Probabilistic Position-Aided Beam Selection for mmWave MIMO Systems

Joseph K. Chege[1], Arie Yeredor[2], and Martin Haardt[1]
[1]Communications Research Laboratory, Ilmenau University of Technology, Ilmenau, Germany
[2]School of Electrical Engineering, Tel Aviv University, Tel Aviv, Israel
Email: {joseph.chege, martin.haardt}@tu-ilmenau.de, ariey@tauex.tau.ac.il

Abstract—Millimeter-wave (mmWave) MIMO systems rely on highly directional beamforming to overcome severe path loss and ensure robust communication links. However, selecting the optimal beam pair efficiently remains a challenge due to the large search space and the overhead of conventional methods. This paper proposes a probabilistic position-aided beam selection approach that exploits the statistical dependence between user equipment (UE) positions and optimal beam indices. We model the underlying joint probability mass function (PMF) of the positions and the beam indices as a low-rank tensor and estimate its parameters from training data using Bayesian inference. The estimated model is then used to predict the best (or a list of the top) beam pair indices for new UE positions. The proposed method is evaluated using data generated from a state-of-the-art ray tracing simulator and compared with neural network-based and fingerprinting approaches. The results show that our approach achieves a high data rate with relatively few training samples and a significantly reduced beam search space. These advantages render it a promising solution for practical mmWave MIMO deployments, reducing the beam search overhead while maintaining a reliable connectivity.

Index Terms—Beam selection, millimeter-wave (mmWave), position-aided, tensors, Bayesian inference, interpretable machine learning (ML).

I. INTRODUCTION

Millimeter-wave (mmWave) communications are a key enabler of next-generation wireless networks, offering high data rates and improved spectral efficiency. However, mmWave signals suffer from high path loss and sensitivity to blockages, necessitating the use of highly directional beamforming to establish reliable communication links [1], [2]. In multiple-input multiple-output (MIMO) systems, selecting the optimal beam pair between the base station (BS) and the user equipment (UE) is critical for maximizing the signal strength and the system performance. For initial beam establishment at the UE, traditional methods include exhaustive and hierarchical beam search. Exhaustive beam search, where the UE measures the quality of all BS beams by sweeping all of its receive beams, suffers from a high overhead and latency, whereas hierarchical beam search [3] reduces the overhead incurred by an exhaustive search, but experiences a degradation in beamforming gain due to its multilevel search [4].

Due to the directional nature of narrow beams and the dependency on line-of-sight (LOS) propagation, knowledge

of UE positions may be exploited to reduce the training overhead in mmWave communications [5]. For example, an inverse fingerprinting method that uses past received power characteristics at UE locations has been shown to reduce beam training overhead in vehicular communications [6]. Machine learning (ML) approaches based on neural networks (NNs) have also been used in this context to learn the complex mapping between UE positions and optimal beam indices [7]. Additional context information such as UE antenna array orientation [8], vehicle size [9], and traffic density [10] has been exploited in ML-based beam selection approaches.

In this paper, we propose a novel and efficient position-aided beam selection method that relies on the statistical relationship between UE positions and optimal beam indices, considering them as random variables drawn from an underlying joint probability mass function (PMF). We model the joint PMF as a low-rank tensor via the canonical polyadic decomposition (CPD), which has been shown to admit a naïve Bayes interpretation in the context of PMF tensor estimation [11]. This significantly reduces the number of parameters, enabling a reliable estimate to be obtained using relatively few training samples. In addition, the parameters of our model consist of probability distributions and are therefore interpretable, in contrast with NN weights and biases. We evaluate the proposed method using realistic data generated from Sionna [12], a state-of-the-art ray tracing simulator. Numerical results demonstrate that our approach significantly reduces the search space while maintaining a higher data rate.

II. SYSTEM MODEL

We consider a mmWave MIMO communication system that operates on the downlink. The communication scenario consists of a fixed BS and a mobile UE, each equipped with a uniform planar array (UPA). The UPAs are placed on the x-y plane and consist of elements $\{M_{\mathrm{T_x}}, M_{\mathrm{T_y}}\}$ for the BS and $\{M_{\mathrm{R_x}}, M_{\mathrm{R_y}}\}$ for the UE. The total number of antenna elements at the BS and UE is $M_{\mathrm{T}} = M_{\mathrm{T_x}} M_{\mathrm{T_y}}$ and $M_{\mathrm{R}} = M_{\mathrm{R_x}} M_{\mathrm{R_y}}$, respectively. The received signal at the UE is given by

$$y_{\mathrm{R}} = \sqrt{P_{\mathrm{T}}} \boldsymbol{w}^{\mathsf{H}} \boldsymbol{H} \boldsymbol{f} s + \boldsymbol{w}^{\mathsf{H}} \boldsymbol{n}, \qquad (1)$$

where $(\cdot)^{\mathsf{H}}$ denotes the conjugate transpose, $\boldsymbol{f} \in \mathbb{C}^{M_{\mathrm{T}}}$ and $\boldsymbol{w} \in \mathbb{C}^{M_{\mathrm{R}}}$ denote the precoding and combining vectors at the BS and the UE, respectively. Furthermore, P_{T} is the transmit

power, $\boldsymbol{H} \in \mathbb{C}^{M_\mathrm{R} \times M_\mathrm{T}}$ is the matrix of channel coefficients, $s \in \mathbb{C}$ is a pilot symbol with unit power, and $\boldsymbol{n} \in \mathbb{C}^{M_\mathrm{R}}$ denotes a zero-mean circularly symmetric complex Gaussian noise vector with variance σ_n^2.

