestation. For instance, \mathbf{f}_{LOS} is clearly a LOS beam, and it represents a straightforward path between the basestation and the green SUV. The NLOS beam \mathbf{f}_{NLOS} , does not point to the red car but to a visible building (i.e., reflector) that leads to the red car. Although simple, these observations highlight an important property that could be utilized to address the beam-training challenge; both the SUV and the building are in the field of view of the basestation

¹Well-calibrated here means the array geometry is known, and no hardware impairments are assumed, more on those issues could be found in [28], [29]




Fig. 1. An illustration of a high-frequency wireless communication system and its environment. “Basestation View” depicts the environment from the basestation perspective, showing some LOS users and possible LOS reflectors. “Top View” shows the invisible part of that environment, i.e., NLOS users, and the beamforming vectors used to serve LOS and NLOS users.

and are qualified as LOS objects. This suggests that detecting those objects and understanding their roles could be key to identifying the choice of beamforming vectors.

A high-frequency wireless system could tap into the advances in the fields of computer vision and machine learning [30] to realize the notion of detecting objects and understand their roles, and perform vision-aided beam prediction. Abstractly, using a machine learning algorithm designed for ViWiComm, the task of predicting beams could be broken down into three stages: *scene analysis*, *object-role identification*, and *decision making*. The first stage is where the machine learning model directly operates on the visual data to extract contextual information. Referring back to Fig. 1, this is loosely equivalent to detecting various objects of importance in the scene, like cars, buses, pedestrians, trees, buildings, etc. Those objects are passed on to the second stage, in which the machine learning attempts to identify the roles of those objects relative to the wireless system; this means labeling the objects in Fig. 1 as candidate users, candidate reflectors, or candidate signal blockages. Such labeling will be discussed under the concept of transmitter identification in Section V. This will highlight that visual data alone is not sufficient for identifying the transmitter in the scene. Hence, the object-role identification stage requires augmenting what has been learned from the previous stage with other sensory data, such as wireless, LiDAR, or GPS data. Recognizing the role of each object gets the machine learning ready to go into the decision-making stage, where it selects the suitable beams to serve the LOS/NLOS users.

C. Flow

This work presents a real-world study on realizing vision-aided beam prediction and demonstrating its ability to tackle

the beam-training challenge. The study is conducted in two major phases: i) Beam-prediction in single-candidate settings, where there is only one candidate user in the visual scene and ii) beam-prediction in multiple-candidate settings, where multiple objects/candidate users exist in the visual scene. The two phases represent the evolution of the ViWiComm framework first proposed in [7] from conception to real-world implementation in realistic multi candidate user settings. This study takes advantage of the recently developed DeepSense 6G dataset [19] that reflect real wireless communication environments. The dataset is designed for multimodal machine learning research in wireless communication. It consists of various scenarios where multimodal sensing and communication data samples are collected using a multi-sensor testbed; see [19] for more information.

III. SYSTEM MODEL

This paper adopts a system model that consists of a basestation, deployed on the sidewalk, and a vehicular mobile user, similar to the system depicted in Fig. 1. The basestation is equipped with a uniform linear array (ULA) with M elements, a standard-resolution RGB camera, and a GPS. For practicality [22], the basestation is assumed to employ analog-only architecture with a single RF chain and M phase shifters. The basestation adopts a predefined local beamforming codebook $\mathcal{F} = \{\mathbf{f}_q\}_{q=1}^Q$, where $\mathbf{f}_q \in \mathbb{C}^{M \times 1}$ and Q is the total number of beams. This codebook spans a ULA field of view of γ° along the azimuth plane. Adopting an OFDM with a cyclic prefix of length D and a number of subcarriers K , the received downlink signal at the mobile unit is given by

$$y_k = \mathbf{h}_k^T \mathbf{f}_x + \beta_k, \quad (1)$$

where $y_k \in \mathbb{C}$ is the received signal at the k th subcarrier, $\mathbf{f} \in \mathcal{F}$ is the selected beamforming vector, $\mathbf{h}_k \in \mathbb{C}^{M \times 1}$ is the channel between the BS and the mobile unit at the k th subcarrier, $x \in \mathbb{C}$ is a transmitted complex symbol that satisfies the following constraint $\mathbb{E}[|x|^2] = P$, where P is a power budget per symbol, and finally β_k is a noise sample drawn from a complex Gaussian distribution $\mathcal{N}_{\mathbb{C}}(0, \sigma^2)$. The choice of the beam is determined using classical beam training, in which the basestation sweeps the codebook \mathcal{F} looking for the optimal vector \mathbf{f}^* . Formally, that sweep could be expressed by

$$\mathbf{f}^* = \underset{\mathbf{f}_q \in \mathcal{F}}{\operatorname{argmax}} \frac{1}{K} \sum_{k=1}^K \log_2 (1 + \operatorname{SNR} |\mathbf{h}_k^T \mathbf{f}_q|^2), \quad (2)$$

where SNR is the signal-to-noise ratio. However, in LOS-dominated system operation (almost single-path channels), (2) could be approximated by

$$\mathbf{f}^* = \underset{\mathbf{f}_q \in \mathcal{F}}{\operatorname{argmax}} \frac{1}{K} \sum_{k=1}^K |\mathbf{h}_k^T \mathbf{f}_q|^2. \quad (3)$$

The channel model adopted throughout this paper is the geometric mmWave channel model. This model choice comes as a result of two facts: i) the model captures the limited scattering property of the mmWave band [18], [31], and ii) the experimental results in this paper are based on real measurements, which are

captured well by the geometric model. The channel vector \mathbf{h}_k in (1) is given by

$$\mathbf{h}_{u,k} = \sum_{d=0}^{D-1} \sum_{\ell=1}^L \alpha_{\ell} e^{-j\frac{2\pi k}{K} d_p (dT_S - \tau_{\ell})} \mathbf{a}(\theta_{\ell}, \phi_{\ell}), \quad (4)$$

where L is number of channel paths, $\alpha_{\ell}, \tau_{\ell}, \theta_{\ell}, \phi_{\ell}$ are the path gains (including the path-loss), the delay, the azimuth angle of arrival, and the elevation angle of arrival, respectively, of the ℓ th channel path. T_S represents the sampling time while D denotes the cyclic prefix length (assuming that the maximum delay is less than DT_S).

IV. SINGLE-CANDIDATE SETTINGS

This section starts the first phase of this study, where vision-aided beam prediction is studied in a real wireless setting with a single candidate user. These settings are the immediate and natural extension to those in [7]. This phase is structured into two parts. The first part includes the formal definition of the single-candidate beam-prediction problem, the proposed deep learning solution to that problem, and a brief discussion on the practical challenges associated with that solution, motivating the transition to the second phase of this study. The second part of this phase presents a detailed evaluation of the proposed solution on the real-world dataset, and it will be discussed in Section VIII-A.

A. Problem Definition

The beam prediction task in single-candidate wireless settings is defined as follows:

Given a wireless communication environment where there is only one possible high-frequency transmitter, the vision-aided beam prediction at the infrastructure is the task of predicting the optimal beam indices from a pre-defined codebook by utilizing a machine learning model and the images captured by the camera installed at the basestation

Formally, the problem can be defined as follows. A dataset of image-beam pairs (samples) is collected from real wireless environments where each pair has an image with a single-candidate transmitter and its best beamforming vector. This dataset could be given by $\mathcal{D}_{\text{task}_1} = \{(\mathbf{X}_u, \mathbf{f}_u^*)\}_{u=1}^U$ where $\mathbf{X}_u \in \mathbb{R}^{H \times W \times C}$ is the RGB image of the u th pair in $\mathcal{D}_{\text{task}_1}$ with height H , width W , and number of channels C , and U is the total number of pairs in the dataset. At any time instant t , the objective of the single-candidate beam prediction task is to find a prediction/mapping function f_{Θ_1} that utilizes the available sensory data $\mathbf{X}[t]$ to predict (estimate) the optimal beam index $\hat{\mathbf{f}}[t] \in \mathcal{F}$ with high fidelity. The mapping function can be formally expressed as

$$f_{\Theta_1} : \mathbf{X}[t] \rightarrow \hat{\mathbf{f}}[t]. \quad (5)$$

In this work, we design a deep learning algorithm to learn a prediction function f_{Θ_1} parameterized by a set of parameters Θ_1 from the dataset $\mathcal{D}_{\text{task}_1}$. The objective of the learned function is to maximize the overall correct prediction probability over all

samples in $\mathcal{D}_{\text{task}_1}$, which could be expressed as follows

$$f_{\Theta_1}^* = \max_{f_{\Theta_1}} \prod_{u=1}^U \mathbb{P}(\hat{\mathbf{f}}_u = \mathbf{f}_u^* | \mathbf{X}_u), \quad (6)$$

where the product is the result of an implicit assumption that the samples of $\mathcal{D}_{\text{task}_1}$ are drawn independently from an unknown distribution $\mathbb{P}(\mathbf{X}, \mathbf{f}^*)$ that models the relation between \mathbf{X} and \mathbf{f}^* .

