

This is the peer reviewed version of the following article:

LIDAR and Position-Aided mmWave Beam Selection with Non-Local CNNs and Curriculum Training / Zecchin, M.; Mashhadi, M. B.; Jankowski, M.; Gunduz, D.; Kountouris, M.; Gesbert, D.. - In: IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY. - ISSN 0018-9545. - 71:3(2022), pp. 2979-2990. [10.1109/TVT.2022.3142513]

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

02/05/2026 04:44

(Article begins on next page)

LIDAR and Position-Aided mmWave Beam Selection with Non-local CNNs and Curriculum Training

Matteo Zecchin*, Mahdi Boloursaz Mashhadi*, Mikolaj Jankowski*, Deniz Gündüz, Marios Kountouris, David Gesbert

Abstract—Efficient millimeter wave (mmWave) beam selection in vehicle-to-infrastructure (V2I) communication is a crucial yet challenging task due to the narrow mmWave beamwidth and high user mobility. To reduce the search overhead of iterative beam discovery procedures, contextual information from light detection and ranging (LIDAR) sensors mounted on vehicles has been leveraged by data-driven methods to produce useful side information. In this paper, we propose a lightweight neural network (NN) architecture along with the corresponding LIDAR preprocessing, which significantly outperforms previous works. Our solution comprises multiple novelties that improve both the convergence speed and the final accuracy of the model. In particular, we define a novel loss function inspired by the knowledge distillation idea, introduce a curriculum training approach exploiting line-of-sight (LOS)/non-line-of-sight (NLOS) information, and we propose a non-local attention module to improve the performance for the more challenging NLOS cases. Simulation results on benchmark datasets show that, utilizing solely LIDAR data and the receiver position, our NN-based beam selection scheme can achieve 79.9% throughput of an exhaustive beam sweeping approach without any beam search overhead and 95% by searching among as few as 6 beams. In a typical mmWave V2I scenario, our proposed method considerably reduces the beam search time required to achieve a desired throughput, in comparison with the inverse fingerprinting and hierarchical beam selection schemes.

Index terms— mmWave beam selection, LIDAR point cloud, non-local convolutional classifier, curriculum training, knowledge distillation.

I. INTRODUCTION

Millimeter wave (mmWave) communication constitutes a fundamental technology in 5G and future networks, which allows to overcome communication bottlenecks of the over-exploited sub-6GHz bands. To overcome the severe propagation impairments of the above-10GHz spectrum, such as high path attenuation and penetration losses, mmWave communication systems employ massive number of antennas at the base station (BS) to form highly directional beams and attain a large beamforming gain. Because of the narrow mmWave beamwidth, extremely precise alignment and tracking procedures are necessary in order to establish a reliable and high throughput communication link. The optimal communication beam can be easily determined with full channel knowledge; however, in the large antenna regime, obtaining an estimate of the high dimensional channel matrix is costly; and hence, beam selection for efficient communication requires iterative search procedures. In vehicular-to-infrastructure (V2I) communications, for which mmWave communication is envisioned to be a key technology [1], beam selection

and tracking are particularly challenging due to the high mobility of the receivers, which leads to reduced beam coherence time [2]. In this scenario, conventional beam selection techniques, such as beam sweeping or multi-level beam selection [3], [4] impose a significant overhead. Therefore, more efficient beam selection techniques that can reduce the cost of iterative search procedure by exploiting contextual information are of great interest.

It has been shown that contextual information from sensors mounted on the vehicles and the infrastructure can be leveraged to reduce the beam selection overhead. For instance, the position information provided by vehicle global positioning system (GPS) can be used to apply an inverse fingerprint approach and query the most prominent mmWave beams [5]. Inertial sensors placed on vehicle's antenna arrays enable efficient antenna element configuration by tracking the orientation of the vehicle [6]. Furthermore, positional and motion information can be jointly processed to further reduce the alignment overhead [7]. From the infrastructure side, a radar located at the BS can help estimate the direction of arrival, which would aid the beam search [8]. Spatial information obtained from out of band measurements was exploited in [9]–[12] where [11], [12] used sub-6GHz channel measurements to train neural network (NN)s for mmWave beamforming. Vision-aided approaches were proposed in [13]–[15]. Base stations equipped with cameras were proposed to employ computer vision and deep learning techniques to predict mmWave blockage and beam strength in [14]. The authors in [15] built a panoramic point cloud from images taken within the cellular coverage area. This point cloud gives a view of the scattering environment, which is then input to a neural network (NN) to predict the optimal beams.

Thanks to recent surge of autonomous driving technologies, high dimensional sensor information is nowadays commonly available also at the vehicle side. For instance, light detection and ranging (LIDAR) is commonly used for autonomous navigation. LIDAR uses a laser to produce a depth map of the environment and surrounding obstacles using delay measurements of the back-scattered signal. Because of the data dimensionality and the lack of analytical models that would relate LIDAR depth map to mmWave beams quality, data-driven methods have been considered to effectively process LIDAR signals as side information for beam search. In [16], [17], a NN architecture was trained over simultaneous LIDAR and ray-tracing channel datasets with a top- k classification metric to identify k beam directions that most probably include the beam resulting in the largest channel gain. In order to reduce the computational cost and NN model size, a simplified classifier architecture that can be trained in a distributed fashion using federated learning was proposed in [18].

This paper builds on the unpublished work of the authors that recently won the “AI/ML in 5G” competition ranking *first* in the “ML for mmWave beam selection” challenge organized by the International Telecommunications Union (ITU) [19], [20]. We propose a convolutional neural network (CNN) architecture along with the corresponding LIDAR preprocessing technique for data-driven mmWave beam selection. The proposed model is trained to exploit LIDAR and positional data in order to identify the best beam directions and reduce the beam search overhead in V2I communication. The specific

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

* Equal contribution.

M. Zecchin, M. Kountouris and D. Gesbert are with the Communication Systems Department, EURECOM, Sophia-Antipolis, France.

M. Boloursaz Mashhadi is with the 5GIC & 6GIC, Institute for Communication Systems (ICS), University of Surrey, UK.

M. Jankowski and D. Gündüz are with the Dept. of Electrical and Electronic Eng., Imperial College London, UK.

The work of M. Zecchin is funded by the Marie Skłodowska Curie action WINDMILL (grant No. 813999).

Mikolaj Jankowski and D. Gündüz received funding from the European Research Council through project BEACON (grant No. 677854).

D. Gesbert's and M. Kountouris's contribution are partially supported from a Huawei France-funded Chair towards Future Wireless Networks.

contributions of this paper in comparison with previous works [16]–[18] can be summarized as follows:

- Inspired by the knowledge distillation (KD) techniques [21], we propose a novel loss function, which not only maximizes the prediction accuracy of the best beam index, but also its corresponding power gain. The proposed loss function improves the beam prediction accuracy specifically for smaller k values achieving considerably higher throughput with significantly reduced beam search overhead.
- We utilize a non-local attention scheme, which improves the beam classification accuracy, specifically for the non-of-sight (NLOS) case. Convolutional classifiers used in previous works [16]–[18] learn local features from the LIDAR input and exploit them for beam classification. We observe that the NN utilizing our proposed non-local attention module considerably benefits from a non-local perception of the LIDAR input, specifically in NLOS scenarios where the mmWaves may be reflected from scatterers located far away.
- We propose a curriculum training strategy, which improves both the convergence speed and the final beam prediction accuracy. We observe that for the samples with no dominant LOS component, the strongest propagation path becomes significantly less predictable as it depends on the location of scatterers and reflectors, which is mainly determined by the traffic conditions. With NLOS samples being more challenging, the proposed curriculum learning strategy starts training with the LOS samples first, and gradually exposes the classifier to more complex NLOS samples. This training strategy achieves faster convergence and improved beam classification accuracy.
- We also present a protocol to integrate the proposed NN-based beam alignment scheme into the existing fifth generation (5G) new radio (NR) standard. We provide comparisons with the hierarchical beam selection [22] and inverse fingerprinting [23] protocols in terms of the time required for beam search. The simulation results show that, in spite of a short initial processing delay, the proposed solution outperforms alternative beam alignment schemes in terms of the achieved throughput ratio. In a typical mmWave V2I scenario, our proposed method considerably reduces the beam search time required to achieve a desired throughput, in comparison with the inverse fingerprinting and hierarchical beam selection schemes.