A. Channel Model

We generate the channel coefficients $\boldsymbol{H}$ using Sionna [12], an open-source software for simulating the physical layer of wireless and optical communication systems. In particular, we employ the Sionna ray tracing (RT) module to simulate physically accurate channel realizations for various UE positions within a given propagation environment. Fig. 1 shows the simulation environment (also referred to as the "scene"), which is the area around the Frauenkirche in Munich, Germany. The scene contains objects made up of different materials such as concrete, brick, wood, glass, metal, among others, resulting in a rich and complex propagation environment. The ray tracing tool simulates a number of rays, each corresponding to the various propagation paths between the BS and each UE position. Each path contains associated information such as the angle of departure (AoD), angle of arrival (AoA), the complex path gain, and the delay. The channel matrix is constructed by applying a narrowband geometric channel model, i.e.,

$$\boldsymbol{H} = \sum_{\ell=1}^{L} \alpha_\ell \boldsymbol{a}_\mathrm{R}(\theta_{\mathrm{R},\ell}, \phi_{\mathrm{R},\ell}) \boldsymbol{a}_\mathrm{T}^\mathsf{H}(\theta_{\mathrm{T},\ell}, \phi_{\mathrm{T},\ell}), \qquad (2)$$

where L denotes the number of paths, α_ℓ the complex path gain, $\theta_{\mathrm{R},\ell}$ and $\theta_{\mathrm{T},\ell}$ the elevation AoA and AoD, and $\phi_{\mathrm{R},\ell}$ and $\phi_{\mathrm{T},\ell}$ the azimuth AoA and AoD for the ℓ-th path. We consider $L = 25$ dominant paths consisting of one line-of-sight and $L - 1$ non-line-of-sight paths arising from reflection and diffraction. In addition, $\boldsymbol{a}_\mathrm{R}(\cdot)$ and $\boldsymbol{a}_\mathrm{T}(\cdot)$ are the receive and transmit array steering vectors, respectively. Let us define the spatial frequencies in the x- and y-directions as $\mu_\mathrm{x}^{(i)} = k\Delta_\mathrm{x} \sin(\theta_i) \cos(\phi_i)$ and $\mu_\mathrm{y}^{(i)} = k\Delta_\mathrm{y} \sin(\theta_i) \sin(\phi_i)$, where $i \in \{\mathrm{R}, \mathrm{T}\}$, $k = 2\pi/\nu$ is the wave number, ν the carrier wavelength, whereas Δ_x and Δ_y denote the spacing between the antenna elements in the x- and y-directions, respectively. The steering vector $\boldsymbol{a}_i(\theta_i, \phi_i) \in \mathbb{C}^{M_\mathrm{x} M_\mathrm{y}}$ is defined as

$$\boldsymbol{a}_i(\theta_i, \phi_i) = \frac{1}{\sqrt{M_\mathrm{x} M_\mathrm{y}}} \left(\boldsymbol{a}_\mathrm{x}^{(i)} \otimes \boldsymbol{a}_\mathrm{y}^{(i)} \right),$$

where $\boldsymbol{a}_\mathrm{x}^{(i)} = \begin{bmatrix} 1 & e^{\mathrm{j}\mu_\mathrm{x}^{(i)}} & \cdots & e^{\mathrm{j}(M_\mathrm{x}-1)\mu_\mathrm{x}^{(i)}} \end{bmatrix}^\mathsf{T}$, $\boldsymbol{a}_\mathrm{y}^{(i)} = \begin{bmatrix} 1 & e^{\mathrm{j}\mu_\mathrm{y}^{(i)}} & \cdots & e^{\mathrm{j}(M_\mathrm{y}-1)\mu_\mathrm{y}^{(i)}} \end{bmatrix}^\mathsf{T}$, and $\otimes$ denotes the Kronecker product. Here, M_x and M_y are the number of antenna elements along the x- and y-axis. In this paper, we set $\Delta_\mathrm{x} = \Delta_\mathrm{y} = \nu/2$.

B. Beamforming Codebooks

We consider analog beamforming with one RF chain at the BS and the UE. For simplicity, we employ discrete Fourier transform (DFT) codebooks for precoding and combining. Let I_f and I_w be the number of precoding and combining beams, while $\bar{\theta}$ and $\bar{\phi}$ are the quantized elevation and azimuth angles, obtained by quantizing the spatial frequencies in the x- and
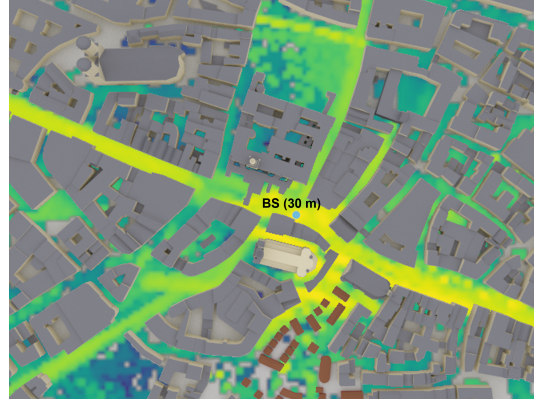