B. Proposed Solution

The proposed deep learning solution leverages the important parallel between vision and high-frequency communications, primarily focusing on Line-of-Sight (LOS) scenarios. This emphasis on LOS is driven by two key factors: frequency dependence and the effectiveness of the vision modality. At higher frequencies (mmWave and sub-THz), the reliance on LOS connections increases significantly as blockages can severely degrade the SNR. Additionally, the primary modality in our approach, vision, is inherently LOS-dependent. Our solution attempts to learn beam prediction by identifying the user's location in the environment. A beam-steering codebook deployed at the basestation induces sectoring of the wireless environment across the azimuth plane [28], which can be projected onto the image plane, resulting in visual sectors that could be regarded as classes [11]. Given the assumption of a single candidate in the environment, the user's location in the image defines the sector to which that user belongs, referred to as the object-sector assignment. Therefore, the prediction function $f_{\Theta}(\mathbf{X})$ should learn such an assignment to predict the best beamforming vector. While our current focus is on LOS scenarios, we acknowledge the importance of Non-Line-of-Sight (NLOS) situations. In NLOS scenarios, visual data can aid in 3D scene understanding, helping to identify the geometry of relevant objects and the location of probable reflective surfaces, potentially narrowing down the beam search space even when the direct path is obstructed.

Based on these considerations, this study proposes a modified residual neural network (ResNet) [33], specifically ResNet-50 as shown in Fig. 2. The architecture is pre-trained on the ImageNet dataset and modified to incorporate a new M -class classifier layer. This choice is rooted in two main facts: the effectiveness of residual learning in building very deep architectures and the strong performance of ResNets in various computer vision tasks. Residual blocks, the fundamental element of ResNets, prevent the performance degradation commonly associated with training deep architectures. Additionally, ResNets have demonstrated good performance in many computer vision tasks, including image classification [33], object detection [34], and semantic segmentation [35]. This versatility makes ResNet a suitable choice for our beam prediction task, which requires both understanding the overall scene and identifying specific objects within it.

²A common assumption in the realm of machine learning, see [32] for example.

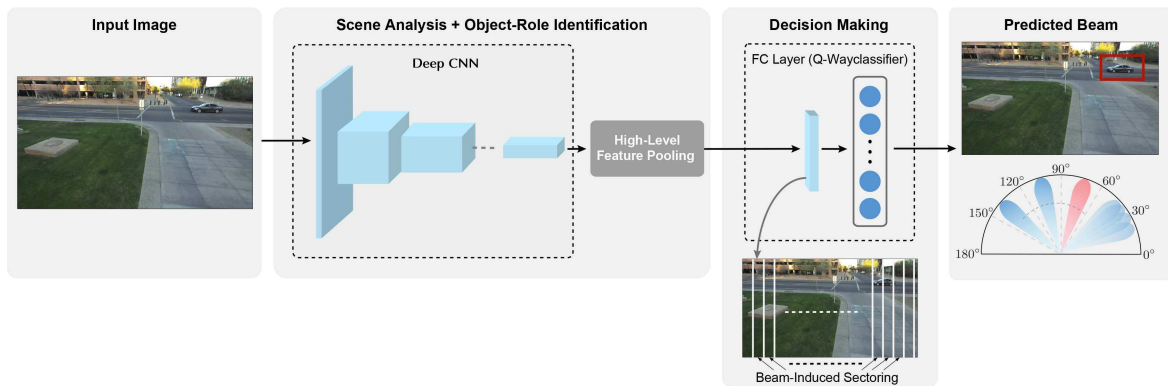


Fig. 2. This figure illustrates the proposed machine learning-based vision-aided beam prediction model that leverages visual data captured at the basestation for mmWave/sub-THz beam prediction in a single-candidate setting.

C. Challenges

The proposed ResNet architecture above is expected to face a critical challenge when deployed in real mmWave environments. The primary reason is the assumption of a single candidate user, i.e., the proposed solution searches for a single candidate user in the scene to assign to a sector. As a result, the proposed solution is incapable of handling a scenario with multiple users. For instance, when two vehicles are present in the image (the wireless environment), and they belong to different visual sectors, it is impossible to predict the beam without identifying the roles of each vehicle (transmitter, moving blockage, or possible reflector). This dilemma with multiple candidates is the building block for the third phase of this study, which is discussed in the following section.

V. MULTI-CANDIDATE SETTINGS

The interest in addressing the dilemma of beam prediction in environments with multi-candidate users is at the center of the second phase of this study. As discussed in Section IV-C, a vision-aided beam prediction algorithm needs, in some way or another, to realize the three-stage process described in Section II-B. The key challenge in doing so is the ability to identify the roles of every candidate in the environment and differentiate the connected users from the other objects in the environment. In order to deal with this challenge, we need to perform what we call *user identification* in the visual scene by leveraging other user attributes (that could also be captured using other sensing modalities). In this section, we investigate this problem and propose a DNN architecture to perform the task of beam prediction in multi-candidate settings. This establishes the bases for enabling vision-aided multi-user communications.

A. Problem Definition

The beam prediction task in multi-candidate wireless settings is defined as follows:

Given a wireless communication environment where there are multiple objects that could visually constitute wireless transmitters, the beam prediction task in multi-candidate settings is defined as

the problem of predicting the optimal beamforming vector from a pre-defined beam codebook using a pool of multimodal data that includes vision.

The task is formally defined as follows: Let \mathcal{V} be a v -tuple of multimodal data samples that includes vision, i.e., $\mathcal{V} = (\mathbf{X}, \mathbf{g}_1, \dots, \mathbf{g}_v)$ where \mathbf{g}_1 to \mathbf{g}_v are vectors containing the other modality data samples. Then, a dataset of $(v + 1)$ -tuples is collected from a real wireless environment $\mathcal{D}_{\text{task}_2} = \{(\mathcal{V}_u, \mathbf{f}_u^*)\}_{u=1}^U$ where U is the total number of samples in $\mathcal{D}_{\text{task}_2}$ and \mathbf{f}_u^* is the optimal beam in \mathcal{F} associated with the u th v -tuple \mathcal{V}_u . At any time instant t , the objective of the single-candidate beam prediction task is to find a prediction/mapping function f_{Θ_2} that utilizes the available sensory data $\mathcal{V}[t]$ to predict (estimate) the optimal beam index $\hat{\mathbf{f}}[t] \in \mathcal{F}$ with high fidelity. The mapping function can be formally expressed as

$$f_{\Theta_2} : \mathcal{V}[t] \rightarrow \hat{\mathbf{f}}[t]. \quad (7)$$

In this work, we design a deep learning algorithm to learn a prediction function f_{Θ_2} parameterized by Θ_2 . This function needs to maximize the probability of correct beam prediction given the image and position data, i.e.,

$$f_{\Theta_2}^* = \max_{f_{\Theta_2}} \prod_{u=1}^U \mathbb{P}(\hat{\mathbf{f}}_u = \mathbf{f}_u^* | \mathcal{V}). \quad (8)$$

Again, similar to the formulation of (6), the product in (8) is a result of an implicit assumption that the samples of $\mathcal{D}_{\text{task}_2}$ are independent and identically distributed, i.e., follow the same unknown joint distribution $\mathbb{P}(\mathcal{V}, \mathbf{f}^*)$.

B. The Choice of Data Modalities for User Identification

From the definition above, the ultimate objective of beam prediction in multi-candidate settings is still the same as that in Section IV-A, developing a deep learning algorithm that predicts optimal beamforming vectors; however, performing the beam-prediction task requires more than visual data in those settings. This is a direct consequence of the fact that multiple candidate users share the same visual traits, see Section II-B and therefore, their roles from the wireless system perspective *cannot be visually determined*. This calls for a secondary source of information

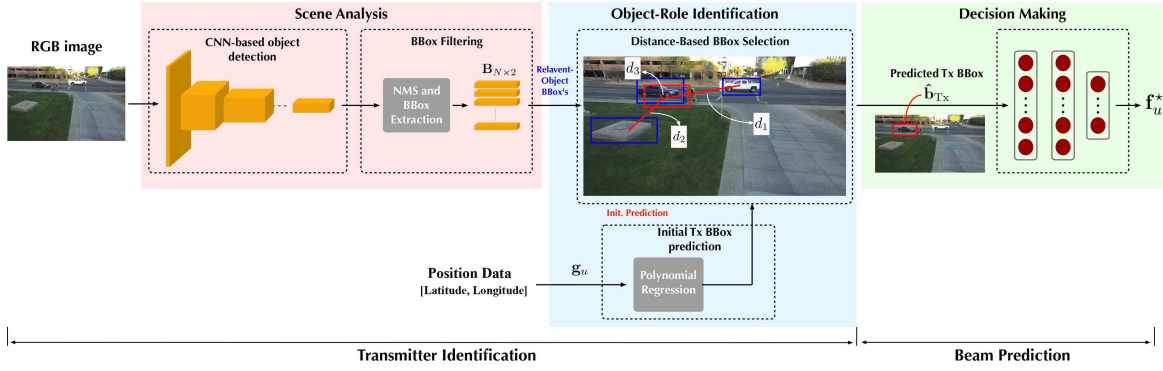


Fig. 3. The figure presents the proposed multi-modal machine learning based beam prediction model that leverages both visual and positional data to predict the optimal beam indices in a multi-candidate settings.

that could augment visual data, and the choice for such source in this study is GPS (i.e., $\mathcal{V} = (\mathbf{X}, \mathbf{g})$ where $\mathbf{g} \in \mathbb{R}^2$ is a vector of latitude and longitude coordinates). It is motivated by two important observations. First, position information is intimately related to visual information; GPS provides x-y coordinates for objects in the 3-dimensional world, and an image is a projection of that 3-dimensional world onto a 2-dimensional plane. Hence, in some sense, the position data complements what is missing in the visual data, which is the sense of distance. The second reason is that position data is lightweight, making them easily exchangeable between a basestation and its candidate user. It is quite important to emphasize at this point that GPS (or position data) is *meant to augment visual data and not replace them*. As shown in [12], position data could be of help to beam prediction, yet they lack a sense of surrounding; position data only indicate where candidate users are, and they do not account for contextual information (shapes and relations between those candidates). For instance, visual data reflect information about the shape and type of candidates (e.g., large vehicle, small vehicle, pedestrians, etc.) and their relation to one another, which provide contextual information to the machine learning algorithm.

