

Beam Index Map Prediction in Unseen Environments from Geospatial Data

Fabian Jaensch, Giuseppe Caire *Fellow, IEEE*, Begüm Demir *Senior Member, IEEE*

Abstract—In 5G, beam training consists of the efficient association of users to beams for a given beamforming codebook used at the base station and the given propagation environment in the cell. We propose a convolutional neural network approach that leverages the position of the base station and geospatial data to predict beam distributions for all user locations simultaneously. Our method generalizes to unseen environments without site-specific training or specialized sensors. The results show that it significantly reduces the number of candidate beams considered, thereby improving the efficiency of beam training. *Index Terms*—Convolutional Neural Networks, Machine Learning, MIMO

I. INTRODUCTION

In 5G/6G (massive) MIMO, large antenna arrays allow the direction of signal energy in specific directions via beamforming. This is particularly important for mmWave and THz communication due to high attenuation, and it also increases data rates and coverage at lower frequencies. Base stations can cover cells or sectors with a family of beam patterns (referred to as beamforming codebook) that have much higher gain and much narrower angle width than conventional sector-wide beams. The efficient association of users to beams is referred to as beam training and is an important aspect of the 5G physical layer. Conventional beam sweeping requires testing a large number of beams in order to find the best combination for a strong communication link and causes significant overhead [1].

Several works have proposed data-driven methods based on machine learning (ML) leveraging additional sensors to reduce the time and signaling overhead required for beam sweeping. In a fixed environment, the GPS geolocation of the user equipment (UE) can be used to determine the optimal beam [2]. To account for dynamic changes such as vehicular movement, prior works have used radar, LiDAR or cameras at the base station (BS) or UE for blockage detection and beam selection [3], [4], [5]. Temporal correlation has also been exploited to predict future beam indices from the current optimal choice [6] (beam tracking).

The aforementioned methods are site-specific, i.e. they rely on the collection of data for a specific environment and the trained models cannot be used in new environments. Consequently, in every new environment, data has to be collected and the training procedure needs to be repeated. In contrast,

[7] investigates applying a model trained in one scenario to another one leveraging transfer learning and incorporating information from cameras, LiDAR and GPS. Furthermore, [8] proposes to encode the geometry of the environment in the form of a point cloud derived from stereo imagery, which is then used as an input. While the resulting model is site-agnostic, it requires multiple photos of the environment from various locations and angles. One reason for the shortage of site-agnostic methods is presumably the lack of appropriate datasets. The largest openly available MIMO channel dataset, DeepMIMO(v3) [9], contains only 20 city environments with 3 base stations each—far below the hundreds of environments required for training environment-agnostic models in related tasks such as radio map prediction [10].

Inspired by our recent works on radio map prediction [10], [11], we propose a different point of view on the problem. In a static environment, the optimal beam index at each UE location is deterministically determined by the environment's geometry and the BS parameters. We consider the problem of predicting a map of optimal beam indices using ML, based on publicly available city height maps or aerial images, together with properly encoded information about the BS. Our contributions are as follows:

- We formulate the beam index prediction problem as an image-to-image translation task and solve it using convolutional neural networks (CNN). The trained models are not site-specific, but rather learn a mapping from the input geometry to the beam indices, so that they can be used in unseen environments.
- Simulation environments and CNN input data are derived from public geospatial databases. Besides the UE location, no further sensor data is required. We release a large synthetic channel dataset with a total of 20151 simulations in 609 city environments, each simulation featuring one BS and tens of thousands of Rx.¹
- The model predicts a beam candidate map for the entire area of interest in a single forward pass. It can be used by several UE and to update beam indices as UEs move.
- We compare the performance of models trained with different loss functions and input data.

II. SYSTEM MODEL AND PROBLEM DEFINITION

A. Channel Model

Consider a uniform rectangular array (URA) at a fixed Tx with $N_a, N_e \in \mathbb{N}$ antennas in azimuth and elevation direction,

F. Jaensch and G. Caire are with the Communications and Information Theory Group, Technische Universität Berlin, 10623 Berlin, Germany.