Fig. 1: Illustration of the ray tracing simulation environment overlaid with a color map showing the coverage of the base station (blue dot). The BS antenna array is placed at a height of 30 m above the terrain surface. The environment depicts the area around the Frauenkirche in Munich, Germany, and is set up using Sionna [12].

y-directions such that $2\pi m_\mathrm{x}/M_\mathrm{x} \leq \mu_\mathrm{x} \leq 2\pi(m_\mathrm{x} + M_\mathrm{x} - 1)/M_\mathrm{x}$, $m_\mathrm{x} \in [0, M_\mathrm{x} - 1]$ and $2\pi m_\mathrm{y}/M_\mathrm{y} \leq \mu_\mathrm{y} \leq 2\pi(m_\mathrm{y} + M_\mathrm{y} - 1)/M_\mathrm{y}$, $m_\mathrm{y} \in [0, M_\mathrm{y} - 1]$, respectively. Note that m_x and m_y are integers. Quantizing the spatial frequencies in this way allows the codebooks for the UPAs to be constructed from a 2D DFT matrix. Thus, the precoding and combining codebooks are given by

$$\begin{aligned} \mathcal{C}_\mathrm{T} &= \left\{ \boldsymbol{a}_\mathrm{T}(\bar{\theta}_\mathrm{T}^{(1)}, \bar{\phi}_\mathrm{T}^{(1)}), \ldots, \boldsymbol{a}_\mathrm{T}(\bar{\theta}_\mathrm{T}^{(I_\mathrm{f})}, \bar{\phi}_\mathrm{T}^{(I_\mathrm{f})}) \right\} \\ \mathcal{C}_\mathrm{R} &= \left\{ \boldsymbol{a}_\mathrm{R}(\bar{\theta}_\mathrm{R}^{(1)}, \bar{\phi}_\mathrm{R}^{(1)}), \ldots, \boldsymbol{a}_\mathrm{R}(\bar{\theta}_\mathrm{R}^{(I_\mathrm{w})}, \bar{\phi}_\mathrm{R}^{(I_\mathrm{w})}) \right\}, \end{aligned} \qquad (3)$$

respectively. Let $i_\mathrm{f} \in \mathcal{F} = \{1, \ldots, I_\mathrm{f}\}$ and $i_\mathrm{w} \in \mathcal{W} = \{1, \ldots, I_\mathrm{w}\}$. Given the precoding beam $\boldsymbol{f}_{i_\mathrm{f}} \in \mathcal{C}_\mathrm{T}$ and the combining beam $\boldsymbol{w}_{i_\mathrm{w}} \in \mathcal{C}_\mathrm{R}$, the received signal strength (RSS) for the beam pair indexed by $(i_\mathrm{f}, i_\mathrm{w})$ is

$$\mathrm{RSS}_{i_\mathrm{f}, i_\mathrm{w}} = \left| \sqrt{P_\mathrm{T}} \boldsymbol{w}_{i_\mathrm{w}}^\mathsf{H} \boldsymbol{H} \boldsymbol{f}_{i_\mathrm{f}} s + \boldsymbol{w}_{i_\mathrm{w}}^\mathsf{H} \boldsymbol{n} \right|^2. \qquad (4)$$

In this paper, we set $I_\mathrm{f} = M_\mathrm{T}$ and $I_\mathrm{w} = M_\mathrm{R}$.

C. Dataset Construction

This paper focuses on an offline learning setting where there is a data collection phase and a training phase. To construct the dataset, UE positions are randomly sampled within the simulation environment. The ray tracing tool simulates multi-path propagation to each position and generates the associated MIMO channels. After generating the channel samples, an exhaustive search is performed to identify the beam pair $(\boldsymbol{f}_{i_\mathrm{f}^*}, \boldsymbol{w}_{i_\mathrm{w}^*})$ that maximizes the RSS at each UE position. The optimal beam pair indices $(i_\mathrm{f}^*, i_\mathrm{w}^*)$ are thus given by

$$(i_\mathrm{f}^*, i_\mathrm{w}^*) = \operatorname*{arg\,max}_{i_\mathrm{f} \in \mathcal{F}, i_\mathrm{w} \in \mathcal{W}} \mathrm{RSS}_{i_\mathrm{f}, i_\mathrm{w}}. \qquad (5)$$

The dataset therefore consists of T samples, each containing the position (x, y, z) of the UE and the optimal beam pair indices $(i_\mathrm{f}^*, i_\mathrm{w}^*)$. In practice, UE positions are obtained using Global Positioning System (GPS) sensors which are available in modern user devices. The positions can be shared with the BS using, for example, sub-6 GHz control channels.