C. Proposed Solution

This subsection presents the proposed solution for beam prediction in a real-wireless environment with multiple transmitting candidates. It proposes a novel approach that utilizes bimodal visual and position data in $\mathcal{D}_{\text{task}_2}$ to predict optimal beamforming vectors. The proposed solution follows the three-stage sequence outlined in Section II-B. More to the point, it breaks down the function $f_{\Theta}(\mathbf{X}, \mathbf{g})$ into two major components: user (transmitter) identification and beam prediction. The first component utilizes visual data to detect relevant objects in the environment (i.e., scene analysis), then identifies the radio transmitter among the objects using position data (i.e., object-role identification). The second component takes in the extracted information about the transmitter and its surrounding objects and predicts the optimal beamforming vector (i.e., decision-making). Fig. 3 shows a schematic of the proposed architecture.

1) *Transmitter Identification:* The goal of the first component of the proposed multi-candidate beam prediction solution is

to identify the candidate transmitter in the image, i.e., transmitter identification. For that task, a two-step architecture is proposed. The first step of the proposed architecture relies on DNNs to produce bounding boxes enclosing relevant objects in the scene. It is performed to detect all the probable transmitting objects in the environment. In the second step, the DNN uses position data to filter out detected candidates that are not the radio transmitter. A deeper look at the two-step DNN architecture is given below.

Bounding box detection: In order to detect the transmitting candidate in real-wireless settings, the first step is to identify all the relevant objects in the scene (scene analysis). A pre-trained object detector is adopted for this purpose. The object detector is modified to detect two classes of objects in the scene, labeled as “Tx (transmitter)” and “No Tx (Distractors)”. The former label encompasses all objects relevant to the wireless system in the scene. For example, in a scene depicting a city street, relevant objects include, but are not limited to, cars, trucks, buses, pedestrians, and cyclists. The other label includes those cases where no relevant objects are present in the scene. The modified object detector is fine-tuned in a supervised fashion using a subset of the manually labeled dataset described in Section VI-A. In this work, a YOLOv3 architecture is selected for the bounding box detection task. It provides accurate detection at a relatively high frame rate, reducing inference latency. During inference, the fine-tuned YOLOv3 model generates bounding boxes for the detected candidates in the scene and their confidence scores. By using those output boxes, the relevant-object matrix $\mathbf{B} \in \mathbb{R}^{N \times 2}$ is constructed such that each row has only the normalized coordinates of the center of a bounding box, see Fig. 3.

Bounding box selection: In this step, both relevant-object matrix \mathbf{B} and position data are utilized to identify the probable transmitter in the scene. This process starts by learning a prediction function that estimates the bounding box center of a transmitter using its position information, encoding the relation between object position in the 3D world and object location in the image. The function is learned using a 3^{rd} -degree polynomial regression model

$$\hat{\mathbf{b}}_{\text{Tx}} = \mathbf{W}^T \phi, \quad (9)$$

where $\hat{\mathbf{b}}_{\text{Tx}} \in \mathbb{R}^{2 \times 1}$ is a vector with an initial prediction of the centers of a transmitter object, and \mathbf{W} is an $A \times 2$ parameter

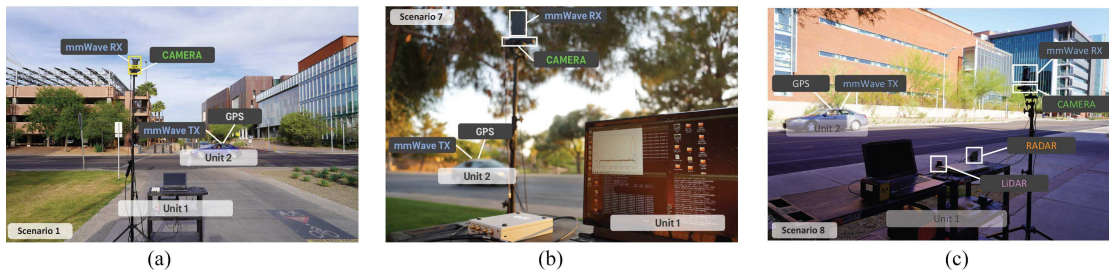


Fig. 4. This figure presents the DeepSense 6G testbed 1 used during the data collection. It consists of a stationary and mobile unit. The stationary basestation (unit1) is equipped with a mmWave receiver (60GHz band) and sensory suite (RGB camera, GPS, radar and LiDAR). The mobile unit, acting as a transmitter, is equipped with a 60GHz quasi-omni antenna and GPS receiver.

matrix, the $A \times 1$ vector ϕ is a feature vector obtained from the 3^{rd} -degree polynomial transformation of the predictors \mathbf{g} , with A denoting the number of unique monomials in a bivariate³ 3^{rd} -degree polynomial, i.e., $A = 9$. The parameter matrix \mathbf{W} is learned from the dataset $\mathcal{D}_{\text{task}_2}$ —more on the training of this model in Section VII—and once it is learned, the model could be used to get an initial estimate on the center of the transmitter bounding box. Since $\hat{\mathbf{b}}_{\text{Tx}}$ is an initial estimate that solely relies on position data, it is not expected to be a final prediction but merely a guide. It is used in conjunction with the relevant-object matrix \mathbf{B} to identify (or select) the object responsible for the radio signal (transmitter). This is done using the nearest neighbor algorithm with an Euclidean distance metric. In other words, the row of \mathbf{B} with the shortest distance to $\hat{\mathbf{b}}_{\text{Tx}}$ is picked as the nearest neighbor and, hence, the predicted user (transmitter).

2) *Beam Prediction*: The second component of the proposed solution is a feed-forward neural network that predicts the optimal beam given the identified transmitter. That this component has enough contextual information from the previous one makes it capable of realizing the third and last stage of the three-stage sequence, which is decision-making. A 2-layer feed-forward neural network is developed to perform that beam prediction task. In particular, the prediction task here is posed as a classification problem. The input to the network is the center coordinates of the identified transmitter, and the output is the best beam in the codebook \mathcal{F} .

VI. DEVELOPMENT DATASET

As the cornerstone of this study is to answer **Q.1** and **Q.2** in Section I, an experimental setup is built around multi-modal real-world sensor measurements collected from real wireless environments as shown in Fig. 4. This is done by utilizing the publicly available DeepSense dataset [19] and constructing a large development dataset with tuples of RGB images, mmWave beams, GPS positions, and bounding boxes for the transmitters and distractors. The considered scenarios from DeepSense and the constructed final development datasets are discussed below.

³It is bivariate because the vector of predictors \mathbf{g} is 2 dimensional, see [32], [36] for more information on polynomial regression.

A. Communication Scenarios and Development Dataset

The DeepSense 6G dataset provides a variety of outdoor wireless communication scenarios with different data modalities [19]. To evaluate the two-beam prediction problems defined in Sections IV and V, we select scenarios 1 to 8 from the DeepSense dataset. They represent eight outdoor wireless environments with vehicles as the main candidate transmitters. Furthermore, the scenarios were collected at different times of the day and in varied weather conditions to increase the overall diversity of the dataset. All eight scenarios contribute a total of ≈ 18000 data samples, each of which is a tuple of RGB image, mmWave received power, and GPS position, and across all of them, the wireless system deploys a beam-steering codebook of 64 beams. In Fig. 5, we present the data samples from the 8 different scenarios of the DeepSense dataset. The proposed beam prediction tasks and their solutions mandate slightly different types of communication scenarios and data modalities. Therefore, the raw dataset passes through a processing pipeline to filter out data samples unrelated to a particular task. This leads to two different development datasets, one for each task. The details of this process are given below:

Single-candidate dataset: The task defined in Section IV requires data collected from the wireless environment with a single candidate in the FoV of the basestation. The first step in generating the development dataset is to filter out samples that do not fit the single-candidate setting from the raw dataset. This is done by manually examining the vision data of each DeepSense scenario. Doing so would show that scenarios 5, 6, 7, and 8 have some multimodal data that pertain to the single-candidate setting. Table I lists those scenarios and the number of samples they contribute to the development dataset. The selected samples undergo a second processing step, in which only visual and wireless data is retained. The task relies on visual data as inputs and optimal beams as targets (labels); therefore, only RGB images and mmWave received power data are picked from the modalities of each scenario. The power data go through one extra step to generate the optimal beams. Each received power vector is first downsampled to 32 elements by selecting every other element in the vector. Since the basestation receives the mmWave signal using an oversampled codebook of 64 pre-defined beams, the downsampling does not affect the

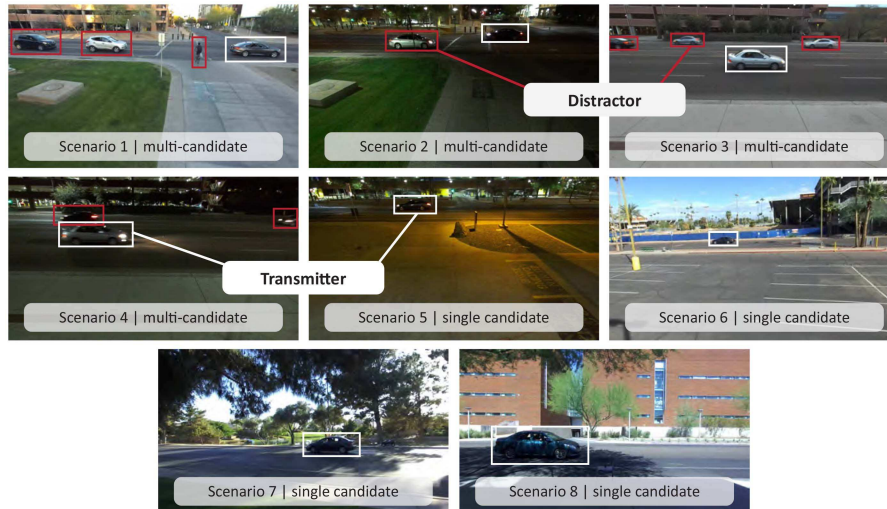


Fig. 5. This figure shows the image samples from the different scenarios in the DeepSense 6G dataset. As shown in this figure, scenarios 1 – 4 are multi-candidate scenarios, i.e., more than one objects-of-interest (vehicles) are usually present in the FoV of the basestation and have been utilized to investigate the multi-candidate beam prediction problem statement. Scenarios 5 – 8 primarily consists of a single object and is useful for evaluating the performance of the proposed sensing-aided single-candidate solution.