B. Demir is with the Remote Sensing Image Analysis Group, Technische Universität Berlin, 10623 Berlin, Germany, and with the BIFOLD - Berlin Institute for the Foundations of Learning and Data, 10623 Berlin, Germany.

Corresponding author: Fabian Jaensch (email: f.jaensch@tu-berlin.de).

¹Our code is published on https://github.com/fabja19/beam_index_prediction and the dataset on <https://zenodo.org/uploads/17290206>.

respectively, placed with $\frac{\lambda}{2}$ -spacing, where $\lambda > 0$ is the carrier wavelength. Assuming the geometric multipath channel model [8], the downlink channel matrix $\mathbb{H}_{\text{Tx}} \in \mathbb{C}^{N_a \times N_e}$ is

$$\mathbb{H}_{\text{Tx}} = \sum_{p=1}^{N_P} c_p e^{j\psi_p} a_{N_a}(\varphi_p^D) \otimes a_{N_e}(\vartheta_p^D), \quad (1)$$

where N_P is the number of paths, $c_p > 0$ and $\psi_p \in [0, 2\pi)$ are the magnitude and the phase of the p -th path coefficient, j is the imaginary unit,

$$a_N(\omega) = [1, e^{j\omega}, \dots, e^{j\omega(N-1)}] \quad \text{for } N \in \mathbb{N}, \omega \in \mathbb{R}, \quad (2)$$

is the steering vector, $\otimes$ denotes the tensor product and $\varphi_p^D, \vartheta_p^D \in [-\pi, \pi]$ are the beamspace angles of departure (AoD). These are calculated from the AoD ϕ_p^D, θ_p^D in a global Cartesian coordinate system as

$$\begin{aligned} \varphi_p^D &= \pi \sin(\theta_p^D) \cos(\phi_p^D), \\ \vartheta_p^D &= \pi \sin(\theta_p^D) \sin(\phi_p^D), \end{aligned} \quad (3)$$

see Fig. 2.

On the Rx side, we model the UE as having a small number of coarse angular “sectors” (e.g., corresponding to four directional antennas around a smartphone chassis). Each arriving path is associated with one such sector according to its azimuth angle of arrival (AoA). Formally, if we divide the azimuth domain $[0, 2\pi)$ into N_r disjoint sectors, each path with azimuth AoA ϕ_p^A contributes through a sector-selection vector

$$b_N(\phi_p^A) \in \{0, 1\}^{N_r}, \quad (4)$$

where $b_N(\phi_p^A)$ is one-hot at the index of the sector containing ϕ_p^A . The complete 3-rd order channel tensor in $\mathbb{C}^{N_a \times N_e \times N_r}$ is then

$$\mathbb{H} = \sum_{p=1}^{N_P} c_p e^{j\psi_p} a_{N_a}(\varphi_p^D) \otimes a_{N_e}(\vartheta_p^D) \otimes b_{N_r}(\phi_p^A). \quad (5)$$

On the Tx side, we assume a beam codebook consisting of beams generated as tensor products of one-dimensional flat-top beams in both dimensions described in [12]. We assume here that the size of the Tx codebook is equal to $N_a N_e$, although this is not necessary in general. The number of beams in each dimension may also be smaller than the number of antennas, depending on the codebook design. For the Rx, since we are interested in finding the sector with the strongest arriving signal, we chose the canonical basis vectors. More concretely, $w^H b_{N_r}(\phi_p^A)$ is either 0 or 1, depending on whether w selects the sector that contains ϕ_p^A or not. For brevity, we refer to all elements of the codebook as ‘beams’, although the Rx vectors do not correspond to physically synthesizable beams. We denote as $\mathcal{B}$ the codebook of size $N_a N_e N_r$ resulting of combining the beams in all dimensions. The application of $\mathbb{H}$ to a beam $u \otimes v \otimes w \in \mathcal{B}$ gives the complex channel gain

$$\begin{aligned} \mathbb{H}(u \otimes v \otimes w) &= \sum_{p=1}^{N_P} c_p e^{j\psi_p} u^H a_{N_a}(\varphi_p^D) v^H a_{N_e}(\vartheta_p^D) w^H b_{N_r}(\phi_p^A). \end{aligned} \quad (6)$$