III. PROPOSED BEAM SELECTION METHOD

The objective of beam selection in the scenario considered in this paper is to assign the optimal beam pair (or a set of beam pair options) to a UE, given its position. This may be viewed as a classification problem in which the input features are the UE positions and the output targets (or labels) are the optimal beam pair indices. ML classifiers based on neural networks attempt to approximate the (possibly highly non-linear) function that maps the features to the labels. In this paper, we propose to view the UE positions and the beam indices as realizations drawn from an unknown joint probability distribution. In contrast to non-linear function approximation, we attempt to estimate the underlying joint PMF of the UE positions and the beam indices from the available data samples, after which the estimated PMF can be used for beam selection at new UE positions. We refer to our proposed method as Probabilistic Position-Aided Beam Selection (ProPABS).

A. Low-Rank Modeling of the Joint PMF

Define $\mathsf{X} = \{X_\mathrm{x}, X_\mathrm{y}, X_\mathrm{f}, X_\mathrm{w}\}$ consisting of $N = 4$ random variables associated with the x-coordinate and y-coordinate of the UE, the precoding beam index, and the combining beam index, respectively. Note that in our simulations, the z-coordinate of the UE (i.e., its height from the ground) is constant for all UE positions and can thus be omitted from the collection of random variables. Let the support of the variables X_x and X_y be uniformly discretized such that $X_\mathrm{x} \in [1, I_\mathrm{x}]$ and $X_\mathrm{y} \in [1, I_\mathrm{y}]$. Furthermore, for ease of notation, define $\mathcal{N} = \{\mathrm{x}, \mathrm{y}, \mathrm{f}, \mathrm{w}\}$ as a collection of subscripts. Then, given that $X_\mathrm{f} \in [1, I_\mathrm{f}]$ and $X_\mathrm{w} \in [1, I_\mathrm{w}]$, the joint PMF of X can be conveniently represented by a four-way tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{I_\mathrm{x} \times I_\mathrm{y} \times I_\mathrm{f} \times I_\mathrm{w}}$, where each element of $\boldsymbol{\mathcal{X}}$ is the joint probability of a given realization of the four random variables. In particular, letting $i_n = 1, \ldots, I_n$ for $n \in \mathcal{N}$, we have that $\boldsymbol{\mathcal{X}}(i_\mathrm{x}, i_\mathrm{y}, i_\mathrm{f}, i_\mathrm{w}) = \mathsf{Pr}(X_\mathrm{x} = i_\mathrm{x}, X_\mathrm{y} = i_\mathrm{y}, X_\mathrm{f} = i_\mathrm{f}, X_\mathrm{w} = i_\mathrm{w})$.

One may choose to estimate $\boldsymbol{\mathcal{X}}$ using a histogram. In this case, the number of free parameters is $(\prod_{n \in \mathcal{N}} I_n) - 1$, which is at least exponential in N (assuming $I_n > 1, \forall n$). Therefore, to reduce the parameter space, we impose a low-rank CPD model (see, e.g., [13]) on $\boldsymbol{\mathcal{X}}$. This decomposition represents $\boldsymbol{\mathcal{X}}$ in terms of a sum of rank-one tensors, i.e.,

$$\boldsymbol{\mathcal{X}} = \sum_{r=1}^{R} \lambda_r \boldsymbol{A}_\mathrm{x}(:, r) \circ \boldsymbol{A}_\mathrm{y}(:, r) \circ \boldsymbol{A}_\mathrm{f}(:, r) \circ \boldsymbol{A}_\mathrm{w}(:, r), \quad (6)$$

where $\circ$ denotes the outer product, R is the smallest number of rank-one components for which such a decomposition exists, $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_R]^\mathsf{T}$ is a *loading vector* whose elements λ_r scale each rank-one component, while $\boldsymbol{A}_n(:, r)$ is the r-th column of the n-th *factor matrix* $\boldsymbol{A}_n \in \mathbb{R}^{I_n \times R}$.

Kargas *et al.* have shown in [11] that such a decomposition of a PMF tensor can be interpreted as a naïve Bayes model with one latent variable H which takes R states. Under this interpretation, $\lambda_r = \mathsf{Pr}(H = r)$, i.e, the prior probability that the latent variable takes the r-th state, while $\boldsymbol{A}_n(:, r) = \mathrm{p}(X_n \mid H = r)$, i.e., the conditional PMF of the

n-th variable, given that $H = r$. Thus, the decomposition in (6) is subject to nonnegativity ($\boldsymbol{\lambda} > \boldsymbol{0}, \boldsymbol{A}_n \geq \boldsymbol{0}, \forall n$) and sum-to-one ($\boldsymbol{1}^\mathsf{T} \boldsymbol{\lambda} = 1, \boldsymbol{1}^\mathsf{T} \boldsymbol{A}_n = \boldsymbol{1}^\mathsf{T}, \forall n$) constraints (also referred to as probability simplex constraints). The total number of free parameters in the CPD representation is $(R - 1) + \sum_{n \in \mathcal{N}} R(I_n - 1)$, which is linear in N rather than exponential, as is the case with histogram estimation. A distinctive feature of such a nonnegative CPD is that it is essentially unique up to a permutation ambiguity among the rank-one components (see, e.g., [13]).

