IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received January 1, 2019, accepted January 22, 2019, date of publication January 31, 2019, date of current version February 22, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2895594

Learning and Data-Driven Beam Selection for mmWave Communications: An Angle of Arrival-Based Approach

CARLES ANTÓN-HARO, (Senior Member, IEEE), AND XAVIER MESTRE, (Senior Member, IEEE)
Centre Tecnològic de Telecomunicacions de Catalunya(CTTC/iCERCA), Parc Mediterrani de la Tecnologia, 08860 Castelldefels, Spain
Corresponding author: Carles Antón-Haro (carles.anton@cttc.es)

This work was supported by the European Commission through the Inclusive Radio Communications (IRACON) COST action.

ABSTRACT This paper investigates how angle-of-arrival (AoA) information can be exploited by deep-/machine-learning approaches to perform beam selection in the uplink of a mmWave communication system. Specifically, we consider a hybrid beamforming setup comprising an analog beamforming (ABF) network with adjustable beamwidth followed by a zero-forcing baseband processing block. The goal is to select the optimal configuration for the ABF network based on the estimated AoAs of the various user equipments. To that aim, we consider 1) two supervised machine-learning approaches: k-nearest neighbors (kNN) and support vector classifiers (SVC); and 2) a feed-forward deep neural network: the multilayer perceptron. We conduct an extensive performance evaluation to investigate the impact of the quality of CSI estimates (AoAs and powers) obtained via the Capon or MUSIC methods, fluctuations in the received power, the size of the training dataset, the total number of analog beamformers in the codebook, their beamwidth, or the number of active users. The computer simulations reveal that performance, in terms of classification accuracy and sum-rate is very close to that achievable via exhaustive search.

INDEX TERMS Machine learning, classification algorithms, MIMO.

I. INTRODUCTION

Machine learning (ML) techniques have been known for a long time as powerful tools for classification and regression (prediction) problems. More recently, deep learning (DL) has emerged with more advanced tools capable of building universal classifiers and/or approximate general functions. Typical problems/scenarios where machine learning methods have been successfully applied include, but are not limited to, image restoration and identification, natural language processing, network security, customer segmentation, predictive maintenance (e.g. for machinery in industrial plants), etc.

Over the last two decades, the application of ML/DL techniques to communication problems has been to a large extent confined to the field of wireless *network optimization* [1]. Consequently, there exists a large body of literature devoted to problems like intelligent resource management, cell association, selection of radio access technologies, or spectrum management, to name a few. More recently, the interest in using ML/DL techniques for problems and functionalities related with the physical layer (PHY) of communication systems (e.g., coding, modulation, detection, equalization, pre-coding, among others) has dramatically increased. And, further, it is generally agreed that enhancing (or even replacing) some PHY functionalities with ML and DL approaches could help achieve the stringent requirements associated with the future releases of 5G [2]. For this reason, O'Shea and Hoydis [3] have started investigating how autoencoders can be used to model an end-to-end communication system comprising the encoding, channel and decoding blocks, for MIMO point-to-point and interference channels. The adoption of feedforward deep neural networks for joint channel estimation and data detection under an MMSE criterion is addressed in [4]. Other research works include the adoption of convolutional neural networks (CNN) to carry out sensing, compression and recovery of channel state information (CSI) in the feedback channel of massive MIMO systems [5]; the use of support vector classifiers (SVC) to perform antenna [6] or beam selection [7], or the exploitation of CNNs for modulation classification tasks [3].

The shift towards communications in mmWave bands makes it feasible to implement transceivers with many more antenna elements. In order to keep hardware cost and

The associate editor coordinating the review of this manuscript and approving it for publication was Mahdi Zareei.

computational complexity at reasonable levels, a number of techniques have been proposed. Hybrid beamforming, for instance, can be used to partition beamforming into the digital and analog (RF) domains and, by doing so, reduce the number of *complete* (i.e., equipped with analog to digital converters) RF chains needed in the base station. In this context, the development of efficient (analog) beam selection techniques has become a hot research area. In [8], for instance, the authors present an iterative scheme where, in each iteration, the beam that contributes less to the sum-rate of a multi-user MIMO system with a discrete lens array and zero-forcing digital pre-coding block is eliminated. In [9], instead, the proposed low-complexity beam selection algorithm exploits the fact that an approximation of the sum-rate can be computed out of the elements of the effective channel matrix at the output of the analog beamforming (ABF) block. The approximation is near optimal in the high-SNR regime. A different approach is adopted in [10], where beam selection is accomplished with a bio-inspired ant colony optimization algorithm of a very reduced complexity. The applicability of ML techniques to beam selection tasks is investigated in [7]. Specifically, the authors model beam selection (in a hybrid receiver architecture with a ZF block) as an instance of a multi-class classification problem, and attempt to solve it with a support vector classifier (SVC). To that aim, the availability of the actual angle-of-arrival (AoA) and amplitude information for each scattering path is assumed. Complementarily, in Klautau*et al* [11] investigate the performance of several ML schemes (SVC, Adaboost, Decision trees, Random forests) but also deep neural networks (DNN) and reinforcement learning for beam pair selection in vehicle-to-infrastructure settings. The proposed method leverages on positioning information (vehicle locations) which is assumed to be supplied by an external entity and be error-free.