B. Objective

The quality of a link established using a beam $u \otimes v \otimes w \in \mathcal{B}$ is measured in terms of the received signal strength (RSS). Since in our scenario the wavelength is significantly shorter than the level of detail of the environment data, the phases ψ_p cannot be accurately estimated. Therefore, we model them as uniformly distributed random variables over $[0, 2\pi]$ and the RSS in Watts is calculated as

$$\begin{aligned} \text{RSS}(u \otimes v \otimes w) &= \mathbb{E} [|\mathbb{H}(u \otimes v \otimes w)|^2] \\ &= \sum_{p=1}^{N_P} c_p^2 |u^H a_{N_a}(\varphi_p^D) v^H a_{N_e}(\vartheta_p^D) w^H b_{N_r}(\phi_p^A)|^2. \end{aligned} \quad (7)$$

This corresponds to a local time-frequency average of the instantaneous power (see [11] for further discussion). For a fixed Tx-Rx pair, the *effective channel tensor* is obtained by calculating (7) for each beam and collecting the values in a 3-rd order tensor that is an element of $\mathbb{R}^{N_a \times N_e \times N_r}$.

Consider a dataset of M channel samples, each corresponding to one Tx-Rx pair, with RSS functions RSS_ℓ , for $\ell = 1, \dots, M$. Note that in our dataset (Section III-A), each channel sample corresponds to a specific environment (city map), a fixed Tx location within that environment, and an Rx location within the Tx’s coverage area. A predictor estimates for each channel sample a probability distribution over the set of possible beams. Let

$$(u \otimes v \otimes w)_\ell^* = \text{argmax}_{u \otimes v \otimes w \in \mathcal{B}} \text{RSS}_\ell(u \otimes v \otimes w) \quad (8)$$

be the optimal beam and denote as $B_\ell^k \subseteq \mathcal{B}$ the set of the $k \in \mathbb{N}$ beams with the highest probability according to the predictor.

Ideally, the optimal beam index should coincide with the index that the predictor assigns the highest probability to. If an exact prediction is infeasible, it is still beneficial to identify a small subset of $k \ll N_a N_e N_r$ high-quality candidate beams to reduce the overhead of beam training. To capture this, we measure the performance of the predictor in terms of top- k accuracy

$$\text{Acc}_k = \frac{1}{M} \sum_{\ell=1}^M \chi_{B_\ell^k}((u \otimes v \otimes w)_\ell^*), \quad (9)$$

for $k \in \mathbb{N}, k < N_a N_e N_r$, where χ_B denotes the indicator function of a set B .

However, for the same spatial location, often more than one beam provides a satisfactory signal strength, since different paths (e.g., from distinct reflections) may yield similar signal strength at the same location. Therefore, as a second evaluation metric, we consider the *throughput ratio (TPR)* [5], [3] defined as

$$T_k = \frac{\sum_{\ell=1}^M \max_{u \otimes v \otimes w \in B_\ell^k} \log_2(1 + \text{SNR}_\ell(u \otimes v \otimes w))}{\sum_{\ell=1}^M \log_2(1 + \text{SNR}_\ell((u \otimes v \otimes w)_\ell^*))}. \quad (10)$$

Here, $\text{SNR}_\ell = \frac{\text{RSS}_\ell}{\sigma^2}$ is the signal-to-noise ratio with noise variance σ^2 , which in turn depends on the signal bandwidth and the noise power spectral density (PSD).

III. BEAM INDEX MAP PREDICTION WITH CNN

Our goal is to predict optimal beam indices—or candidate sets—for all locations in a given city environment using a CNN. We formulate the task as image-to-image translation, mapping environment and Tx information to a spatial distribution over beam indices.