TABLE I
SINGLE-CANDIDATE AND MULTI-CANDIDATE DATASET

Dataset	Location	Time of Day	Number of Samples		Speed Distribution	
			Training	Validation	Mean (km/h)	StdDev (km/h)
Single-Candidate	Tyler Parking (Scenario 5)	Night	1817	736	18.34	10.11
	Lot-59 Parking (Scenario 6)	Day	812	348	17.97	7.38
	Downtown Chandler (Scenario 7)	Day	641	275	43.12	5.73
	Bio-design (Scenario 8)	Day	3077	1320	48.92	5.35
Multi-Candidate	McAllister Ave. (Scenario 1)	Day	1904	816	19.35	2.95
	McAllister Ave. (Scenario 2)	Night	2197	942	31.35	2.14
	Rural Rd. (Scenario 3)	Day	1120	477	66.14	16.71
	Rural Rd. (Scenario 4)	Night	1363	592	13.31	2.34

total area covered by the beams. This implies that the effective size of the codebook in this paper is $Q = 32$. Then, out of the 32 elements per vector, the index of the beam with maximum received power is selected as the optimal beam (as described in Section III and given by (3)). The final outcome of this pipeline is a dataset $\mathcal{D}_{\text{task}_1} = \{(\mathbf{X}, \mathbf{f}^*)_u\}_{u=1}^{U_1}$ where $U_1 \approx 9000$.

Multi-candidate dataset: The multi-candidate task requires data samples collected from a wireless environment with multiple candidates. Thus, similar to the single-candidate dataset, the raw dataset is examined to identify the multi-candidate samples. This reveals that Scenarios 1, 2, 3, and 4 are more suited for multi-candidate settings, see Table I for more details. As described in Section V, the development dataset for this task requires the preparation of data tuples of RGB images, GPS position, and mmWave optimal beams. This means the extra step from the single-candidate dataset is also applied here to obtain the optimal beam indices for each data sample, i.e., downsampling and using (3). The final outcome of this simple pipeline is a dataset $\mathcal{D}_{\text{task}_2} = \{(\mathcal{V}_u, \mathbf{f}^*)_u\}_{u=1}^{U_2}$ where $U_2 \approx 9500$.

The two development datasets above are used to train and evaluate the performance of the proposed solutions. Both datasets are further divided into training and validation sets with a split of 70 – 30%. The details of the experimental setup are provided in the next section, while Section VIII is devoted to discussing the performance of the proposed solutions.

VII. EXPERIMENTAL SETUP

This section first presents an overview of the model training process and the hyper-parameters utilized to train the proposed machine learning model for both single-candidate and multi-candidate settings. Next, we discuss the metric used to evaluate the beam prediction performance of the proposed solution. All the experiments are performed on a single NVIDIA Quadro RTX 6000 GPU using the PyTorch deep learning framework.

Network Training: In Section IV-B and Section V-C, we proposed two different machine learning-based models for the single-candidate and the multi-candidate settings, respectively. Both the proposed solutions utilize the cross-entropy loss with Adam optimizer to train the models. The detailed hyper-parameters used to fine-tune the models are presented in Table II.

Single-Candidate: As described in Section IV-B, the proposed vision-aided beam prediction solution adopts an ImageNet pre-trained ResNet-50 object classification model. The model is further modified by removing the last output layer and replacing it with a fully-connected layer with $M = 32$ neurons. The proposed model is trained in a supervised manner with a dataset $\mathcal{D}_{\text{task}_1}$, comprising RGB images and its corresponding ground-truth beam index.

Multi-Candidate: The multi-candidate proposed solution consists of two major components, namely, i) transmitter

TABLE II
DESIGN AND TRAINING HYPER-PARAMETERS

Parameters	ResNet-50	MLP
Batch Size	32	32
Learning Rate	1×10^{-4}	1×10^{-2}
Learning Rate Decay	epochs 4 and 8	epochs 20 and 40
Learning Rate Reduction Factor	0.1	0.1
Dropout	0.3	0.3
Total Training Epochs	15	50
Number of Output Nodes (M)	32	32

identification and ii) beam prediction. The transmitter identification component consists of a pre-trained YOLOv3 model, which is further fine-tuned in a supervised fashion using a subset of manually labeled dataset $\mathcal{D}_{\text{task}_2}$. During inference, the YOLOv3 model is used to detect all the relevant objects and extract the bounding boxes of those objects. The proposed solution utilizes the user's positional data to select the most probable bounding box. After identifying the probable transmitter, the bounding box center coordinates, $\hat{\mathbf{b}}_{\text{Tx}}$, are then provided as input to the proposed feed-forward neural network. The proposed machine learning model is then trained in a supervised fashion using the ground-truth bounding-box coordinate and the corresponding beam index.

Evaluation Metric: , we present the details of the metric adopted to evaluate the efficacy of the proposed solution. The primary metric adopted for evaluation is the top- k accuracy. The top- k accuracy is defined as the percentage of the test samples where the ground-truth beam is within the top- k predicted beams. In this work, we utilize the top-1, top-2, top-3, top-4, and top-5 accuracies to compute the prediction performance of the proposed solution. In Section VIII, we present the in-depth evaluation of the proposed solution for both single-candidate and multi-candidate settings.

VIII. EXPERIMENTAL RESULTS

The performance of the proposed ViWiComm beam prediction solutions is studied in this section. The discussion is divided into two parts. The first will discuss the performance of the proposed DNN in single-candidate settings. It will highlight the main advantages and shortcomings of ViWiComm for beam prediction, setting the stage for the discussion on the multi-candidate setting. The second part will focus on the multi-candidate settings and show how the proposed solution can handle the beam prediction task in those practical settings.

A. Single-Candidate

The beam prediction performance of the proposed solution is studied from two different standpoints, machine learning and wireless communication. The first perspective includes experiments that evaluate the performance of the proposed DNN architecture from different machine-learning angles, e.g., performance per location, number of data samples needed for training, etc. Those experiments use machine learning metrics such as top- k accuracy, where $k \in \{1, 2, 3, 4, 5\}$ is the accuracy

rank, and confusion matrices. The second perspective attempts to translate the results of the machine learning evaluation into results pertaining to the wireless system performance, such as studying the implications of beam-prediction failure on the wireless receive power.

1) *Machine Learning Perspective:* The proposed DNN architecture is evaluated on the single-candidate dataset. This evaluation investigates three important questions: i) Does the proposed DNN perform consistently across different wireless environments? ii) Are there any advantages to training the proposed DNN on several environments simultaneously? Furthermore, iii) how many data samples are needed to learn the beam prediction in single-candidate settings, in general?

Could the DNN have consistent performance across different environments (scenarios)? This question attempts to identify whether the proposed DNN can achieve similar single-candidate beam-prediction performance across different wireless environments or not. Fig. 6 addresses the question by training and testing the DNN on each scenario individually and two combinations of scenarios. The first thing one might observe from the figure is how volatile the top-1 performance looks across different environments, ranging from $\approx 67\%$ to $\approx 84\%$. This could be attributed to the difference between the wireless environments. The four scenarios represent four different physical locations and two different times of day, i.e., scenario 5 represents a wireless environment operating at night. The fluctuation might initially seem alarming, for it suggests a sense of uneven DNN performance. However, a closer look at the top-3 and 5 performances negates that suggestion; for *it shows that the DNN could produce consistent or nearly uniform performance*. From the figure, the DNN registers almost the same accuracies across locations and times of day, ranging from $\approx 92\%$ to $\approx 99\%$ for top-3—the implications of this will be further explored in Section VIII-A-2.