Thanks to the above ideas, our NN classifier significantly outperforms previous works [16]–[18]. Utilizing the benchmark Raymobtime dataset [24], [25], our solution achieves top-1, top-5 and top-10 beam selection accuracies of 59.5%, 87.0%, and 92.2%, respectively. In a mmWave communication system with 256 possible beam pairs, our LIDAR-based approach achieves 95% of the available throughput, only by searching among the 6 most probable beams suggested by the NN classifier, greatly reducing the beam search space and the corresponding beam selection overhead. Our classifier harnesses, on average, 79.9% of the available throughput without any beam search at all, just by utilizing the LIDAR and position information. Finally, we show that the proposed NN classifier can be further simplified utilizing effective NN pruning techniques without significant loss of its performance while reducing the computational and storage costs for practical deployment.

The content of the paper is organized as follows: Sec. II introduces the system model. Sec. III illustrates our proposed NN model for mmWave beam selection utilizing LIDAR data. Simulation results are reported in Sec. IV. Sec. V presents a protocol for NN-based beam alignment to be integrated into the existing 5G NR standard. Finally, Sec. VI concludes the paper. The simulation code is

publicly available at: <https://github.com/MatteoEURECOM/LIDAR-mmWave-Beam-Selection>.

II. SYSTEM MODEL

We consider a downlink orthogonal frequency-division multiplexing (OFDM) mmWave system using analog beamforming, where the BS located on the street curb serves a vehicle in its coverage area utilizing N_c subcarriers. Both the transmitter and the receiver ends are equipped with antenna arrays with a single radio frequency (RF) chain and fixed complex beam codebooks, which we denote by $\mathcal{C}_t = \{\mathbf{f}_i\}_{i=1}^{C_t}$ and $\mathcal{C}_r = \{\mathbf{w}_j\}_{j=1}^{C_r}$, respectively. The downlink channel matrix from the BS to the vehicle over the n 'th subcarrier is denoted by \mathbf{H}_n .

For each precoder and combiner vector pair $(i, j) \in \mathcal{C}_t \times \mathcal{C}_r$, the resulting channel gain at subcarrier n is determined by $\mathbf{w}_j^H \mathbf{H}_n \mathbf{f}_i$, where $(\cdot)^H$ denotes the conjugate transpose. Without loss of generality, we assume a unit transmit power over each subcarrier and a unit noise variance. The receive power summed over all subcarriers for the transmitter-receiver codebook pair (i, j) is

$$\mathbf{G}_{i,j} = \sum_{n=1}^{N_c} |\mathbf{w}_j^H \mathbf{H}_n \mathbf{f}_i|^2, \quad (1)$$

and the optimal pair of precoding and combining vectors that maximizes the receive signal-to-noise ratio (SNR) is given by

$$(i^*, j^*) = \underset{(i,j)}{\operatorname{argmax}} \mathbf{G}_{i,j}. \quad (2)$$

Without side information, the transmitter and receiver need to perform an exhaustive search through all $\mathcal{C}_t \times \mathcal{C}_r$ beam pairs in order to identify (i^*, j^*) . Our goal is to infer a small subset of k beam pairs $\mathcal{S}_k \subset \mathcal{C}_t \times \mathcal{C}_r$ exploiting the available position and LIDAR data, such that $(i^*, j^*) \in \mathcal{S}_k$. This results in a reduction of $\frac{k}{C_t \times C_r}$ in the search space of the beam selection procedure. Two metrics to gauge the quality of \mathcal{S}_k as a function of its size k , are the top- k accuracy and top- k throughput ratio. The top- k accuracy is formally defined as

$$A(k) = \mathbb{E} [\mathbb{1}\{(i^*, j^*) \in \mathcal{S}_k\}], \quad (3)$$

where $\mathbb{1}\{(i^*, j^*) \in \mathcal{S}_k\}$ is the indicator function associated to the event that the best beam index is in the top- k classifier output. With the unit noise variance and transmit power assumptions, the top- k throughput ratio defined as

$$T(k) = \frac{\mathbb{E} [\max_{(i,j) \in \mathcal{S}_k} \log_2(1 + \mathbf{G}_{i,j})]}{\mathbb{E} [\log_2(1 + \mathbf{G}_{i^*,j^*})]} \quad (4)$$

is used as a complementary performance metric, where all expectations are with respect to the inherent randomness introduced by the vehicles' positions, channel realizations, and LIDAR measurements. Note that the top- k throughput ratio is a very informative metric for the problem at hand. In fact, the numerator represents the throughput that can be achieved (at a zero dB transmit SNR) by searching only among the top- k beams suggested by the NN; while the denominator is a normalizing factor representing the maximum throughput achievable by an exhaustive beam sweeping approach.

In the next section, we propose a novel NN architecture that jointly processes location information along with LIDAR data in order to solve the top- k classification task; and therefore, to find the most promising mmWave beams to establish a reliable communication link with reduced beam search overhead.




Fig. 1: Preprocessing of the LIDAR point cloud.




Fig. 2: The proposed CNN model architecture.

III. CNN-BASED BEAM SELECTION UTILIZING LIDAR DATA

We propose a novel CNN-based beam selection scheme, where connected vehicles utilize measurements from their LIDAR sensors along with their location data to reduce the beam search overhead required to establish a mmWave link with a nearby BS.

A. LIDAR Preprocessing

Raw LIDAR data is in the form of large point cloud measurements $\mathcal{P} = \{(x_i, y_i, z_i)\}_{i=1}^{|\mathcal{P}|}$, where each triplet (x_i, y_i, z_i) represents coordinates of an obstacle point measured by the LIDAR sensor. To avoid excessive computations on large point clouds, we preprocess the raw LIDAR data to get a simplified representation of the coverage area of the BS, which is then input to the classifier CNN. We assume that each vehicle knows the location of the BS and its coverage area, and divides the coverage area into a 2D grid of equal squares. We then produce a top-view representation of the coverage area setting grid entries to 1 whenever at least one point in \mathcal{P} lies within that grid square, and to 0 otherwise. We also embed the location of the vehicle and the BS into this representation by setting the grid value of the square accommodating the BS and the vehicle to -1 and -2 , respectively. Fig. 1 shows one such preprocessing step. We note that discarding the z -axis causes certain information loss; however, we found this loss not to affect the accuracy of the beam classification task considerably. On the other hand, it allows us to reduce the complexity of our NN model significantly.