A. Dataset and Processing




Fig. 1: ((a)) City map with overlaid nDSMs of buildings in red, vegetation in green, color intensity correlating with height, and Tx position in blue, ((b)) aerial image, ((c)) ground truth optimal beam indices, ((d))–((f)) radio maps depicting the RSS per UE location for three different beams

We use our 3D model of Berlin described in [11], based on open LiDAR data with corresponding aerial images. Environment information is encoded in *normalized digital surface models (nDSMs)* for the building and vegetation classes, depicting the height of objects per pixel. Each 256×256 nDSM corresponds to a 256 m^2 region at 1m resolution, see Fig. 1 for an example. For each Tx (placed on rooftop edges) and Rx (on a grid of 1m resolution), both with isotropic antennas, we simulate full path data (coefficients, delays, AoD, AoA) using the ray-tracing software Wireless InSite [13]. In total, 20151 simulations have been carried out in 609 city environments at a carrier frequency of 3.9GHz.

For each grid location (pixel), the effective channel tensor is obtained from the path data as described in Section II, with parameters $N_a = 8, N_e = N_r = 4$. The orientation of the arrays is chosen so that in azimuth direction they point in straight line away from the building where the Tx is located and the boresight is tilted downwards at a 45° angle from the horizontal plane passing through the Tx. Postprocessing revealed that the $1 \times 1 \text{ m}$ spatial resolution combined with the channel tensor dimensions incurred excessive memory usage, implying impractical memory requirements during training. To resolve this issue, we downscale beam index maps and channel

tensors by a factor of 4 in each spatial dimension.²




Fig. 2: AoD in the coordinate system of the array. The intercepts dx', dy' of the projection onto the $x' - y'$ -plane with the x' - and y' -axes correspond to the angles φ, θ up to the factor π in (3).

B. Deep Learning Approach

Tx position and orientation are encoded using one-hot maps, distance maps, and azimuth maps as in [11]. We also adopt the UNetDCN architecture from there with $C = 32, d = 5$, previously shown effective for learning radio maps.

This is a semantic segmentation task with the particularity that near-optimal classes may still yield high signal quality. To account for this, we train the model with several different objective loss functions in separate experiments and compare the performance. In all cases, we use standard supervised learning with backpropagation. The number of channels of the output tensor and the last activation function are adapted to the requirements of the loss function. We evaluate both joint classification over $N_a N_e N_r$ beam indices and separate classification over azimuth, elevation, and Rx sectors (*sep*), when applicable.

We consider the following loss functions:

- *Cross Entropy (CE)*: The standard loss function for classification and segmentation problems.
- *Cross Entropy with Probabilities (CEP)*: The effective channel tensor in dB is normalized and used as a class probability distribution for CE, similar to label smoothing [14].
- *Wasserstein (WS)*: Following the intuition that predicting an adjacent beam to the optimal one should result in lower loss than predicting a beam that is further away, we define beam distances via Euclidean distance between index tuples. We calculate the Wasserstein loss [15] between the probability distribution defined by the output logits and the one-hot target distribution of the optimal index using these distances.
- *Index Regression (IR)*: Prediction of the optimal index is considered as a regression problem, following again the intuition that subsequent beam indices correspond to

²To assess the impact of downscaling, we additionally generated high-resolution beam index maps for a small subset of samples. Interpreting the ranking of beams in each pixel of the downscaled effective channel tensor as an estimate for the corresponding 4×4 block in the high-resolution map, we observed a top-1 accuracy of 0.72 and a top-1 TPR of 0.88, when referring to the values in the high-resolution map as the ground truth.

Model	Top-1	Top-2	Top-4	Top-8	Top-16	Top-32
CE	0.611	0.736	0.834	0.909	0.963	0.992
CE sep	<i>0.583</i>	<i>0.701</i>	<i>0.784</i>	<i>0.852</i>	<i>0.909</i>	<i>0.956</i>
CEP	0.431	0.614	<u>0.748</u>	<u>0.844</u>	<u>0.921</u>	<u>0.971</u>
CEP sep	0.461	0.619	0.728	0.813	0.893	0.951
WS	0.347	0.412	0.459	0.506	0.547	0.591
WS sep	0.412	0.497	0.550	0.603	0.672	0.748
IR sep	0.429	0.431	0.443	0.457	-	-
GR	0.400	0.539	0.637	0.713	0.800	0.887
GR sep	<u>0.484</u>	0.616	0.699	0.775	0.855	0.919

TABLE I: Top- k accuracy. Best results for each k highlighted in **boldface**, *italics* and underlined.