B. Estimating the Joint PMF

In this subsection, we describe how to estimate the model parameters $\{\boldsymbol{\lambda}, \boldsymbol{A}_\mathrm{x}, \boldsymbol{A}_\mathrm{y}, \boldsymbol{A}_\mathrm{f}, \boldsymbol{A}_\mathrm{w}\}$ from training data using Bayesian inference [14]. Define the dataset consisting of T i.i.d. samples of the UE positions (x, y) and the optimal beam indices $(i_\mathrm{f}^*, i_\mathrm{w}^*)$ as $\mathcal{D} = \{\boldsymbol{d}_t\}_{t=1}^T$, where $\boldsymbol{d}_t = [x_t, y_t, i_{\mathrm{f},t}^*, i_{\mathrm{w},t}^*]^\mathsf{T}$ and $(\cdot)^\mathsf{T}$ denotes the transpose operator. Under the naïve Bayes interpretation, each sample $\boldsymbol{d}_t$ is associated with a realization of the latent variable $H \in [1, R]$. We define a *local* latent variable $\boldsymbol{z}_t = [z_{1,t}, \ldots, z_{R,t}]^\mathsf{T}$ such that $z_{r,t} = 1$ if $H = r$ and $z_{r,t} = 0$ otherwise. It follows that $\mathsf{Pr}(z_{r,t} = 1) = \mathsf{Pr}(H = r) = \lambda_r$.

From the definition of $\boldsymbol{z}_t$, we can write $\mathrm{p}(\boldsymbol{z}_t \mid \boldsymbol{\lambda}) = \prod_{r=1}^R \lambda_r^{z_{r,t}}$. The likelihood of $\mathcal{D}$ given the parameters $\boldsymbol{\Theta} = \{\boldsymbol{Z}, \boldsymbol{\lambda}, \boldsymbol{A}_\mathrm{x}, \boldsymbol{A}_\mathrm{y}, \boldsymbol{A}_\mathrm{f}, \boldsymbol{A}_\mathrm{w}\}$ is

$$\mathrm{p}(\mathcal{D} \mid \boldsymbol{\Theta}) = \mathrm{p}(\mathcal{D} \mid \boldsymbol{Z}, \boldsymbol{A}_\mathrm{x}, \boldsymbol{A}_\mathrm{y}, \boldsymbol{A}_\mathrm{f}, \boldsymbol{A}_\mathrm{w}) \cdot \mathrm{p}(\boldsymbol{Z} \mid \boldsymbol{\lambda}), \quad (7)$$

where $\boldsymbol{Z} = \{\boldsymbol{z}_t\}_{t=1}^T$, $\mathrm{p}(\boldsymbol{Z} \mid \boldsymbol{\lambda}) = \prod_{t=1}^T \mathrm{p}(\boldsymbol{z}_t \mid \boldsymbol{\lambda})$ and

$$\mathrm{p}(\mathcal{D} \mid \boldsymbol{Z}, \boldsymbol{A}_\mathrm{x}, \boldsymbol{A}_\mathrm{y}, \boldsymbol{A}_\mathrm{f}, \boldsymbol{A}_\mathrm{w}) =$$
$$\prod_{t=1}^T \prod_{r=1}^R \left(\boldsymbol{A}_\mathrm{x}(x_t, r) \boldsymbol{A}_\mathrm{y}(y_t, r) \boldsymbol{A}_\mathrm{f}(i_{\mathrm{f},t}^*, r) \boldsymbol{A}_\mathrm{w}(i_{\mathrm{w},t}^*, r) \right)^{z_{r,t}}.$$

To ensure that the probability simplex constraints are satisfied, we assign Dirichlet prior distributions (see, e.g., [15]) for $\boldsymbol{\lambda}$ and $\boldsymbol{A}_n(:, r), \forall n, r$, i.e.,

$$\mathrm{p}(\boldsymbol{\lambda}) \propto \prod_{r=1}^R \lambda_r^{\alpha_{\lambda,r}}, \quad \mathrm{p}\left(\boldsymbol{A}_n(:, r)\right) \propto \prod_{i_n=1}^{I_n} \left(\boldsymbol{A}_n(i_n, r) \right)^{\alpha_{n,r,i_n}},$$

where $\alpha_{\lambda,r}$ and α_{n,r,i_n} are the hyperparameters of the priors.

The posterior distribution of the model parameters given the data is usually found using Bayes' theorem, $\mathrm{p}(\boldsymbol{\Theta} \mid \mathcal{D}) = \mathrm{p}(\mathcal{D}, \boldsymbol{\Theta}) / \int_{\boldsymbol{\Theta}} \mathrm{p}(\mathcal{D}, \boldsymbol{\Theta})$, where

$$\mathrm{p}(\mathcal{D}, \boldsymbol{\Theta}) = \mathrm{p}(\mathcal{D} \mid \boldsymbol{\Theta}) \cdot \mathrm{p}(\boldsymbol{\lambda}) \cdot \prod_{n \in \mathcal{N}} \prod_{r=1}^R \mathrm{p}\left(\boldsymbol{A}_n(:, r)\right). \quad (8)$$

However, in our case, exact Bayesian inference is infeasible due to the high-dimensional integral in the denominator of Bayes' theorem. Therefore, in this work, we adopt the variational Bayes for PMF estimation (VB-PMF) algorithm proposed in [14] to approximate $\mathrm{p}(\boldsymbol{\Theta} \mid \mathcal{D})$. VB-PMF employs co-ordinate ascent variational inference [15], which is guaranteed to converge to a local optimum. Point estimates $\{\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{A}}_\mathrm{x}, \hat{\boldsymbol{A}}_\mathrm{y}, \hat{\boldsymbol{A}}_\mathrm{f}, \hat{\boldsymbol{A}}_\mathrm{w}\}$ are obtained by computing conditional expectations with respect to the estimated posteriors for each parameter.