Transmit *antenna selection* (TAS) is also an effective way of reducing the number of complete RF chains needed. Joung [6] propose to solve the combinatorial TAS problem with ML techniques: support vector classifier or k-Nearest Neighbors (kNN) [12]. Interestingly, their symbol error rate performance is acceptable when compared with the benchmark, an optimization-based technique based on singular-value decomposition, at a much lower complexity. Complementarily, [13] focuses on MIMO wiretap channels and resorts to SVM and Naive Bayes classifiers as a means to maximize the secrecy performance. Finally, Mukherjee and Hottinen [14] adopt a multi-armed bandit framework (reinforcement learning) to maximize throughput in MIMO point-to-point channels via antenna subset selection.

A. CONTRIBUTION

In this paper, we investigate how angle-of-arrival (AoA) information can be exploited by a number of deep and machine learning-based analog beam selection schemes for the uplink of a multi-user MIMO mmWave communication system. As in [7], [8], and [10], we adopt a hybrid beamforming architecture comprising an analog beamforming (ABF) network plus a ZF baseband (digital) beamforming block (BBF). Unlike [7] and [11] and many other works, here AoA (and received power) information is not assumed to be known or externally supplied. Instead, it is estimated by the same system via Capon (minimum variance) or MUSIC (MUltiple SIgnal Classification) spectral estimation [15] which is far more realistic. In this context, we investigate whether (i) the resolution provided by the aforementioned methods suffices to perform data-based analog beam selection; and (ii) such estimates can be successfully obtained from a filtered version of the received signal. To perform beam selection, we resort to a number of ML/DL approaches this including support vector classifiers (like in [7] and [11]) but also consider other schemes like k-nearest neighbors (that was used in [6] for transmit antenna selection) and the multi-layer perceptron (MLP). The rationale behind this choice is the different degrees of complexity/sophistication exhibited by those three schemes.. Unlike in many previous works, we propose to use a codebook of analog beamformers with adjustable beamwidth (rather than phased-arrays or DFT-based codebooks). And, more importantly, we analyze the interplay between codebook design and the performance exhibited by the ML/DL-based analog beam selection methods. As benchmarks, we consider a scheme where the optimal beam is identified via exhaustive search, and random beam selection. Differently from our previous work [16], here we (i) investigate how a super-resolution spectral estimation method like MUSIC can help mitigate the deficiencies of the Capon method (i.e., low-resolution of AoA estimates) that we used in the past; (ii) modify the construction of the codebook of analog beamformers in such a way that the search is restricted to the ones that are more frequently used; and, (iii) we conduct an extensive performance evaluation of the proposed ML/DL schemes, mostly in terms of classification accuracy (i.e., percentage of correct beam selection decisions) and sum-rate. Specifically, we investigate the impact of CSI estimates (namely, estimated AoAs and powers) and fluctuations in the received power, the actual CSI estimation method (Capon, MUSIC) as discussed above, the size of the training dataset, the total number of analog beamformers in the codebook, the beamwidth of such beamformers and, finally, the number of active users.

II. SIGNAL AND SYSTEM MODEL

Consider the uplink of a multi-user SIMO system operating in the mmWave band. As shown in Fig. 1, the base station (BS) is equipped with $N_{\rm BS}$ antennas and serves K mobile terminals (or user equipments, UE). The BS adopts a hybrid beamforming architecture, comprising an analog beamforming (ABF) and a digital/baseband (BBF) beamforming block. The former performs a spatial pre-filtering of the received signal, whereas the latter is aimed at removing inter-user interference. The ABF network provides full connectivity among the set of receive antennas and the $N_{\rm RF}$ radiofrequency chains. Hereinafter, we assume that the number of active users does not exceed the number of available RF chains (i.e., $K \leq N_{\rm RF}$).