B. NN Architecture

Fig. 2 shows our proposed NN architecture for LIDAR-aided mmWave beam selection, which is composed of 6 convolutional layers and 5 linear layers. Each layer is followed by the rectified linear unit (ReLU) activation, except for the last layer, which is followed by softmax activation to output beam predictions. Our architecture

consists of separate branches to process LIDAR and GPS inputs. The LIDAR branch comprises 6 convolutional layers followed by a linear layer to extract features from the preprocessed LIDAR input as in Sec. III-A. The output features from the LIDAR branch are then concatenated with (x, y) location coordinates of the vehicle from the GPS input. We discard the BS coordinates, as these are fixed on the whole dataset and do not need to participate in the training process.

The concatenated feature vector is then input to 4 linear layers. The first three linear layers include 64 neurons and the last one outputs a 256-element vector corresponding to the number of possible beam pairs, i.e., $|C_t| \cdot |C_r| = 256$. Although these linear layers increase the complexity of our NN architecture, in Sec. III-F we use pruning techniques to effectively reduce the memory and computation requirements of the proposed NN.

C. Loss Function

To define the loss function, we denote by $\mathbf{y} \in \mathbb{R}_+^{(|C_t| \cdot |C_r|)}$ the vectorized version of the mmWave power gain matrix \mathbf{G} obtained by the following bijective map

$$\mathbf{y}_{(i-1) \cdot |C_r| + j} = \mathbf{G}_{i,j}. \quad (5)$$

For each \mathbf{y} vector, we also denote by $\bar{\mathbf{y}}$ the vector \mathbf{y} normalized to unit Euclidean norm and by \mathbf{y}^* the unitary vector that is non-zero on the component corresponding to the largest entry of \mathbf{y} (ties broken arbitrarily). Then, for $\beta \in [0, 1]$, we train our model by minimizing the following loss function

$$\ell(\mathbf{y}, \hat{\mathbf{y}}) = (1 - \beta)\mathcal{H}(\mathbf{y}^*, \hat{\mathbf{y}}) + \beta\mathcal{H}(\bar{\mathbf{y}}, \hat{\mathbf{y}}) \quad (6)$$

where $\hat{\mathbf{y}}$ is the model prediction and $\mathcal{H}(\cdot)$ denotes the empirical cross entropy that, for two non-negative unit norm vectors p and q in \mathbb{R}^d

is defined as

$$\mathcal{H}(p, q) = - \sum_{i=1}^d p_i \log(q_i). \quad (7)$$

The first term in (6) is a standard multi-class cross entropy loss that enforces the NN to predict the indices of the beam associated to the strongest mmWave beam. The second term, instead, drives the NN to match the output of each neuron in the last layer to the normalized power gain of the corresponding mmWave beam pair. This is achieved by treating \mathbf{y} as a vector of “soft labels” and minimizing the corresponding empirical cross entropy loss. This last term is specifically effective in top- k classification for $k > 1$, where it is advisable not to output exclusively the best beam pair, but also accurately predict k competing candidate beam pairs. Finally, the value $\beta \in [0, 1]$ provides a tradeoff between the two terms. The idea of combining two training objectives as in Eq. (6) resembles the KD technique [21]. KD is a popular model compression technique that aims at instilling the knowledge of a large classifier network, termed as *teacher*, into a lightweight *student* classifier. This is achieved by augmenting the original labels of a dataset by the soft prediction of the teacher model. This additional training objective has been shown to improve the performance of the student model, in some cases even outperforming the teacher model, and to act as a regularization term [21], [26]. In the context of mmWave beam selection, we show experimentally that by exploiting the soft labels, we are able to obtain similar gains and improve the predictive capabilities of the trained model.

D. Curriculum Training with LOS/NLOS Samples

In the absence of LOS, the predictability of the strongest propagation paths greatly decreases as a consequence of the prominent dependency on the relative positions of scatterers and reflectors. However, the presence of neighbouring moving obstacles renders NLOS condition frequent in vehicular type of communication. The difficulty of the prediction task in NLOS scenarios is so high compared to the LOS case that, data-driven methods tend to be biased towards the easier LOS samples to the detriment of the NLOS performance. In order to address this performance imbalance, we propose to adapt the sample distribution during training time so that the challenging instances are less likely during the initial phases of learning. This procedure is motivated by curriculum learning [27], [28], which suggests to expose the training process to easier instances at the initial phases and to gradually increase the difficulty of the tasks. To apply this strategy it is necessary to first define *scoring* and a *pacing* functions: the former assigns a level of difficulty to each sample while the latter determines at which rate the transition should be made from easier samples to harder ones during the learning process.

For the task at hand, a natural measure of difficulty is the absence of LOS and an effective way to modulate the difficulty of the learning task consists in changing the probability of NLOS samples during the training epochs. Hence, denoting by P the feature distribution from which the original training dataset \mathcal{D} is generated, we exploit a biased sampling scheme to generate skewed dataset \mathcal{D}_λ , whose hardness is proportional to the rejection coefficient $\lambda \in [0, 1]$. In particular, the set of instances \mathcal{D}_λ is created from \mathcal{D} by independently removing each NLOS sample with probability $1 - \lambda$. As a result, \mathcal{D}_λ represents a sample drawn from the following distribution

$$P_\lambda \propto (1 - q)P_{LOS} + q\lambda P_{NLOS} \quad (8)$$

where P_{LOS} is the feature distribution conditioned on the presence of LOS, P_{NLOS} the distribution conditioned on its absence, and q is the probability of NLOS condition under the original distribution P .




Fig. 3: Block diagram of the non-local attention module.

The pacing function is represented by a sequence $\{\lambda_i\}$ that for each epoch i determines the probability that a NLOS sample is accepted for training. As shown below, a properly chosen sequence $\{\lambda_i\}$ can improve both the convergence speed and the accuracy of the final solution compared to the unstructured and randomized sampling of training instances.

E. Non-local Attention for Improved NLOS Performance

The convolutional layers in Fig. 2 learn and exploit local features from the LIDAR input and use them for beam classification. However, the classifier can benefit from non-local perception of the coverage area specifically in the NLOS cases. In order to extract non-local perception of the coverage area from the LIDAR input in an efficient manner, we use a non-local attention module introduced in [29], [30]. We later show through simulations that the non-local attention module further improves the performance. The general input-output relation for non-local attention is given by

$$\mathbf{O}_i = \mathbf{I}_i + \frac{1}{\eta(\mathbf{I})} \sum_{\forall j} \phi(\mathbf{I}_i, \mathbf{I}_j) \psi(\mathbf{I}_j), \quad (9)$$

in which i is the space-time index of an output position whose response is to be computed, j is the index that enumerates all possible positions, \mathbf{I} is the input, and \mathbf{O} is the output of the same size as \mathbf{I} . A pairwise function ϕ computes a scalar representing the relationship between i and all j . The unary function ψ computes a representation of the input signal at position j . The response is normalized by the factor $\eta(\mathbf{I})$. We tried various popular choices for $\eta(\cdot)$, $\phi(\cdot)$ and $\psi(\cdot)$ functions (refer to [30] for further details), and found that the best performing ones in our case are the embedded Gaussian function for $\phi(\mathbf{I}_i, \mathbf{I}_j)$, given by $\phi(\mathbf{I}_i, \mathbf{I}_j) = \exp([\mathbf{W}_{\phi_1} \mathbf{I}_i]^T [\mathbf{W}_{\phi_2} \mathbf{I}_j])$, $\psi(\mathbf{I}_j) = \mathbf{W}_\psi \mathbf{I}_j$, and $\eta(\mathbf{I}) = \sum_{\forall j} \phi(\mathbf{I}_i, \mathbf{I}_j)$, as will be discussed later in Sec. IV. Here, \mathbf{W} s are trainable weight matrices. Fig. 3 provides the block diagram for our non-local attention module, where 1×1 convolutions implement \mathbf{W} weight multiplications, and Softmax activation implements the exponential function. Refer to [30] for more details on attention modules.