What is the gain of learning from multiple scenarios? Even more interesting than the consistency of the DNN performance is the impact of combining scenarios on that performance. The bars of “Day Combined” and “Total” in Fig. 6 indicate two intertwined facts: i) A model can be shared across environments and ii) combining data samples could help improve the top-1 performance on each scenario. When the DNN is trained on a combined dataset (whether combining scenarios having the same time of day or combining all), it achieves a top-1 performance that is *better* than the top-1 performance of three of the individual scenarios. It is observed that in the “Total” case, the top-1 accuracy is higher than those of individual scenarios, i.e., scenarios 5, 6, and 7. *This not only indicates that a single model could be trained for multiple scenarios simultaneously but that combining scenarios can even help in improving the learning process*. At first glance, one could be tempted to attribute this improvement to the unbalanced data contribution of each scenario, see Table I; scenario 8 with its 1320 validation samples may bias the top-1 performance of the combined dataset. However, this is not the case, and the improvement is a result of the improved learning process for the DNN. Fig. 7(a) corroborates this conclusion. It compares the top-1 accuracy of the DNN when it is trained on each scenario individually and on all four scenarios combined, i.e., training on individual scenarios

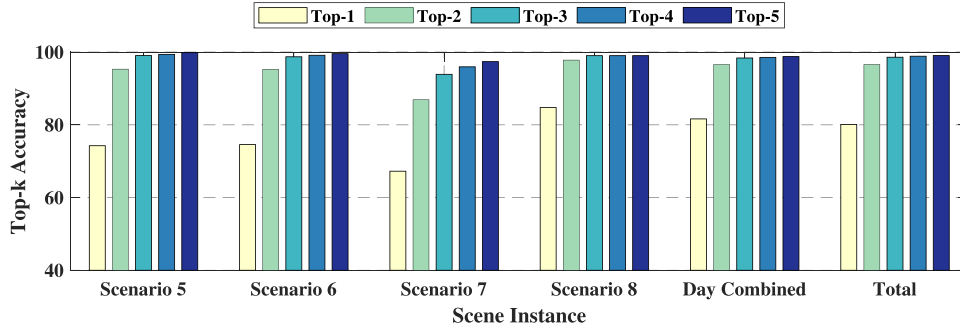


Fig. 6. Top-k accuracies ($k \in \{1, 2, 3, 5\}$) of the proposed DNN across all four single-candidate scenarios and two choices of combining, combining day time scenarios and combining all four scenarios.

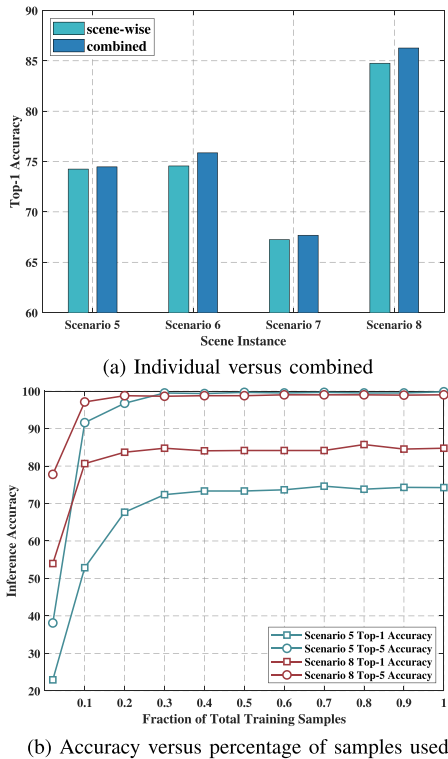


Fig. 7. A deeper dive into the performance of the proposed DNN. (a) highlights the value of combining data samples from different scenarios. (b) sheds some light on how many training samples are required to get the performances in Fig. 6. (a) Individual versus combined. (b) Accuracy versus percentage of samples used.

separately and testing on individual scenarios or training on all scenarios together and testing on individual scenarios.

How many training samples are needed? This interesting question could be seen as a natural follow-up to the previous discussion, for it ponders the computational cost of that performance. Fig. 7(b) provides an answer to that question. It plots the top-1 and top-5 accuracies versus the number of training samples used for two scenarios. An obvious observation from the figure is that approximately 30% of the total training samples is enough to achieve the reported top-1 performances in Fig. 6 and even less than 30% is needed for top-5. Under the surface of this observation lies a more interesting takeaway. The figure

consolidates the earlier conclusion on the role of combining scenarios in achieving improved performance. Adding more data samples (more than the 30%) does not have much of an impact on the performance of the DNN. This means the improvement observed after combining the scenarios is actually a consequence of an improved learning process and not an increased number of data samples.

2) *Wireless Communication Perspective:* This perspective focuses on the implications of the performance of the proposed DNN on the wireless system. More specifically, it attempts to answer two critical questions: i) What are the implications of predicting the wrong beam? Moreover, ii) how much of an impact does mis-prediction have on the wireless system?

What are the implications of mis-predictions? The previous results in Fig. 6 indicate that the prediction of the proposed solution deviates from the optimal beam between $\approx 16\%$ to $\approx 33\%$ of the times (based on top-1 accuracy). This may seem concerning at first, yet, as discussed earlier, a closer look at top-3 or top-5 accuracies shows a much slimmer margin of error, $\approx 8\%$ to $\approx 1\%$. This means that a wireless system may not need to rely exclusively on the proposed DNN, obsoleting classical beam training. Instead, *the system could use the top-3 or top-5 beams in conjunction with some lightweight beam training*; it trains the wireless user on the predicted top-3 or top-5 beams to determine the optimal one, which adds a level of robustness to the system operation.

What if only the top-1 prediction is used? This is an interesting question as it encourages a dive into the effect of mis-prediction on the wireless system. Fig. 8 is one way to address that question. Its subfigures are obtained on the case of training and testing the DNN on all scenarios at the same time (case of “Total” in Fig. 6). The confusion matrix, Fig. 8(a), suggests that even when the DNN misses, its top-1 prediction is likely to be one of the neighboring beams, e.g., if the optimal beam is 15, the DNN is most likely to predict beams between 14 and 16. Below the surface, such mis-prediction is not costly; neighboring beams are expected to achieve reasonable wireless performance. This is verified by Fig. 8(b). It shows a scatter plot for the top-1 received power versus the ground-truth received power. The figure indicates that the neighborhood of the optimal beam registers a similar level of received power. Hence, predicting

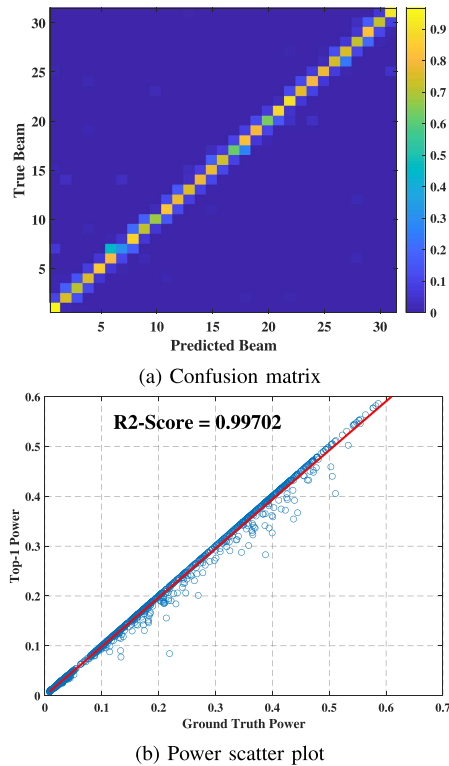


Fig. 8. (a) shows the confusion matrix of beam prediction. (b) illustrates the relation between the groundtruth received power and the top-1 predicted received power. (a) Confusion matrix. (b) Power scatter plot.

a neighboring beam achieves almost the same received power as the optimal beam. This is quantified on the figure using the R^2 -score.

B. Multi-Candidate Beam Prediction

This section focuses on the practical case of beam prediction in multi-candidate settings. As mentioned in Section V-C, the proposed solution has two components, transmitter identification, and beam prediction. Hence, the performance evaluation here is divided into three subsections. The first one evaluates the performance of the transmitter identification component, and the other two evaluate the beam prediction performance in a similar way to that presented in Section VIII-A, namely from machine learning and wireless communication perspectives.

1) *Transmitter Identification*: As described in Section V-C, this component performs scene analysis and object-role identification. In particular, it attempts to identify the target user in the visual scene (from the other objects/distractors). The performance of transmitter identification is evaluated on the multi-candidate dataset before the component is integrated with beam prediction. Fig. 9 shows seven confusion matrices for transmitter identification, four for individual scenarios and three for different combined scenarios. Overall, the matrices display high true positive and negative rates, which do not go below $\approx 90\%$. This indicates a good precision-recall performance regardless of the location or time of day a scenario represents. For instance, when all four scenarios are combined (case of “All” in

the figure), the proposed DNN achieves a precision of $\approx 91\%$ at a recall of 93%. Such good precision-recall performance translates to high confidence in the proposed DNN to identify transmitters in various wireless environments and lighting conditions.

2) *Machine Learning Perspective*: The two-component DNN is evaluated on the multi-candidate dataset, and the focus of this evaluation is still the same as that of Section VIII-A-1.

Could the DNN have consistent performance across scenarios? Addressing this question is even more interesting in multi-candidate settings than it is in single-candidate settings, as they are the epitome of practical wireless environments. Fig. 10 shows the top- k beam prediction accuracies over all four scenarios and three different combinations of them. The first thing one could observe there is a slight performance degradation across all variants of the top- k metric. The primary reason behind the reduction in model performance is *the multi-candidate nature of these scenarios* (i.e., various relevant objects appear in the RGB image). It makes the beam-prediction task much more challenging than single-candidate. Identifying the user (transmitter), in these cases, is no longer straightforward, as evident from the confusion matrices in Fig. 9. However, a closer look at the top-3 and 5 accuracies highlights the following. Compared to the single-candidate settings, top-3 accuracies fluctuate within a range of 10% (a 3% increase over that in single-candidate), and top-5 accuracies register a range of $\approx 5\%$ (a 4% increase over that of single-candidate). *These numbers, overall, are quite encouraging because they indicate a fairly consistent performance considering the increased challenge in multi-candidate settings.*

What is the gain of learning from multiple scenarios? Combining scenarios and training the proposed DNN does not result in the same improvement in performance seen in the single-candidate case. In fact, combining seems to yield the same or slightly degraded performance as opposed to individual scenarios, as shown in the bars “Day Comb.,” “Night Comb.,” and “Total” of Fig. 10. The reason for this could be traced back to the bottleneck of the proposed DNN, which is the transmitter identification component. Owing to its reliance on position data that are inherently noisy⁴ and the presence of multiple candidate transmitters in an image, a transmitter might not be correctly identified. Despite how rarely this happens, when it does, it leads to significant mis-prediction. This clarifies that mis-identifying the transmitter often leads to critical performance loss as the predicted beam is not in the neighborhood of the optimal one.