An important advantage of the VB-PMF algorithm is that the rank R of $\mathcal{X}$, which is usually unknown in practical scenarios, is estimated as part of the inference procedure. In particular, thanks to the sparsity-promoting property of the Dirichlet distribution, setting $\alpha_{\lambda,r} < 1, \forall r$ ensures that the estimated posterior distribution for $\boldsymbol{\lambda}$ (i.e., the PMF of H) is sparse [15]. We choose $\alpha_{\lambda,r} = 10^{-6}, \forall r$. On the other hand, we set $\alpha_{n,r,i_n} = 1, \forall n, r, i_n$ to allow the columns $\boldsymbol{A}_n(:,r)$ (i.e., the conditional PMFs) to explore the space of all possible distributions. The algorithm is initialized with a rank greater than the unknown true rank (e.g., the maximum possible rank required for CPD uniqueness [13]), whereas the hyperparameters of the posterior distributions are initialized randomly. Since the joint PMF tensor is low-rank, some elements of $\hat{\boldsymbol{\lambda}}$ will be close to zero after convergence. The corresponding rank-one components can then be removed, resulting in automatic rank detection. Due to space limitations, we refer the interested reader to [14] for more details.

C. PMF-Based Beam Selection

Given a test UE position $(x_{\text{test}}, y_{\text{test}})$, we would like to predict the optimal precoding and combining beam indices $(\hat{i}_{\text{f}}^*, \hat{i}_{\text{w}}^*)$. The parameters $\{\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{A}}_{\text{x}}, \hat{\boldsymbol{A}}_{\text{y}}, \hat{\boldsymbol{A}}_{\text{f}}, \hat{\boldsymbol{A}}_{\text{w}}\}$ represent an estimate of the joint PMF of all UE positions, precoding beam indices, and combining beam indices. With the joint PMF at hand, we can obtain $(\hat{i}_{\text{f}}^*, \hat{i}_{\text{w}}^*)$ by maximizing the posterior distributions of the beam indices given the test position. For a classification problem, such an estimate, referred to as the maximum *a posteriori* (MAP) estimate, minimizes the probability of misclassification [15]. Thus, we have

$$(\hat{i}_{\text{f}}^*, \hat{i}_{\text{w}}^*) = \arg\max_{i_{\text{f}} \in \mathcal{F}, i_{\text{w}} \in \mathcal{W}} \mathsf{Pr}(i_{\text{f}}, i_{\text{w}} \mid x_{\text{test}}, y_{\text{test}}) \quad (9a)$$

$$= \arg\max_{i_{\text{f}} \in \mathcal{F}, i_{\text{w}} \in \mathcal{W}} \sum_{r=1}^{R} \mathsf{Pr}(H=r)\mathsf{Pr}(x_{\text{test}}, y_{\text{test}}, i_{\text{f}}, i_{\text{w}} \mid r), \quad (9b)$$

where (9b) is obtained after applying Bayes' rule and noting that the maximization is independent of the marginal probability $\mathsf{Pr}(x_{\text{test}}, y_{\text{test}})$. Recall that under the naïve Bayes model, the variables $\{X_{\text{x}}, X_{\text{y}}, X_{\text{f}}, X_{\text{w}}\}$ are independent when conditioned on the latent variable H. Thus, the objective in (9b) is expressed in terms of the estimated parameters as

$$\sum_{r=1}^{R} \hat{\lambda}_r \hat{\boldsymbol{A}}_{\text{f}}(i_{\text{f}}, r) \hat{\boldsymbol{A}}_{\text{w}}(i_{\text{w}}, r) \hat{\boldsymbol{A}}_{\text{x}}(x_{\text{test}}, r) \hat{\boldsymbol{A}}_{\text{y}}(y_{\text{test}}, r). \quad (10)$$

The estimated posterior distribution $\mathsf{p}(X_f, X_w \mid x_{\text{test}}, y_{\text{test}})$ is used to produce a top-N_{b} beam pair list by selecting the indices corresponding to the N_{b} largest posterior probabilities.

IV. RESULTS AND DISCUSSION

We consider the simulated outdoor environment depicted in Fig. 1. The transmit power P_{T}, the carrier frequency f_c, and the carrier bandwidth B are set to $30\,\text{dBm}$, $26\,\text{GHz}$, and $10\,\text{MHz}$, respectively. Furthermore, we assume that the noise variance is given by $\sigma_n^2 = -174 + 10\log_{10} B$ dBm. The BS and the UE are equipped with 8×8 and 2×2 UPAs, respectively. Thus, the number of antennas is $M_{\text{T}} = 64$ and

$M_{\text{R}} = 4$. The positions of the UEs are randomly sampled within the BS coverage area and realistic channel responses between the BS and each position are generated using the Sionna ray tracing module [12]. For each UE position, the RSS is calculated using all precoder-combiner combinations according to (4). The optimal beam pair indices are then obtained using (5).