F. Network Pruning

To reduce the computational and memory footprint of the proposed model, we introduce an additional network pruning [31], [32] step. Network pruning is a method for reducing the computational complexity and the size of NNs by finding redundant neurons and removing them, based on some saliency measure. In this work we

employ a straightforward approach of removing neurons or filters with the lowest L_1 -norm of the weights, which already leads to maintaining a satisfactory performance with low network complexity. We leave the exploration of more effective pruning methods for future work. Since the majority of the weights in our network are contained in the later linear layers the convolutional part of the network is already lightly parameterized. In this work we explore two particular strategies for network pruning: *unstructured* and *structured* pruning. Unstructured pruning removes the weights with the smallest magnitude, whereas structured pruning removes entire neurons with the lowest mean magnitude of weights. Unstructured pruning usually leads to better gains, as it is able to prune more parameters, while maintaining satisfactory performance.

IV. NUMERICAL EVALUATIONS

In this section, we provide a series of experiments highlighting the performance improvements and benefits that each of the proposed techniques and architecture designs can attain. Subsequently, we evaluate the proposed solution with all of the above enhancements, compare it against the state of the art and showcase its superiority. For performance comparisons, we use the top- k classification accuracy and the top- k throughput ratio as defined in Sec. II.

A. Dataset

All experiments are carried out using the benchmark Raymobtime dataset [24], [25], which contains synthetic data for LIDAR and position-aided mmWave beam selection in a V2I communication scenario, where vehicles connect to a road side unit (RSU). The dataset is generated using a range of simulators including tools for transforming satellite images into realistic 3D urban scenes, mobility simulator for generating traffic data, and tools for generating ray-tracing mmWave channels. The resulting dataset contains the average received powers $G_{i,j}$ calculated for each transmitter/receiver beam pair, together with the simultaneous 3D point cloud \mathcal{P} collected by the LIDAR sensor. In this work, we utilize s008 and s009 subsets of the Raymobtime dataset, which are based on the data simulated for one of the streets in Rosslyn, Virginia. Each of the subsets considers a carrier frequency of 60GHz and contains 2086 and 2000 episodes (time-instances), respectively. The episodes were collected in the 30-second intervals of time. We train all models using s008, whereas we utilize s009 for testing.

B. Choice of β

We optimize the NN parameters by minimizing the proposed loss function (6). Note, however, that the proposed loss function is parameterized by $\beta \in [0, 1]$ weighting the training signal coming from the hard and soft labels, respectively, represented by the one-hot-encoded best beam index \mathbf{y}^* and the normalized channel gain vector $\bar{\mathbf{y}}$. Choosing a proper value for β is necessary in order to strike a good balance between these two pulling forces. For small β values, the objective function resembles the standard multi-class cross-entropy loss and the trained NN tends to act as a myopic classifier that tries to correctly predict the best beam index regardless of the performance of the other beams. On the other extreme, very large values of β drive the output of the model to match the normalized channel gains associated with different beams and to potentially trade the best beam prediction accuracy for this purpose. We evaluate the effect of β by training the model proposed in Sec. III-B for $\beta \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$ and we report the final top- k accuracy and throughput ratio averaged over 10 runs for each β value in Fig. 4. We observe that a value of $\beta = 0.8$ yields a trained model that combines the best of the two above mentioned




Fig. 4: Average top- k accuracy and throughput ratio, with 95% confidence intervals, for $k \in \{1, 5, 10\}$ averaged over 10 training runs for different β values.

TABLE I: Performance of curriculum, anti-curriculum and standard training procedures.

Strategy	$A(1)$	$A(5)$	$T(1)$	$T(5)$
Curr.	58.1%	86.6%	77.9%	94.1%
Standard	57.1%	86.0%	75.9%	92.9%
Anti-curr.	53.8%	85.1%	71.9%	91.2%

behaviours. In fact, for this choice of β , the predictor attains the highest top-1 accuracy, outperforming even the models trained for smaller β values that prioritize this metric. At the same time, for larger k , its performance is indistinguishable from the network trained with $\beta = 1$. The same conclusion holds for the top- k throughput ratio metric. The model trained with the optimal β value is able to provide 76% top-1 throughput ratio, while the ratios for the two worst performing values are 69% ($\beta = 0$) and 73% ($\beta = 1$).

Fig. 4 also shows that incorporating in the loss function our proposed KD term not only improves the accuracy and throughput ratio, but with a proper choice of β it also helps to reduce the variance. For example, without the second term (i.e., for $\beta = 0$), the top-1 accuracy shows a relatively wider confidence interval. This means that, although we get a good average beam prediction accuracy, there exist settings in which our prediction accuracy is unsatisfactory. A tighter confidence interval ensures that such settings are infrequent. According to this figure, $\beta = 0.8$ is an excellent choice that provides not only the best accuracy and throughput, but also a small variance.

C. Curriculum Training

While it is usual practice to train NNs using batches of data sampled uniformly at random from the training dataset, in the following we illustrate the benefits of biasing the sampling procedure in order to obtain a sequence of training samples with an increasing level of difficulty. As illustrated in Sec. III-D, this is achieved by employing a sample rejection strategy based on the presence or absence of LOS components. Specifically, we train the proposed NN architecture for 45 epochs decreasing the NLOS rejection probability $1 - \lambda$ from 1 by steps of 0.2 every 9 epochs until reaches 0. In this




Fig. 5: Evolution of the top- k accuracy over the training epochs for curriculum, standard and anti-curriculum training approaches. The curves are averaged over 10 training runs.

TABLE II: Performance improvement by the proposed non-local attention approach.

$\phi(\cdot)$	$A(1)$	$A(5)$	$T(1)$	$T(5)$
Embedded	58.3%	86.7%	78.7%	94.3%
Gaussian	57.9%	86.5%	77.7%	94.3%
Dot	57.7%	86.7%	77.7%	94.1%
Without NLA	57.1%	86.0%	75.9%	92.9%

manner, the first batches contain only LOS samples, whereas during the last 9 epochs the ratio between LOS and NLOS will be the same as the one in the original unbiased empirical distribution. We also consider the opposite strategy, namely exposing the NN to batches of hard samples first. This is obtained by the same sampling procedure but with the role of NLOS and LOS swapped. We term this alternative anti-curriculum as it starts from the hardest instances. As a natural baseline, we also consider the standard unbiased sampling procedure. For each case, we train the same model architecture proposed in Sec. III-B using the loss function with $\beta = 0.8$. Distinct training dynamics result in different convergence time and final accuracy values. In Fig. 5, we plot the evolution of the accuracy metrics, averaged over 10 repetitions, of the standard, curriculum and anti-curriculum learning procedures. In terms of convergence time the curriculum learning strategy outperforms both the standard and anti-curriculum sampling schemes as it quickly plateaus to higher accuracy levels. The final performance, also averaged over 10 runs, is reported in Table I. The curriculum learning strategy improves by 2% the top-1 throughput ratio and by 1.2% the top-5 throughput ratio compared to standard learning. On the other hand, anti-curriculum learning has a detrimental effect on the performance, resulting in a performance loss of 4% and 1.7% in terms of top-1 and top-5 throughput ratio, respectively.