How many training samples are needed? The previous discussion has extended the findings from single-candidate beam prediction to multi-candidate cases. More specifically, the proposed DNN displays consistent performance across different scenarios, and the model can be trained on multiple scenarios simultaneously. In parallel to the single-candidate discussion, it is important to establish how many training samples are required to get those results. Fig. 11(b) presents two examples for training on scenarios 1 and 4. The top-1 accuracies for both examples require approximately 50% of the training samples in each scenario to achieve the reported performance in Fig. 10. Compared

⁴GPS data do not guarantee zero positioning error. They provide latitude and longitude information within a certain error range.




Fig. 9. The confusion matrices on all case studies in the multi-candidate settings. Each matrix quantifies the likelihood of identifying a transmitter in a group of candidates with similar visual traits.


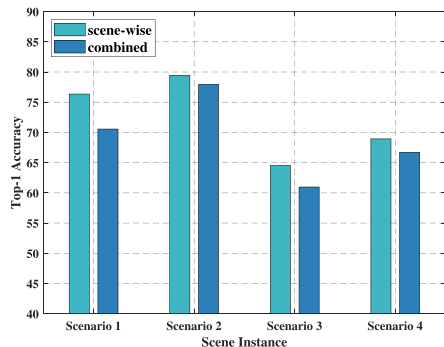
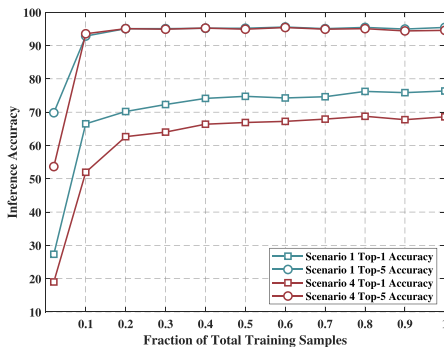


Fig. 10. Top-k accuracies ($k \in \{1, 2, 3, 5\}$) of the proposed DNN across all four multi-candidate scenarios and three choices of combining, combining day time scenarios, night time scenarios, and combining all four.



(a) Individual versus combined



(b) Accuracy versus percentage of samples used

Fig. 11. A deeper dive into the performance of the proposed DNN. (a) explores the impact of combining data samples from different scenarios. (b) sheds some light on how many training samples are required to get the performances in Fig. 6. (a) Individual versus combined. (b) Accuracy versus percentage of samples used.

with the need for 30% of the samples in the single-candidate settings, they highlight the difficulty of the beam prediction in multi-candidate settings.

What is the impact of transmitter identification accuracy on beam prediction accuracy? The proposed multi-candidate beam prediction solution involves first utilizing the visual data and position data to identify the transmitter (user) in the scene, i.e., extract the bounding box coordinate of the transmitter. These bounding box coordinates are then provided as input to the next stage, where we predict the optimal beam index. However, the transmitter identification stage is highly dependent on the quality of the GPS data. Therefore, to truly study the impact of noisy GPS data and the transmitter identification stage on our overall beam prediction accuracy, we conduct a new experiment. In this experiment, instead of inputting the predicted bounding box coordinates, we provide the ground-truth coordinates of the transmitter. This removes the variability of the first stage. Fig. 12 presents a bar plot comparing the top-1 accuracy of our approach when using the ground-truth bounding box coordinates (Ground-truth Bbox) and the predicted bounding box coordinates (Predicted Bbox) for each multi-candidate scenario. The results show that, as expected, the beam prediction performance increases by approximately 1%-10% for the scenarios when using the ground-truth bounding box coordinates. The most significant improvement is observed in scenario 3, where the top-1 accuracy increases from 64.55% to 71.51%, a difference of 10.78%. Scenario 4 also shows a notable improvement of 5.41%, with the top-1 accuracy increasing from 68.93% to

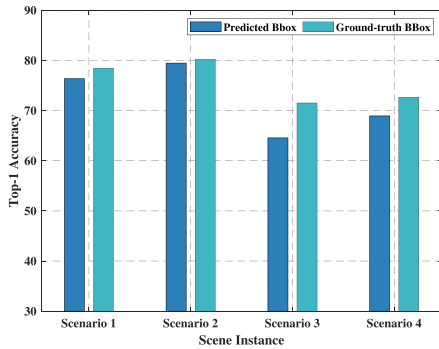


Fig. 12. Comparison of top-1 accuracy for beam prediction using predicted bounding boxes (Predicted Bbox) versus ground-truth bounding boxes (Ground-truth BBox) across four different scenarios.

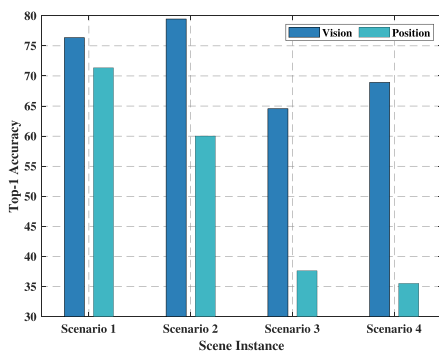


Fig. 13. Top-1 accuracy comparison between the proposed vision and position-based approach and the position-only solution for beam prediction in multi-candidate scenarios across four scene instances.

72.66%. Scenarios 1 and 2 exhibit smaller improvements of 2.74% and 0.94%, respectively. These findings highlight two important points: i) accurate position data can help improve the top-1 beam prediction accuracy, and ii) even with noisy position data, the transmitter identification stage and beam prediction stage perform well, and the increase in accuracy is marginal in certain scenarios. This suggests that while the quality of GPS data plays an important role, the proposed approach is robust enough to handle noisy position data.

How does the proposed multi-modal approach compare to the position-only solution in multi-candidate scenarios? We propose a novel approach that combines both vision and position information. Our method first uses the position data to identify the correct transmitter among the candidates and then predicts the optimal beams based on the bounding box coordinates of the identified transmitter. To evaluate the performance of our proposed approach, we compare its top-1 accuracy to that of the position-only solution adapted from the previously published study [37] across four diverse scenarios, as shown in Fig. 13. In [37], the authors employed a feed-forward neural network with three hidden layers consisting of 256 nodes each. The network was trained for 60 epochs using cross-entropy loss and Adam optimizer with an initial learning rate of 0.01 and a multi-step learning rate scheduler that multiplied all learning parameters by 0.2 at epochs 20 and 40. It is important to note

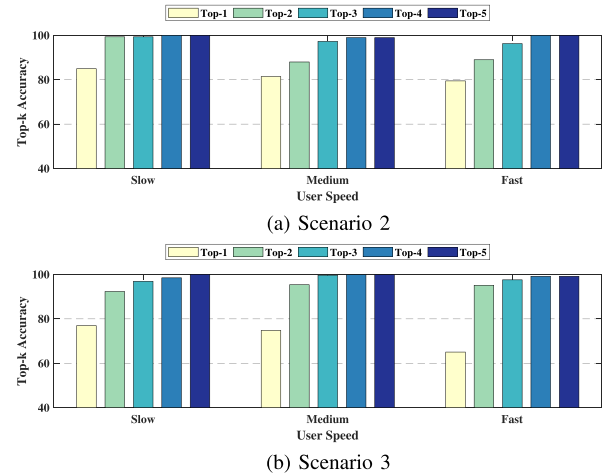


Fig. 14. Beam prediction accuracy versus transmitter speed categories for (a) Scenario 2 and (b) Scenario 3. The plots show Top-1 to Top-5 accuracies for low, medium, and high-speed transmitters, demonstrating the impact of speed on prediction performance in different environments.

that the results presented here for the position-only solution are based on the same scenario data and with a codebook size of 32 beams, identical to the setup used in our paper. In scenario 1, our vision and position-based approach achieves a top-1 accuracy of 76.35%, outperforming the position-only accuracy by 7.02%. The performance gap becomes more significant in scenario 2, where our approach surpasses the position-only accuracy by 32.36%. The advantage of incorporating visual information becomes even more evident in the challenging scenarios 3 and 4. In these environments, the position data alone struggles to accurately predict the optimal beams. In contrast, our vision and position-based approach attain much higher top-1 accuracies, representing substantial improvements over the position-only solution. These results highlight the robustness of our proposed approach in handling multi-candidate scenarios.