Model	Top-1	Top-2	Top-4	Top-8	Top-16	Top-32
CE	0.974	0.987	0.994	0.998	0.999	1.000
CE sep	<i>0.963</i>	<i>0.980</i>	<i>0.989</i>	<i>0.993</i>	<i>0.996</i>	<i>1.000</i>
CEP	0.789	0.904	0.968	<u>0.990</u>	<u>0.997</u>	<u>0.999</u>
CEP sep	0.875	0.941	0.970	0.983	0.993	0.998
WS	0.686	0.730	0.749	0.758	0.767	0.789
WS sep	0.888	0.932	0.954	0.964	0.972	0.983
IR sep	0.862	0.870	0.879	0.882	-	-
GR	0.866	0.929	0.967	0.981	0.991	0.997
GR sep	<u>0.917</u>	<u>0.960</u>	<u>0.9774</u>	0.988	0.994	0.998

TABLE II: Top- k TPR. Best results for each k highlighted in **boldface**, *italics* and underlined.

illuminated areas which are spatially close. This intuition is meaningful only if we consider azimuth, elevation and Rx beams separately.

- *Gain Regression (GR)*: The model is trained to predict the effective channel tensor in dB scale with MSE loss. The predictions are ranked according to the order of the estimated RSS values.

IV. RESULTS

A. Experimental Setup

We use about 80% of the dataset for training, 10% for validation (learning rate reduction, early stopping and selection of the best model state) and 10% for testing. Locations with path gain below -147 dB across all beams are excluded, as they imply $\text{SNR} < 0$ and are irrelevant for communication. This is discussed in more detail in [10], [11]. This threshold assumes 10 MHz bandwidth, -174 dBm/Hz noise PSD, 23dBm transmit power, and 0dB noise figure, following [11].

B. Numerical Results on Loss Functions

The performance of models trained with the loss functions described in III-B is compared in Tables I, II and the corresponding curves are shown in Fig. 4. The results show that standard CE loss over all beam combinations yields the best performance in terms of both evaluation metrics considered, and for all k . From the curves for WS and IR sep in Fig. 4, we observe that a higher accuracy does not necessarily imply a higher TPR as well. A high TPR of over 0.99 is already achieved for $k = 4$, although a high accuracy of at least 0.9 or 0.99 requires $k = 8$ or $k = 32$ beams, respectively. This indicates that suboptimal beams often yield comparable RSS, due to overlapping beam footprints, reflections, and sidelobes — consistent with findings in [5], [6].




Fig. 3: Training and validation losses for CE. Dots indicate learning rate reduction; the cross marks the best epoch in terms of validation loss.

In Fig. 3, we show the loss curves during training and Fig. 5 provides some samples from the dataset together with the areas in which the top- k predicted beams contain the optimal one. Here, we used the model trained with CE. The last column also shows the areas with line-of-sight (LoS) to the Tx. LoS areas (dominant path is direct) are shown in dark blue areas with attenuated LoS due to trees in light blue, areas without a direct path in red. Performance is strongest in LoS regions, as expected, but high top-4/top-8 coverage is also achieved in NLoS regions.




Fig. 4: Top- k accuracy (solid) and TPR (dashed) plotted vs k

C. Numerical Results on Input Data

Following [11], we retrain models on either aerial imagery only or imagery with an unclassified nDSM (combined height of buildings and vegetation) using CE loss. These inputs are more widely available, broadening applicability. The results in Table III and IV show that also without explicit access to the classes and even without height information, the model retains high prediction accuracy.




Fig. 5: ((a)) City map (explained in Fig. 1), ((b)) aerial image, ((c)) ground truth optimal beam indices, ((d))–((g)) areas where the optimal beam index is contained in the set of top- k predicted beams in green, others in red, for $k = 1, 2, 4, 8$, ((h)) areas with a straight path from the Tx in blue, darker blue if this is the strongest among all arriving paths.

Inputs	Top-1	Top-2	Top-4	Top-8	Top-16	Top-32
Baseline	0.611	0.736	0.834	0.909	0.963	0.992
Image	0.586	0.706	0.804	0.888	0.952	0.989
Image+nDSM	0.606	0.733	0.832	0.909	0.962	0.991

TABLE III: Top- k accuracy for modified inputs. Baseline is CE in Table I.