D. Non-local Attention

In the following we consider augmenting the NN architecture of Sec. III-B by adding one non-local attention module located after the fifth convolutional layer of the LIDAR processing branch. This design choice represents the best trade-off between the performance gain and additional computational burden. In fact, even if multiple non-local blocks generally lead to better performance, our NN is reasonably shallow and it does not display a significant improvement when more than one non-local block is instantiated. Additionally, we did not observe any meaningful variation in the final performance by placing the non-local block at different depths in the network. On the other hand, the number of floating point operations (FLOPs) required for

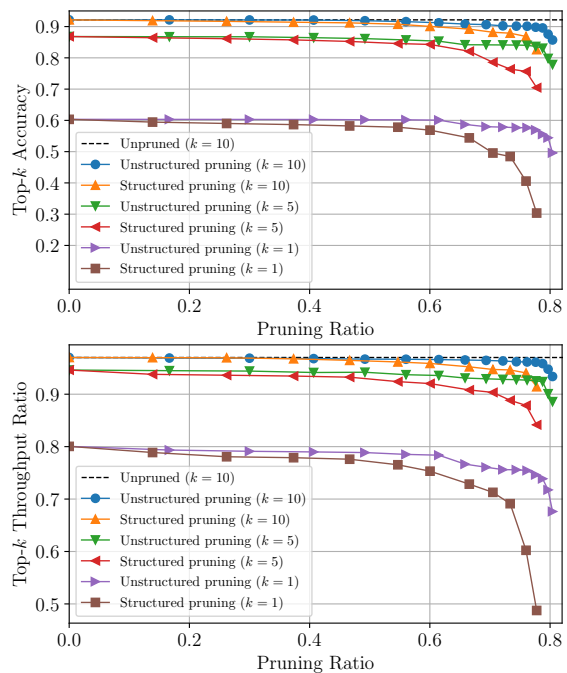


Fig. 6: Accuracy and throughput ratio as a function of pruning ratio.

a forward pass through the non-local module depends on its position within the NN. In fact, high level representations at deeper stages of the NN have lower dimension resulting in modest computational effort when processed by the non-local block. In particular, the output of the fifth convolutional layer is a tensors of size $5 \times 50 \times 5$, that is further sub-sampled using a max pooling kernel of size $1 \times 2 \times 2$. This operation preserves the non-local behaviour but it translates in computational cost of only 250k floating point operations.

We test three popular choices for the non-local operations denoted by $\phi(\cdot)$: Gaussian, embedded Gaussian and dot product pair-wise functions. For each, we train the model minimizing the proposed loss for $\beta = 0.8$ and the standard sampling procedure. In this way we are able to assess the gain that the non-local attention module brings independently of curriculum learning. We average the final performance over 10 runs and report the average results in Table II. Across the three different choices of $\phi(\cdot)$ we do not observe great

variability, with the non-local attention using embedded Gaussian slightly outperforming the others. Nonetheless, compared to the same model without non-local attention we measure an improvement of 2.8% and 1.4% in terms of top-1 and top-5 throughput ratios, respectively.

E. Pruning

We consider network pruning to further reduce the computational and storage cost of deploying the proposed NN solution. Before we prune any of the weights, we first pretrain the network, following the strategy described in Sec. III-B. After the pretraining is finished, we run multiple iterations of pruning by first discarding a portion of the remaining weights with the lowest L_1 -norm and fine-tuning the network by following the same training strategy that we used for pretraining. We evaluate both the unstructured and structured pruning methodologies introduced in Sec. III-F. The results are reported in Fig. 6, where we plot the top- k accuracy and top- k throughput ratio of the pruned model against the pruning ratio. The pruning ratio is defined here as the fraction of the weights removed from the unpruned model.

The unstructured pruning achieves better accuracy and throughput ratio at larger pruning rates, compared to structured pruning. This is due to the fact that unstructured pruning can be more precise in removing weights since it works by removing single weights rather than entire columns or rows of the weight matrix. Nevertheless, in unstructured pruning, weight matrices of the linear layers are only partially sparsified, which is computationally sub-optimal compared to removing entire rows or columns as structured pruning does. In both cases we find that the pruned model maintains excellent predictive performance even for large pruning ratios. In fact, it is possible to obtain a top-10 throughput ratio above 96.5%, despite pruning 60% of the weights in the case of structured pruning, and 95.9% with unstructured pruning. A graceful performance degradation also happens in terms of accuracy. Therefore, we conclude that pruning represents a viable option to further reduce the computational and storage footprint of learned models while maintaining excellent predictive capabilities.

F. Comparison with LIDAR and Position-Aided Neural Beam Selection Techniques

In the following we evaluate and compare the proposed final solution that includes all the above improvements against two state-of-the-art data-driven approaches. The first baseline model we consider is the 13-layer NN presented in [16], [17] and the second one is the lighter version recently proposed in [18] that comprises 8 layers. In Table III we compare the accuracy and throughput achieved by these models along with their number of trainable parameters and FLOPs. These results are averaged over 10 Monte Carlo rounds of training and reported with the corresponding 95% confidence intervals. For a fair comparison, we have reported results for our unpruned architecture. The number of trainable parameters and FLOP count directly relate to the computational and memory footprint of the models, and therefore, are of interest in order to assess their deployment feasibility. Compared to the larger architecture in [16], [17], our model requires only 7.6% of the trainable parameters and 2.5% of the floating point operations to perform the forward pass. At the same time our model is comparable to [18] in terms of computational and storage cost; the difference can be further reduced using the pruning techniques proposed in Sec. III-F. This renders the proposed solution also applicable in distributed training scenarios with memory and computationally constrained devices. Our model consistently outperforms both baselines in terms of top- k




Fig. 7: Accuracy and throughput ratio performance curves comparing the proposed solution and previous work.

accuracy and throughput ratio metrics, which are reported in Fig. 7 for $k \in [0, 30]$. Our model yields a striking 79.9% top-1 throughput ratio, harnessing a great portion of the available rate without any search procedure. The two competing baselines attain only 46.1% and 70.3% top-1 throughput ratios respectively. At the same time, in order to ensure a 95% expected throughput ratio our model needs to sweep on average 6 beams, greatly reducing the beam search overhead. As a comparison, the two baselines requires 28 and 11 beams, respectively. In terms of NLOS performance, our model outperforms both alternatives providing an average top-1 throughput ratio of 79.0%, improving upon the two baselines by 25% and 4%, respectively. Finally, our proposed model achieves a tighter 95% confidence interval in comparison with [16]–[18] specifically for smaller k . This is very favourable in practice as it ensures more reliable performance guarantees for different instances of training and deployment of our model.

V. NN-BASED BEAM ALIGNMENT PROTOCOL

In the previous section we showed that our solution outperforms state-of-the-art NN-based baselines that infer the most prominent beam directions leveraging LIDAR and position information. In this section, we present a protocol for NN-based beam alignment to be integrated into the existing 5G NR standard, through which we show the feasibility of the proposed solution in practical scenarios considering the delay and computation overhead associated with NN-based inference. We also provide comparisons with alternative beam alignment protocols.

A. Integration with 5G NR

The beam sweeping procedure in 5G NR is based on synchronization signal blocks (SSB), which are used to probe the channel

TABLE III: Performance comparison between the proposed NN model and previous works.