What is the impact of transmitter speed on beam prediction accuracy, and how can we address real-time inference challenges? Given the dynamic nature of wireless environments, it is important to consider the impact of transmitter speed on beam prediction accuracy. To analyze this, we estimate the speed of each transmitter by calculating the distance traveled between three consecutive samples: the previous sample, the current sample, and the next sample. We use the Haversine formula to compute the distance between the latitude and longitude coordinates of these samples. Considering a sampling time of 100 ms between each sample, we calculate the speed by dividing the total distance traveled by 200 ms (the time span of the three samples). For each scenario, we calculate the speed mean (μ) and standard deviation (σ). Using these statistics, we divide the transmitters into three speed categories: i) low-speed transmitters with speeds less than or equal to $\mu - \sigma$; ii) high-speed transmitters with speeds greater than or equal to $\mu + \sigma$; and iii) medium-speed transmitters with speeds between those of low- and high-speed transmitters. Fig. 14 presents the beam prediction accuracy versus the transmitter speed category for scenarios 2 and 3. In both scenarios, we observe a trend of

decreasing top-1 accuracy as transmitter speed increases. For scenario 2, the accuracy ranges from 84.89% for low-speed to 79.47% for high-speed transmitters, while for scenario 3, it ranges from 76.92% to 65.04%. The impact of speed on accuracy is more pronounced in scenario 3, particularly for high-speed transmitters. Despite the decrease in top-1 accuracy, the top-5 accuracy remains consistently high across all speed categories in both scenarios, often reaching or approaching 100%. This indicates that our approach maintains robustness even for faster-moving transmitters, as the optimal beam is typically included within the top 5 predictions. It is important to note that real-time inference presents additional challenges beyond accuracy. The entire process—from data capture and processing to inference—introduces latency, which is further exacerbated by increased mobility requiring more frequent beam training. To address this latency concern and achieve real-time or near real-time beam predictions, a promising solution is to proactively predict future beams. By leveraging additional sensing modalities like vision and position to understand mobility patterns and anticipate probable blockages, we could predict optimal beams 300–500 ms in the future. This proactive approach has the potential to mitigate latency issues and maintain high accuracy even in highly mobile scenarios.

3) *Wireless Communication Perspective:* Next, we investigate the performance from wireless perspective and draw some insights about the system operation.

What are the implications of mis-predictions? The findings in Section VIII-B-2 indicate a slight dip in the beam prediction accuracies, especially top-1 accuracy. The implications of that on the wireless system performance could be explored further with the confusion matrix in Fig. 16(a). The first thing to observe from the figure is that most predictions identify the optimal beam or a beam in its neighborhood. *This means that robust beam prediction could still be achieved with lightweight beam training and the top-5 predictions of the DNN*, as in the single-candidate settings. However, in contrast to the confusion matrix in Fig. 8(a), there is a higher likelihood of seeing predictions far from the ground truth. This is evident from the number of bright spots scattered around the diagonal of the matrix. They are a direct consequence of the catastrophic mis-predictions resulting from misidentifying the transmitter.

How does the codebook size impact beam prediction accuracy? This question aims to understand the effect of varying the number of beams in the codebook on the performance of our proposed DNN. While our primary results utilize a 32-beam codebook, we conducted an additional analysis to compare this with a 64-beam codebook. Fig. 15(a) addresses this question by presenting the top-1 accuracy for both codebook sizes across scenarios 1–4. The figure reveals a consistent difference in accuracy between the two codebook sizes. For example, in scenario 1, the top-1 accuracy decreases from 76.35% with 32 beams to 58.01% with 64 beams. This reduction in accuracy is observed across all tested scenarios. The performance difference can be attributed to the reduced complexity of the classification task with fewer beams. However, it is crucial to note that this difference in beam index prediction accuracy does

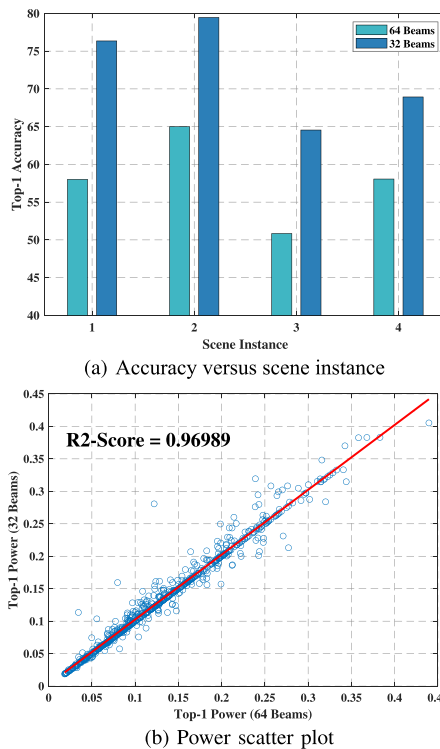


Fig. 15. (a) Comparison of beam prediction accuracy for codebook sizes 32 and 64 across four scenarios. (b) comparing the receive power of the top-1 predicted beam for both codebook sizes. (a) Confusion matrix. (b) Power scatter plot.

not necessarily translate to a significant difference in receive power performance. To investigate this further, we conducted an additional analysis comparing the receive power of the top-1 predicted beam for both codebook sizes, focusing on scenario 1 as a representative case. As shown in Fig. 15(b), there is a strong positive correlation (R^2 -Score = 0.96989) between the power values predicted using 32-beam and 64-beam codebooks. This high correlation indicates that despite the difference in beam index accuracy, both codebook sizes tend to predict similar receive power levels. This analysis highlights the complex trade-offs between prediction accuracy, computational efficiency, and beam resolution in practical implementations.

What if only the top-1 prediction is used? The impact of mis-predictions in multi-candidate settings becomes clear when a wireless system relies only on top-1 predictions. Figs. 10 and 16(a) have hinted to the phenomenon, but neither has explored its direct impact on the wireless system performance, and as such, Fig. 16(b) attempts to bridge that gap and round out this analysis. It shows a power scatter plot similar to that presented in Fig. 8(b). The plot shows a wider cluster of blue points, resulting in a smaller R^2 -score compared to the single-candidate settings. This score directly reflects the impact of catastrophic mis-prediction; the phenomenon produces beam predictions with very low received power. For instance, the lower half of the y-axis (Top-1 Power) shows some examples where the received power of a top-1 predicted beam is

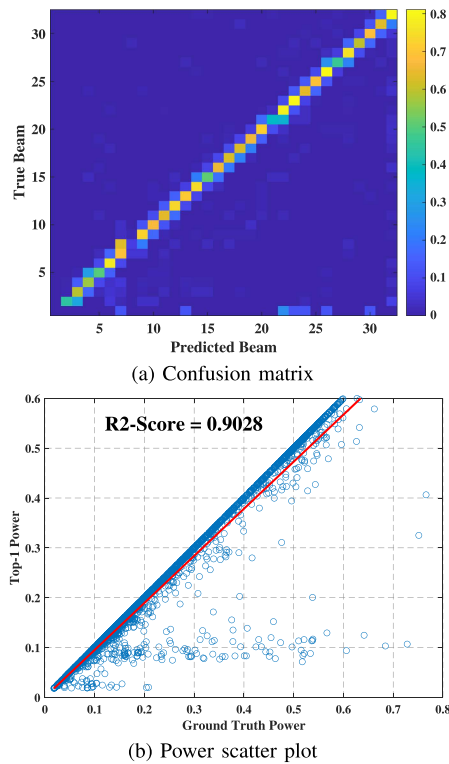


Fig. 16. (a) shows the confusion matrix of beam prediction. (b) illustrates the relation between the groundtruth receive power and the top-1 predicted receive power. (a) Confusion matrix. (b) Power scatter plot.

≈ 0.1 . In contrast, the ground truth beam produces received power in the range of 0.3 to 0.7, i.e., ground truth power is 3 to 7 times higher than the predicted beam. *Such observation emphasizes the importance of augmenting a ViWiComm system with lightweight beam training in practical wireless environments.*

IX. CONCLUSION

This paper presents a machine learning framework specifically designed to address the challenges of realistic scenarios in highly mobile mmWave/sub-THz wireless communication systems with multiple probable transmitting candidates. The proposed solution utilizes visual and positional data to predict optimal beam indices, effectively reducing the beam training overhead associated with adjusting narrow beams of large antenna arrays. Experimental evaluation on the DeepSense 6G dataset demonstrates that the proposed solution can achieve between 67 – 84% top-1 and close to 100% top-5 beam prediction accuracy for single-user scenarios and 65 – 80% top-1 and approximately 95% top-5 beam prediction accuracy for multi-object scenarios while accurately identifying the probable transmitting candidate with over 93% accuracy. An important area for future work is ensuring the generalizability of the trained model to unseen scenarios, further enhancing its applicability in real-world mmWave/sub-THz wireless communication networks.