The dataset $\mathcal{D}$ consists of 2×10^3 samples, of which 80% are used for training and 20% for testing. Moreover, 50 experiments are conducted, each involving the shuffling of the entire dataset and its division into training and testing sets, with the final result being the average across all experiments. For ProPABS, the x- and y-coordinates are discretized into bins of size $5\,\text{m}$, yielding $I_{\text{x}} = 280$ and $I_{\text{y}} = 179$, respectively. Therefore, our model contains 15,719 trainable parameters (cf. Section III-A). We initialize the VB-PMF algorithm with $R = 30$ components. The estimated ranks obtained from 50 experiments, along with their frequency of occurrence, are $\hat{R} = 3$ (4%), $\hat{R} = 4$ (34%), $\hat{R} = 5$ (50%), and $\hat{R} = 6$ (12%). Then, we compare ProPABS with two position-aided beam selection approaches:

1) Inverse multipath fingerprinting [6] builds a database of top beam pair indices for each UE position through exhaustive search. Beam pairs are ranked by their likelihood of being optimal (i.e., highest RSS) and used during testing. A $5\,\text{m}$ bin size is used to discretize UE positions, as in our method.

2) A fully connected neural network (NN) maps UE positions to optimal beam pairs (e.g, [7]). It has three hidden layers with 18, 36, and 50 neurons using sigmoid activations, a 2-neuron input layer, and a 256-neuron output layer, totaling 15,644 trainable parameters—comparable to ProPABS. The NN is trained for 50 epochs using the Adam optimizer [16].

Each beam selection method produces a candidate list $\mathcal{S}$ of the top-N_{b} beam pairs and the pair resulting in the highest RSS is selected, i.e.

$$(\hat{i}_{\text{f}}^*, \hat{i}_{\text{w}}^*) = \arg\max_{(\hat{i}_{\text{f}}, \hat{i}_{\text{w}}) \in \mathcal{S}} \text{RSS}_{\hat{i}_{\text{f}}, \hat{i}_{\text{w}}}. \quad (11)$$

To evaluate the performance of the beam selection methods, we consider the power loss probability metric defined in [6]. Let the set of all possible beam pairs be $\mathcal{B}$. Then, the power loss probability is given by [6]

$$P_{\text{pl}}(c, \mathcal{S}) = \mathsf{Pr}\left[\max_{(i_{\text{f}}, i_{\text{w}}) \in \mathcal{B}} \text{RSS}_{i_{\text{f}}, i_{\text{w}}} > c \cdot \text{RSS}_{\hat{i}_{\text{f}}^*, \hat{i}_{\text{w}}^*}\right] \quad (12)$$

where $c \geq 1$. In our experiments, we evaluate the $0\,\text{dB}$ ($c = 1$) and $3\,\text{dB}$ ($c = 2$) power loss probabilities. In addition, we compute the achievable rate of the selected beam pair, defined as

$$R_{\text{a}} = \log_2\left[1 + \left(\left|\sqrt{P_{\text{T}}} \boldsymbol{w}_{\hat{i}_{\text{w}}^*}^{\mathsf{H}} \boldsymbol{H} \boldsymbol{f}_{\hat{i}_{\text{f}}^*} s\right|^2 / \sigma_n^2\right)\right]. \quad (13)$$

Fig. 2 presents the results arising from our numerical simulations. Fig. 2a shows the average power loss probability as a function of the number of beam pairs searched. ProPABS consistently achieves a lower power loss probability compared
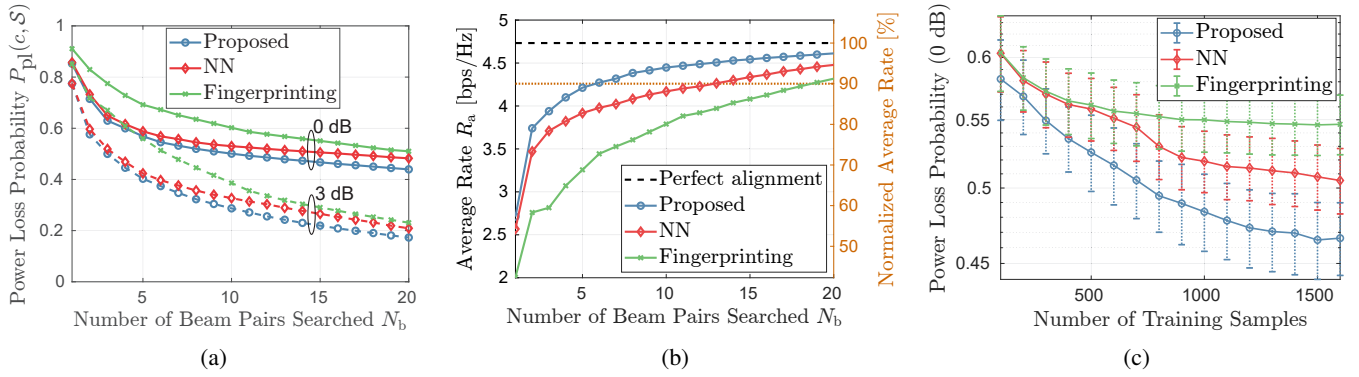
Fig. 2: Performance evaluation of the proposed beam selection method compared to two baselines.