Inputs	Top-1	Top-2	Top-4	Top-8	Top-16	Top-32
Baseline	0.974	0.987	0.994	0.998	0.999	1.000
Image	0.970	0.984	0.992	0.997	0.999	1.000
Image + nDSM	0.972	0.986	0.994	0.997	0.999	1.000

TABLE IV: Top- k TPR for modified inputs. Baseline is CE in Table II.

V. CONCLUSION

We demonstrate that CNN-based models can significantly reduce beam training overhead using only public geospatial data and Tx/UE positions. As the model predicts beam indices for all locations, inference is required only once per area, not per UE position. This also enables multi-user applicability from a single inference pass.

One limitation of our experimental setup is that all environments are static. For future work, it would be interesting to incorporate changes in the environment such as cars moving along the streets. Since inference with the trained models is very fast, the beam index map could be updated periodically.

While our method generalizes well across places in the same city, we expect it to be less accurate in cities with a very different distribution of buildings and vegetation. Retraining on a broader dataset spanning multiple cities would likely mitigate this limitation.

REFERENCES

- [1] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi, “A tutorial on beam management for 3gpp nr at mmwave frequencies,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 173–196, 2019.
- [2] M. Arvinte, M. Tavares, and D. Samardzija, “Beam management in 5g nr using geolocation side information,” in *2019 53rd Annual Conference on Information Sciences and Systems (CISS)*, 2019, pp. 1–6.
- [3] M. Dias, A. Klautau, N. González-Prelcic, and R. W. Heath, “Position and lidar-aided mmwave beam selection using deep learning,” in *2019 IEEE 20th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2019, pp. 1–5.
- [4] A. Klautau, N. González-Prelcic, and R. W. Heath, “Lidar data for deep learning-based mmwave beam-selection,” *IEEE Wireless Communications Letters*, vol. 8, no. 3, pp. 909–912, 2019.
- [5] 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. 87 479–87 493, 2019.
- [6] S. Jiang, G. Charan, and A. Alkhateeb, “Lidar aided future beam prediction in real-world millimeter wave v2i communications,” *IEEE Wireless Communications Letters*, vol. 12, no. 2, pp. 212–216, 2023.
- [7] J. Gu, B. Salehi, S. Pimple, D. Roy, and K. R. Chowdhury, “Tune: Transfer learning in unseen environments for v2x mmwave beam selection,” in *ICC 2023 - IEEE International Conference on Communications*, 2023, pp. 1658–1663.
- [8] W. Xu, F. Gao, S. Jin, and A. Alkhateeb, “3d scene-based beam selection for mmwave communications,” *IEEE Wireless Communications Letters*, vol. 9, no. 11, pp. 1850–1854, 2020.
- [9] A. Alkhateeb, “Deepmimo: A generic deep learning dataset for millimeter wave and massive mimo applications,” 2019. [Online]. Available: <https://arxiv.org/abs/1902.06435>
- [10] R. Levie, c. Yapar, G. Kutyniok, and G. Caire, “Radiounet: Fast radio map estimation with convolutional neural networks,” *IEEE Trans. Wireless Commun.*, vol. 20, no. 6, pp. 4001–4015, 6 2021.
- [11] F. Jaensch, G. Caire, and B. Demir, “Radio map prediction from aerial images and application to coverage optimization,” *IEEE Transactions on Wireless Communications*, pp. 1–1, 2025.
- [12] F. Pedraza, J. C. Hauften, F. Jaensch, S. Li, and G. Caire, “Distributed beam alignment in sub-thz d2d networks,” p. 8, 2024. [Online]. Available: <https://arxiv.org/abs/2412.16015>
- [13] Remcom. Wireless insite. [Online]. Available: <https://www.remcom.com/wireless-insite-em-propagation-software/>
- [14] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2818–2826.
- [15] C. Frogner, C. Zhang, H. Mobahi, M. Araya-Polo, and T. Poggio, “Learning with a wasserstein loss,” in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, ser. NIPS’15. Cambridge, MA, USA: MIT Press, 2015, p. 2053–2061.