Model	$A(1)$	$T(1)$	$A(5)$	$T(5)$	$A(10)$	$T(10)$	FLOP count	# params.
[16], [17]	$31.5 \pm 2.6\%$	$46.1 \pm 2.6\%$	$71.9 \pm 2.2\%$	$76.1 \pm 1.9\%$	$83.9 \pm 0.9\%$	$86.1 \pm 0.8\%$	179.01×10^6	403677
[18]	$52.3 \pm 1.9\%$	$70.3 \pm 2.6\%$	$85.3 \pm 0.9\%$	$90.8 \pm 1.5\%$	$91.1 \pm 0.3\%$	$94.7 \pm 0.6\%$	1.72×10^6	7462
Proposed	$59.5 \pm 0.5\%$	$79.9 \pm 0.8\%$	$87.0 \pm 0.3\%$	$94.6 \pm 0.8\%$	$92.2 \pm 0.2\%$	$96.9 \pm 0.6\%$	4.55×10^6	30872

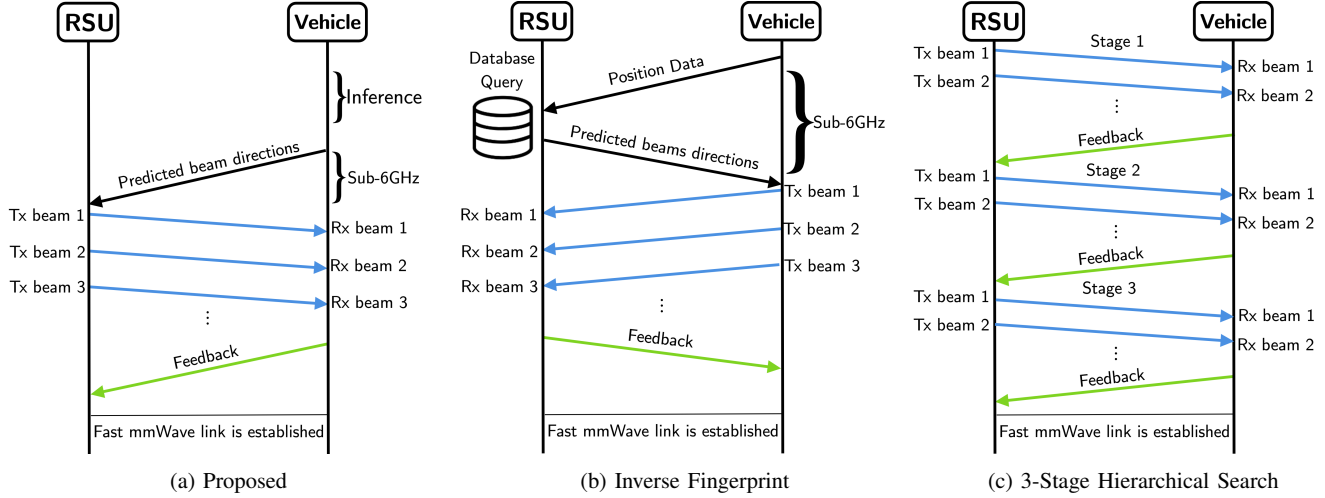


Fig. 8: Timing diagrams of (a) the proposed NN-based beam alignment protocol using LIDAR and position information, (b) the inverse fingerprinting scheme using position information only [23] and (c) a 3-stage hierarchical search [22].

when using each specific beam. In order to sweep multiple angular directions, a sequence of SSBs are time-multiplexed into a SSB burst. The length of SSB bursts is fixed to $T_{br} = 5\text{ms}$ and the number of SSBs that fit into a single burst depends on the operating frequency. In particular, for communication under 3GHz, the number of SSBs in a burst is 4, while it is 8 for communication at frequencies in the 3 to 6 GHz range, and can be up to 64 SSBs in a single burst for higher frequencies [33]. In our scenario, the operating frequency is 60GHz, leading to a single SSB duration of $T_{SSB} = T_{br}/64 = 0.078\text{ms}$.

Fig. 8a depicts our design that integrates the proposed NN-based beam alignment procedure with the existing beam search protocol of 5G NR. In the proposed design, the vehicle that wishes to communicate, predicts a subset of best mmWave beams using its pre-trained NN module fed with its readily available LIDAR and position data. Afterwards, it sends to the RSU the transmitter beam indices corresponding to the NN-predicted subset of candidate beams. These candidate beam indices are transmitted over a sub-6GHz link which is assumed to be available. Upon reception at the RSU side, the standard 5G NR SSB-based beam sweeping procedure is performed over the NN-predicted subset of candidate beams. Once the search is complete, the vehicle reports its measurements to the RSU and the mmWave link is established. The total time associated with the beam search process in this design is given by

$$T_{NN} = T_{Inf} + T_{CBC} + K \times T_{SSB} + T_{FB}, \quad (10)$$

where T_{NN} is the total time required for NN-based beam alignment, K is the number of candidate beams suggested by the NN, and T_{Inf} , T_{CBC} , and T_{FB} denote the time associated with NN inference, communicating the NN-predicted candidate beams over the sub-6GHz link, and the feedback, respectively. The overhead introduced by the NN-based approach amounts to the inference time and the time required to communicate the NN-predicted candidate beams over the sub-6GHz link. Measuring the beam inference time (including data pre-processing and forward path of the NN) using a consumer grade CPU, we have $T_{Inf} \leq 1\text{ms}$. We expect even smaller inference times

in case of dedicated hardware. The T_{CBC} and T_{FB} terms depend on the data rate of the sub-6GHz link, but are typically very small (i.e., $T_{CBC}, T_{FB} \leq 0.1\text{ms}$).

B. Comparison with Alternative Beam Alignment Protocols

Herein, we compare the performance of the proposed protocol against alternative beam alignment protocols, and show that the benefits of processing LIDAR data by the NN to predict the candidate beams justifies the initial inference overhead it introduces.

1) **Inverse fingerprinting:** The inverse fingerprinting beam alignment protocol proposed in [23], and depicted in Fig. 8b, exploits the position data to query a site-specific database, which is available at the RSU, and contains channel measurements along with the vehicle position information. According to the inverse fingerprinting protocol, the vehicle sends the position information to the RSU. Upon reception, the RSU determines the candidate beam directions using the previously measured channel fingerprints from users in the same location. The RSU then sends back the candidate SSB beam sweeping signals. Once all the candidate beam directions are probed, the RSU updates the local database and feedbacks the measurements to the vehicle in order to establish a fast mmWave link. As a result, the total time T_{IF} required for the fingerprinting is given by

$$T_{IF} = T_{GPS} + T_{Qr} + T_{CBC} + K \times T_{SSB} + T_{FB}, \quad (11)$$

where K is the number of candidate beams returned by querying the database, and T_{GPS} , T_{Qr} , T_{CBC} , and T_{FB} denote the time associated with communicating the GPS coordinates, querying the database at the RSU, communicating the resulting candidate beams, and the feedback, respectively. In Eq. (11), the T_{GPS} and T_{CBC} terms depend on the data rate of the sub-6GHz link and T_{Qr} depends on the processing power available for querying the database.

2) **Hierarchical beam selection:** An alternative class of algorithms carry out coordinated search without making use of any side information. The simplest scheme in this class is the coordinated

exhaustive search that sequentially tests all possible beams in a predefined order. A more efficient alternative is the coordinated hierarchical beam search which leverages multi-resolution codebooks at the receiver and transmitter side to sweep beams with an increasing angular resolution as depicted in Fig. 8c. Hierarchical beam search greatly reduces the search space but it is known to be prone to compounding errors that can mislead the search.