REFERENCES

- [1] A. Alkhateeb, S. Alex, P. Varkey, Y. Li, Q. Qu, and D. Tujkovic, "Deep learning coordinated beamforming for highly-mobile millimeter wave systems," *IEEE Access*, vol. 6, pp. 37328–37348, 2018.
- [2] T. S. Rappaport et al., "Wireless communications and applications above 100 GHz: Opportunities and challenges for 6G and beyond," *IEEE Access*, vol. 7, pp. 78729–78757, 2019.
- [3] G. J. Sutton et al., "Enabling technologies for ultra-reliable and low latency communications: From PHY and MAC layer perspectives," *IEEE Commun. Surveys Tut.*, vol. 21, no. 3, pp. 2488–2524, Third Quarter 2019.
- [4] J. G. Andrews et al., "What will 5G be?," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, Jun. 2014.
- [5] J. G. Andrews, T. Bai, M. N. Kulkarni, A. Alkhateeb, A. K. Gupta, and R. W. Heath, "Modeling and analyzing millimeter wave cellular systems," *IEEE Trans. Commun.*, vol. 65, no. 1, pp. 403–430, Jan. 2017.
- [6] M. Alrabeiah, A. Hredzak, Z. Liu, and A. Alkhateeb, "ViWi: A deep learning dataset framework for vision-aided wireless communications," in *Proc. IEEE 91st Veh. Technol. Conf.*, 2020, pp. 1–5.
- [7] M. Alrabeiah, A. Hredzak, and A. Alkhateeb, "Millimeter wave base stations with cameras: Vision-aided beam and blockage prediction," in *Proc. IEEE 91st Veh. Technol. Conf.*, 2020, pp. 1–5.
- [8] G. Charan, M. Alrabeiah, and A. Alkhateeb, "Vision-aided 6G wireless communications: Blockage prediction and proactive hand-off," *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 10193–10208, Oct. 2021.
- [9] M. Alrabeiah, U. Demirhan, A. Hredzak, and A. Alkhateeb, "Vision aided URLLC communications: Proactive service identification and coexistence," in *Proc. IEEE 54th Asilomar Conf. Signals Syst. Comput.*, 2020, pp. 174–178.
- [10] M. Alrabeiah, J. Booth, A. Hredzak, and A. Alkhateeb, "ViWi vision-aided mmWave beam tracking: Dataset, task, and baseline solutions," 2020, *arXiv:2002.02445*.
- [11] G. Charan and A. Alkhateeb, "User identification: The key enabler for multi-user vision-aided wireless communications," *IEEE Open J. Commun. Soc.*, 2023.
- [12] G. Charan, T. Osman, A. Hredzak, N. Thawdar, and A. Alkhateeb, "Vision-position multi-modal beam prediction using real millimeter wave datasets," in *Proc. 2022 IEEE Wireless Commun. Netw. Conf.*, 2022, pp. 2727–2731.
- [13] G. Charan et al., "Towards real-world 6G drone communication: Position and camera aided beam prediction," in *Proc. 2022 IEEE Glob. Commun. Conf.*, 2022, pp. 2951–2956.
- [14] U. Demirhan and A. Alkhateeb, "Radar aided 6G beam prediction: Deep learning algorithms and real-world demonstration," in *Proc. 2022 IEEE Wireless Commun. Netw. Conf.*, 2022, pp. 2655–2660.
- [15] S. Jiang, G. Charan, and A. Alkhateeb, "Lidar aided future beam prediction in real-world millimeter wave V2I communications," *IEEE Wireless Commun. Lett.*, vol. 12, no. 2, pp. 212–216, Feb. 2023.
- [16] Z. Ying, H. Yang, J. Gao, and K. Zheng, "A new vision-aided beam prediction scheme for mmwave wireless communications," in *Proc. IEEE 6th Int. Conf. Comput. Commun.*, 2020, pp. 232–237.
- [17] Y. Tian, G. Pan, and M. S. Alouini, "Applying deep-learning-based computer vision to wireless communications: Methodologies, opportunities, and challenges," *IEEE Open J. Commun. Soc.*, vol. 2, pp. 132–143, 2021.
- [18] M. Alrabeiah and A. Alkhateeb, "Deep learning for TDD and FDD massive mimo: Mapping channels in space and frequency," in *Proc. IEEE 53rd Asilomar Conf. Signals Syst. Comput.*, 2019, pp. 1465–1470.
- [19] A. Alkhateeb et al., "Deepsense 6G: A large-scale real-world multi-modal sensing and communication dataset," *IEEE Commun. Mag.*, vol. 61, no. 9, pp. 122–128, Sep. 2023.
- [20] J. Wang et al., "Beam codebook based beamforming protocol for multi-gbps millimeter-wave WPAN systems," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1390–1399, Oct. 2009.
- [21] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh, "Millimeter wave beamforming for wireless backhaul and access in small cell networks," *IEEE Trans. Commun.*, vol. 61, no. 10, pp. 4391–4403, Oct. 2013.
- [22] M. Alrabeiah and A. Alkhateeb, "Deep learning for mmwave beam and blockage prediction using sub-6GHz channels," *IEEE Trans. Commun.*, vol. 68, no. 9, pp. 5504–5518, Sep. 2020.
- [23] Y. Wang, A. Klautau, M. Ribero, A. C. K. Soong, and R. W. Heath, "Mmwave vehicular beam selection with situational awareness using machine learning," *IEEE Access*, vol. 7, pp. 87479–87493, 2019.

- [24] A. Alkhateeb, S. Jiang, and G. Charan, "Real-time digital twins: Vision and research directions for 6G and beyond," *IEEE Commun. Mag.*, 2023.
- [25] V. M. De Pinho, M. L. R. De Campos, L. U. Garcia, and D. Popescu, "Vision-aided radio: User identity match in radio and video domains using machine learning," *IEEE Access*, vol. 8, pp. 209619–209629, 2020.
- [26] N. Abuzainab, M. Alrabeiah, A. Alkhateeb, and Y. E. Sagduyu, "Deep learning for THz drones with flying intelligent surfaces: Beam and handoff prediction," in *Proc. 2021 IEEE Int. Conf. Commun. Workshops*, 2021, pp. 1–6.
- [27] G. Charan, M. Alrabeiah, and A. Alkhateeb, "Vision-aided dynamic blockage prediction for 6G wireless communication networks," in *Proc. 2021 IEEE Int. Conf. Commun. Workshops*, 2021, pp. 1–6.
- [28] M. Alrabeiah, Y. Zhang, and A. Alkhateeb, "Neural networks based beam codebooks: Learning mmwave massive MIMO beams that adapt to deployment and hardware," *IEEE Trans. Commun.*, vol. 70, no. 6, pp. 3818–3833, 2022.
- [29] Y. Zhang, M. Alrabeiah, and A. Alkhateeb, "Reinforcement learning of beam codebooks in millimeter wave and terahertz MIMO systems," *IEEE Trans. Commun.*, vol. 70, no. 2, pp. 904–919, 2021.
- [30] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>
- [31] A. Alkhateeb, G. Leus, and R. W. Heath, "Limited feedback hybrid precoding for multi-user millimeter wave systems," *IEEE Trans. Wireless Commun.*, vol. 14, no. 11, pp. 6481–6494, Nov. 2015.
- [32] C. M. Bishop, *Pattern Recognition and Machine Learning*. Berlin, Germany: Springer, 2006.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [34] K. He, G. Gkioxari, P. Dollar, and R. Girshick, "Mask R-CNN," in *Proc. Int. Conf. Comput. Vis.*, 2017, pp. 2961–2969.
- [35] J. Dai, K. He, and J. Sun, "Instance-aware semantic segmentation via multi-task network cascades," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3150–3158.
- [36] T. K. Keyes, "Applied regression analysis and multivariable methods," *Technometrics*, vol. 43, no. 1, pp. 101–101, 2001, doi: [10.1198/tech.2001.s552](https://doi.org/10.1198/tech.2001.s552).
- [37] J. Morais, A. Bchboodi, H. Pezeshki, and A. Alkhateeb, "Position-aided beam prediction in the real world: How useful GPS locations actually are?," in *Proc. 2023 IEEE Int. Conf. Commun.*, 2023, pp. 1824–1829.



Gouranga Charan received the B.Tech. degree in instrumentation engineering from Indian Institute of Technology Kharagpur, India, in 2015, and the M.S. and Ph.D. degrees in electrical engineering from Arizona State University, Tempe, AZ, USA, in 2021 and 2024, respectively. Between 2015 and 2017, he was an IC Design Engineer with Broadcom Inc., Bengaluru, India. He is currently a Postdoctoral Research Scholar with Arizona State University, USA. He has completed three research internships with Nokia Bell Labs in Murray Hills, NJ, META (formerly Facebook) in Redmond, WA, and Apple, Seattle, WA. His research interests include studying the different applications of deep learning in computer vision, wireless communications, and wireless sensing, with my primary focus on sensing-aided wireless communication.



Muhammad Alrabeiah received the Bachelor of Engineering degree from King Saud University (KSU), Riyadh, Saudi Arabia, in 2010, the Master of Applied Sciences degree in electrical engineering from McMaster University, Hamilton, ON, Canada, in 2015, and the Ph.D. degree in electrical engineering from the school of Electrical, Computer, and Energy Engineering Arizona State University, Tempe, AZ, USA, in 2021. He holds a Lecturer position with the Electrical Engineering Department, KSU, and he was the recipient of the KSU's scholarship for graduate-studies abroad. His research interests include Machine Learning, with his recent work focusing on the applications of deep learning in computer vision, wireless communications, and wireless sensing.



Tawfik Osman received the B.S. degree in electrical and electronic engineering from Ashesi University, Berekuso, Ghana, in 2020 and the M.S.E. degree in electrical engineering from Arizona State University, Tempe, AZ, USA, (under a joint program), in 2021. He is currently working toward the Ph.D. degree with Arizona State University under the supervision of Prof. Ahmed Alkhateeb. His research interests include wireless communications, wireless prototyping with software-defined radio, and machine learning.



Ahmed Alkhateeb received the B.S. (distinction with honor) and the M.S. degrees in electrical engineering from Cairo University, Giza, Egypt, in 2008 and 2012, respectively, and the Ph.D. degree in electrical engineering from The University of Texas at Austin, Austin, TX, USA, in 2016. Between 2016 and 2017, he was a Wireless Communications Researcher with the Connectivity Lab, Facebook, in Menlo Park, CA. He joined Arizona State University in 2018, where he is currently an Assistant Professor with the School of Electrical, Computer and Energy Engineering. He has held R&D internships with FutureWei Technologies, Chicago, IL, USA, and Samsung Research America, Dallas, TX, USA. His research interests are in the broad areas of wireless communications, communication theory, signal processing, machine learning, and applied math. He was the recipient of the 2012 MCD Fellowship from The University of Texas at Austin, 2016 IEEE Signal Processing Society Young Author Best Paper Award for his work on hybrid precoding and channel estimation in millimeter wave communication systems, and the 2021 NSF CAREER Award.