to the baseline approaches. The average training times for ProPABS, the NN, and fingerprinting were 1.2 s, 36.75 s, and 0.23 s, respectively. Fig. 2b presents the average achievable rate and the corresponding average rate normalized by the rate obtained with perfect beam alignment (i.e., using exhaustive search). We observe that ProPABS already achieves about 90 % of the maximum achievable rate by searching only the top $N_{\mathrm{b}} = 6$ beam pairs. Note that for the setup considered, there are $M_{\mathrm{T}} M_{\mathrm{R}} = 256$ possible beam pair combinations, demonstrating that ProPABS greatly reduces the search space while achieving a considerably high average rate. In comparison, the NN and the fingerprinting method need to search around $N_{\mathrm{b}} = 13$ and $N_{\mathrm{b}} = 19$ beam pairs, respectively, to achieve a similar normalized average rate. Fig. 2c shows the average 0 dB power loss probability as a function of the number of training samples, ranging from 100 to 1600. For this evaluation, we set $N_{\mathrm{b}} = 19$ to ensure a normalized average rate greater than 90 % for all beam selection methods, and average over 100 trials. We observe that ProPABS yields the best performance compared to the NN and fingerprinting approaches across all training dataset sizes. Therefore, ProPABS achieves an acceptable average rate with relatively few training samples, demonstrating its efficiency.

V. CONCLUSION

This paper presents an efficient beam selection method for mmWave MIMO systems that employs a low-rank probability mass function (PMF) tensor model with interpretable parameters. By estimating the joint PMF tensor of discretized UE positions and beam indices using a Bayesian approach, the proposed method effectively captures the statistical relationship between them, enabling a reduction in the beam search space while maintaining a high achievable rate. Compared to approaches based on neural networks and fingerprinting, our method achieves a better performance in terms of the power loss probability and the achievable rate with relatively few training samples. The effectiveness of the proposed approach is demonstrated using data generated from a state-of-the-art ray tracing simulator, ensuring realistic evaluation conditions. Since this paper only considers a single stream, an immediate future step is to extend the proposed approach to multi-streaming. Additionally, online beam selection could be con-

sidered to allow continuous adaptation to the communication environment.

REFERENCES

[1] T. S. Rappaport *et al.*, "Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.
[2] W. Hong *et al.*, "The Role of Millimeter-Wave Technologies in 5G/6G Wireless Communications," *IEEE Journal of Microwaves*, vol. 1, no. 1, pp. 101–122, 2021.
[3] Z. Xiao, T. He, P. Xia, and X.-G. Xia, "Hierarchical Codebook Design for Beamforming Training in Millimeter-Wave Communication," *IEEE Transactions on Wireless Communications*, vol. 15, no. 5, pp. 3380–3392, 2016.
[4] M. Q. Khan, A. Gaber, P. Schulz, and G. Fettweis, "Machine Learning for Millimeter Wave and Terahertz Beam Management: A Survey and Open Challenges," *IEEE Access*, vol. 11, pp. 11 880–11 902, 2023.
[5] 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. IEEE International Conference on Communications*, 2023.
[6] V. Va, J. Choi, T. Shimizu, G. Bansal, and R. W. Heath, "Inverse Multipath Fingerprinting for Millimeter Wave V2I Beam Alignment," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 5, pp. 4042–4058, 2018.
[7] Y. Heng and J. G. Andrews, "Machine Learning-Assisted Beam Alignment for mmWave Systems," *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 4, pp. 1142–1155, 2021.
[8] S. Rezaie, E. de Carvalho, and C. N. Manchón, "A Deep Learning Approach to Location- and Orientation-Aided 3D Beam Selection for mmWave Communications," *IEEE Transactions on Wireless Communications*, vol. 21, no. 12, pp. 11 110–11 124, 2022.
[9] Y. Wang, A. Klautau, M. Ribero, M. Narasimha, and R. W. Heath, "MmWave Vehicular Beam Training with Situational Awareness by Machine Learning," in *Proc. IEEE Globecom Workshops*, 2018.
[10] K. Satyanarayana, M. El-Hajjar, A. A. M. Mourad, and L. Hanzo, "Deep Learning Aided Fingerprint-Based Beam Alignment for mmWave Vehicular Communication," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 11, pp. 10 858–10 871, 2019.
[11] N. Kargas, N. D. Sidiropoulos, and X. Fu, "Tensors, Learning, and "Kolmogorov Extension" for Finite-Alphabet Random Vectors," *IEEE Transactions on Signal Processing*, vol. 66, no. 18, pp. 4854–4868, 2018.
[12] J. Hoydis, S. Cammerer, F. Ait Aoudia, A. Vem, N. Binder, G. Marcus, and A. Keller, "Sionna: An open-source library for next-generation physical layer research," *arXiv preprint*, 2022.
[13] T. G. Kolda and B. W. Bader, "Tensor Decompositions and Applications," *SIAM Rev.*, vol. 51, pp. 455–500, 2009.
[14] J. K. Chege, M. J. Grasis, A. Yeredor, and M. Haardt, "Bayesian Estimation of a Probability Mass Function Tensor with Automatic Rank Detection," in *Proc. 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, Herradura, Costa Rica, 2023.
[15] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
[16] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd International Conference on Learning Representations, ICLR*, San Diego, CA, USA, 2015.