In Fig. 9 we compare the above beam alignment schemes and report the fraction of throughput obtained as the sweeping progresses. In our comparisons, we consider all the components affecting the search time of each protocol as provided in Eqs. (10) and (11). In particular, we benchmark the algorithms over the s009 portion of data while we use the s008 dataset to train the neural network model and to build the site-specific database for inverse fingerprints. We assume that a sub-6GHz link is available with a data rate of 0.1 Mb/sec and that each coordinate is quantized into 16 bits. We also endow the RSU and the vehicle with consumer-grade computing capabilities and accordingly set $T_{inf} = 0.75$ ms and $T_{Qr} = 0.3$ ms.

In Fig. 9, we increase the number of candidate beams K to achieve different points on the throughput ratio versus time curves, for our proposed approach and inverse fingerprinting. For hierarchical beam search, we consider the RSU and vehicle equipped with a DFT-based three-level codebook as in [22]. We denote the codebooks at the RSU by $\{C_t^1, C_t^2, C_t^3\}$ and at the vehicle by $\{C_r^1, C_r^2, C_r^3\}$. We fix the corresponding codebook sizes to be, respectively, $\{8, 16, 32\}$ at the RSU and $\{2, 4, 8\}$ at the vehicle. The RSU and vehicle first sweep the beams in $C_t^1 \times C_r^1$ and then refine the search over the angular sector that results in the largest channel gain using the subset of beams in $C_t^2 \times C_r^2$. The refinement procedure is then repeated using the beams in $C_t^3 \times C_r^3$ to produce the final beam. For this specific choice of codebooks the search procedure ends after sweeping through 24 beam pairs. As expected, the throughput ratio sharply improves when switching from a lower resolution to a higher resolution one.

Overall, the delay introduced by the data-driven methods is very modest, and, at the same time, utilizing sensor information yields high throughput beam pairs with minimal search. For this reason, despite the initial offset due to the processing of the side information and the transmission of the predicted best beam directions, the inverse fingerprinting and the proposed method outperform both the exhaustive and hierarchical search methods. In this scenario, our proposed approach almost halves the search time compared to the inverse fingerprinting scheme to attain a 90% throughput ratio. This shows the benefits of using the additional LIDAR side information in comparison with the inverse fingerprinting scheme that only uses the position data.

VI. CONCLUSIONS

We have proposed a supervised learning scheme for efficient mmWave beam selection that exploits side information in the form of LIDAR and position data to reduce the beam search overhead. Our approach significantly outperformed the state of the art in terms of beam classification accuracy and resulting throughput. We have introduced a non-local attention block to improve the performance, specifically for the more challenging NLOS scenarios. Additionally, we have proposed a curriculum learning strategy and a novel loss function inspired by knowledge distillation, which improved the training speed and the accuracy of the final solution. Our NN-based beam selection scheme was able to harness almost 79.9% of the available throughput without any beam search, just by utilizing the LIDAR and location side information. In a typical mmWave V2I scenario, our proposed method considerably reduces the beam search time required to achieve a desired throughput, in comparison with the inverse fingerprinting and hierarchical beam selection schemes. Our




Fig. 9: Throughput ratio versus beam sweeping time. The proposed and the inverse fingerprint schemes curves are obtained evaluating the schemes for different values of K . They introduce a modest overhead to the overall beam training but they are able to perform a more efficient informed search compared to the hierarchical and exhaustive sweeping algorithms.

NN architecture can be further simplified by pruning up to 60% of its parameters without considerable performance loss, hence, rendering it suitable for practical deployment.


REFERENCES

- [1] J. Choi, V. Va, N. Gonzalez-Prelcic, R. Daniels, C. R. Bhat, and R. W. Heath, "Millimeter-wave vehicular communication to support massive automotive sensing," *IEEE Communications Magazine*, vol. 54, no. 12, pp. 160–167, 2016.
- [2] V. Va, J. Choi, and R. W. Heath, "The impact of beamwidth on temporal channel variation in vehicular channels and its implications," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 6, pp. 5014–5029, 2016.
- [3] L. Wei, Q. Li, and G. Wu, "Initial access techniques for 5G NR: Omni/beam SYNC and RACH designs," in *2018 International Conference on Computing, Networking and Communications (ICNC)*. IEEE, 2018, pp. 249–253.
- [4] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh, "Multilevel millimeter wave beamforming for wireless backhaul," in *2011 IEEE Globecom Workshops*. IEEE, 2011, pp. 253–257.
- [5] 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.
- [6] M. Brambilla, M. Nicoli, S. Savaresi, and U. Spagnolini, "Inertial sensor aided mmWave beam tracking to support cooperative autonomous driving," in *2019 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2019, pp. 1–6.
- [7] I. Mavromatis, A. Tassi, R. J. Piechocki, and A. Nix, "Beam alignment for millimetre wave links with motion prediction of autonomous vehicles," in *Antennas, Propagation RF Technology for Transport and Autonomous Platforms 2017*, 2017, pp. 1–8.
- [8] N. González-Prelcic, R. Méndez-Rial, and R. W. Heath, "Radar aided beam alignment in mmWave V2I communications supporting antenna diversity," in *2016 Information Theory and Applications Workshop (ITA)*, 2016, pp. 1–7.
- [9] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-Wave beam steering without in-band measurement," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 2416–2424.
- [10] A. Ali, N. González-Prelcic, and R. W. Heath, "Millimeter wave beam-selection using out-of-band spatial information," *IEEE Transactions on Wireless Communications*, vol. 17, no. 2, pp. 1038–1052, 2018.
- [11] M. Alrabeiah and A. Alkhateeb, "Deep learning for mmWave beam and blockage prediction using sub-6 GHz channels," *IEEE Transactions on Communications*, vol. 68, no. 9, pp. 5504–5518, 2020.

- [12] I. Chafaa, R. Negrel, E. V. Belmega, and M. Debbah, "Federated channel-beam mapping: from sub-6ghz to mmwave," in *IEEE WCNC 2021 Workshop: Distributed Machine Learning*, 2021.
- [13] 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. 209 619–209 629, 2020.
- [14] M. Alrabeiah, A. Hredzak, and A. Alkhateeb, "Millimeter wave base stations with cameras: Vision-aided beam and blockage prediction," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*. IEEE, 2020, pp. 1–5.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] M. B. Mashhadi, M. Jankowski, T.-Y. Tung, S. Kobus, and D. Gündüz, "Federated mmWave beam selection utilizing LIDAR data," *IEEE Wireless Communications Letters*, vol. 10, no. 10, pp. 2269–2273, 2021.
- [19] *AI/ML in 5G*, accessed April 2021, <https://www.itu.int/en/ITU-T/AI/challenge/2020/Pages/default.aspx>.
- [20] *AI/ML in 5G*, accessed April 2021, <https://www.itu.int/en/ITU-T/AI/challenge/2020/Pages/results.aspx>.
- [21] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *NIPS 2014 Deep Learning Workshop.*, 2014.
- [22] S. Noh, M. D. Zoltowski, and D. J. Love, "Multi-resolution codebook and adaptive beamforming sequence design for millimeter wave beam alignment," *IEEE Transactions on Wireless Communications*, vol. 16, no. 9, pp. 5689–5701, 2017.
- [23] 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, 2017.
- [24] *Raymobtime*, accessed November 2020, <https://www.lasse.ufpa.br/raymobtime/>.
- [25] A. Klautau, P. Batista, N. González-Prelcic, Y. Wang, and R. W. Heath, "5G MIMO data for machine learning: application to beam-selection using deep learning," in *2018 Information Theory and Applications Workshop (ITA)*, 2018, pp. 1–9.
- [26] P. I. Yonglong Tian, Dilip Krishnan, "Contrastive representation distillation," in *Proceedings of the International Conference on Learning Representations, ICLR 2020, April 2020, Addis Ababa, Ethiopia*.
- [27] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 41–48.
- [28] G. Hacohen and D. Weinshall, "On the power of curriculum learning in training deep networks," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2535–2544.
- [29] A. Buades, B. Coll, and J. . Morel, "A non-local algorithm for image denoising," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 2, 2005, pp. 60–65 vol. 2.
- [30] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7794–7803.
- [31] Y. LeCun, J. S. Denker, and S. A. Solla, "Optimal brain damage," in *Advances in Neural Information Processing Systems*, 1990, pp. 598–605.
- [32] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both weights and connections for efficient neural networks," in *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, 2015, pp. 1135–1143.
- [33] E. Dahlman, S. Parkvall, and J. Skold, *5G NR: The next generation wireless access technology*. Academic Press, 2020.



Zecchin Matteo (Student Member, IEEE) received the BSc degree in Information Engineering from Università di Padova, Padova, Italy, in 2017 and the MSc degree in Telecommunications Engineering from Università di Padova and National Taiwan University in 2019, both cum laude. He is currently an Early Stage Research in the H2020 Project Windmill and a Ph.D student within the Department of Communication Systems, EURECOM, Sophia Antipolis, France.



Mahdi Boloursaz Mashhadi [S'14-M'20] received a dual major B.S. degree (Hons.) in electrical and industrial engineering and the M.S. and Ph.D. degrees in electrical engineering from Sharif University of Technology (SUT), Tehran, Iran, in 2011, 2013 and 2018, respectively. He is currently a lecturer (assistant professor) at the 5GIC & 6GIC, Institute for Communication Systems (ICS), University of Surrey, UK. He previously worked as a post-doctoral research associate at the Intelligent Systems and Networks (ISN) research group, Imperial College London, UK, and as a visiting research associate at the University of Central Florida (UCF), and Queens' University, Canada. His research interests include the areas of wireless communications, machine learning, and sparse/statistical signal processing. His awards include the Best Paper Award at the EWDTS 2012 conference, and several national, regional, and IEEE grants. He serves as a reviewer for IEEE Transactions on Wireless Communications, IEEE Transactions on Communications, and IEEE Transactions on Signal Processing.



Mikolaj Jankowski (Graduate Student Member, IEEE) received the B.Sc. degree in control and robotics from the Warsaw University of Technology in 2016, and the M.Sc. degree in control systems from Imperial College London in 2018. He is currently pursuing the Ph.D. degree with Intelligent Systems and Networks Group, under the supervision of Prof. Gündüz and Prof. Mikolajczyk. His main research areas are machine learning, computer vision, image processing and information theory. He is currently working on data compression optimized

for various machine learning tasks and applying machine learning to improve wireless communications.



Deniz Gündüz [S'03-M'08-SM'13] received the B.S. degree in electrical and electronics engineering from METU, Turkey in 2002, and the M.S. and Ph.D. degrees in electrical engineering from NYU Tandon School of Engineering (formerly Polytechnic University) in 2004 and 2007, respectively. After his PhD, he served as a postdoctoral research associate at Princeton University, as a consulting assistant professor at Stanford University, and as a research associate at CTTC in Barcelona, Spain. In Sep. 2012, he joined the Electrical and Electronic Engineering

Department of Imperial College London, UK, where he is currently a Professor of Information Processing, and serves as the deputy head of the Intelligent Systems and Networks Group. He is also a part-time faculty member at the University of Modena and Reggio Emilia, Italy, and has held visiting positions at University of Padova (2018-2020) and Princeton University (2009-2012).

His research interests lie in the areas of communications and information theory, machine learning, and privacy. Dr. Gündüz is a Distinguished Lecturer for the IEEE Information Theory Society (2020-22). He is an Area Editor for the IEEE Transactions on Communications and the IEEE Journal on Selected Areas in Communications (JSAC). He also serves as an Editor of the IEEE Transactions on Wireless Communications. He is an organiser and co-chair of the 2021 London Symposium on Information Theory. He is the recipient of the IEEE Communications Society - Communication Theory Technical Committee (CTTC) Early Achievement Award in 2017, a Starting Grant of the European Research Council (ERC) in 2016, and several best paper awards.



David Gesbert (Fellow, IEEE) received the Ph.D. degree from the Ecole Nationale Supérieure des Telecommunications, France, in 1997. He is currently a Professor and the Head of the Communication Systems Department, EURECOM. From 1997 to 1999, he was with the Information Systems Laboratory, Stanford University. He was the Founding Engineer of Iospan Wireless Inc., a Stanford spin off pioneering MIMO-OFDM (currently Intel). Before joining EURECOM in 2004, he was with the Department of Informatics, University of Oslo, as

an Adjunct Professor. He has published about 340 articles and 25 patents, some of them winning the 2019 ICC Best Paper Award, the 2015 IEEE Best Tutorial Paper Award (Communications Society), the 2012 SPS Signal Processing Magazine Best Paper Award, the 2004 IEEE Best Tutorial Paper Award (Communications Society), the 2005 Young Author Best Paper Award for Signal Processing Society journals, and the paper awards at conferences 2011 IEEE SPAWC and 2004 ACM MSWiM. He has been the Technical Program Co-Chair for ICC2017. He was named a Thomson-Reuters Highly Cited Researchers in computer science. He was on the Advisory Board of the Huawei European Research Institute. He is a Board Member for the OpenAirInterface (OAI) Software Alliance. In 2015, he was awarded the ERC Advanced Grant "PERFUME" on the topic of smart device communications in future wireless networks. Since early 2019, he has been heading the Huawei-funded Chair on Advanced Wireless Systems Towards 6G Networks. In 2020, he was awarded funding by the French Interdisciplinary Institute on Artificial Intelligence for a Chair in the area of AI for the future IoT.



Marios Kountouris [S'04-M'08-SM'15] received the diploma degree in electrical and computer engineering from NTUA, Athens, Greece in 2002 and the M. S. and Ph.D. degrees in electrical engineering from Télécom Paris, France in 2004 and 2008, respectively. He is currently a Professor and Chair PI on Advanced Wireless Systems at EURECOM, France. Prior to his current appointment, he has held positions at Huawei Paris Research Center, France, CentraleSupélec, France, The University of Texas at Austin, USA, and Yonsei University, S. Korea. He

was awarded a consolidator grant by the European Research Council (ERC) in 2020 on goal-oriented semantic communication. He has served as Associate Editor for the IEEE Transactions on Wireless Communications, the IEEE Transactions on Signal Processing, and the IEEE Wireless Communication Letters. He has received several awards, including the 2020 IEEE ComSoc Young Author Best Paper Award, the 2016 IEEE ComSoc CTTC Early Achievement Award, the 2013 IEEE ComSoc Outstanding Young Researcher Award for the EMEA Region, the 2012 IEEE SPS Signal Processing Magazine Award, the IEEE SPAWC 2013 Best Paper Award, and the IEEE Globecom 2009 CT Best Paper Award.