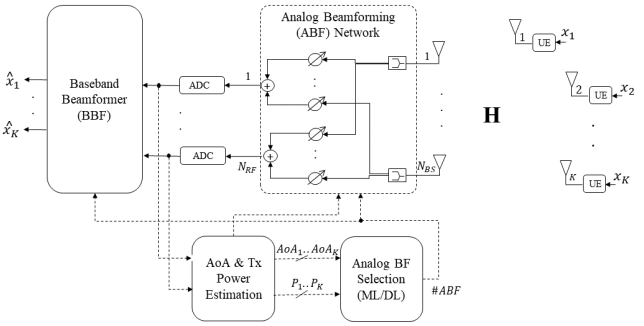

FIGURE 1. System Model.

The received signal at the BS after RF and baseband combining can be expressed as

$$\mathbf{z} = \mathbf{W}^H \mathbf{B}^H \mathbf{H}\mathbf{x} + \mathbf{W}^H \mathbf{B}^H \mathbf{n} \tag{1}$$

where $\mathbf{H} \in \mathbb{C}^{N_{\rm BS} \times K}$ is the channel matrix, $\mathbf{x} \in \mathbb{C}^{K \times 1}$ denotes the transmit signal, with its elements fulfilling $\mathbb{E}\left[\|x_k\|^2\right] = P^{\rm t}$ for $k = 1 \ldots K$; vector $\mathbf{n} \in \mathbb{C}^{N_{\rm BS} \times 1}$ is zero-mean i.i.d. AWGN noise with variance σ^2, i.e., $\mathbf{n} \sim \mathcal{CN}\left(\mathbf{0}, \sigma^2 \mathbf{I}_{N_{\rm BS}}\right)$; and, finally, $\mathbf{W} \in \mathbb{C}^{N_{\rm RF} \times K}$ and $\mathbf{B} \in \mathbb{C}^{N_{\rm BS} \times N_{\rm RF}}$ stand for the BBF and ABF matrices, respectively. For the channel, we adopt a geometric model with L_p scattering paths, which is widely-used in mm-Wave communications. Consequently, for each of the columns of the channel matrix $\mathbf{H} = [\mathbf{h}_1 \ldots \mathbf{h}_K]$, we have that

$$\mathbf{h}_k = \sqrt{\frac{1}{L_p \rho_k}} \sum_{i=1}^{L_p} \alpha_{k,i} \mathbf{a}_{\rm BS}\left(\theta_{k,i}^{\rm BS}\right), \tag{2}$$

where $\alpha_{k,i}$ is the complex gain of the i-th path with $\mathbb{E}\left[|\alpha_{k,i}|\right] = 1$, ρ_k is the path-loss between the BS and the k-th user, and $\theta_{k,i}^{\rm BS} \in [-\pi/2, \ldots, \pi/2]$ denotes the AoA of the i-th path of user k. Vector $\mathbf{a}_{\rm BS}(\theta_{k,i}^{\rm BS})$ is the antenna array response at the BS. Hereinafter, we assume uniform linear arrays (ULA) and, hence, the $N_{\rm BS} \times 1$ vector can be expressed as

$$\mathbf{a}_{\rm BS}\left(\theta_{k,i}^{\rm BS}\right) = \left[1, e^{j\frac{2\pi}{\lambda} d \sin\left(\theta_{k,i}^{\rm BS}\right)}, \ldots, e^{j(N_{\rm BS}-1)\frac{2\pi}{\lambda} d \sin\left(\theta_{k,i}^{\rm BS}\right)}\right] \tag{3}$$

where λ is the signal wavelength, and d is the distance between antenna elements (in the sequel, we assume $d = 0.5\lambda$). Throughout this work, the term SNR stands for the per-antenna signal-to-noise ratio.

III. ANALOG AND DIGITAL BEAMFORMING STRATEGIES

The proposed hybrid beamforming architecture for the BS comprises an analog and baseband (digital) beamforming blocks that are described in this section.

A. ANALOG BEAMFORMING (ABF)

Let $\mathcal{B}$ denote a pre-defined codebook of $B = |\mathcal{B}|$ different configurations of the ABF network (referred to, in the sequel, as *analog beamformers*), namely,

$$\mathcal{B} = \left\{\mathbf{B}_1, \ldots, \mathbf{B}_{|\mathcal{B}|}\right\} \tag{4}$$

with $\mathbf{B}_i \in \mathbb{C}^{N_{\rm BS} \times N_{\rm RF}}$ denoting each of those analog beamformers. Since we are interested in performing a spatial pre-filtering, we let the j-th column vector $\mathbf{b}_{i,j}$ in $\mathbf{B}_i = \left[\mathbf{b}_{i,1}, \ldots, \mathbf{b}_{i,N_{\rm RF}}\right]$ point at a different direction θ_d out of a set of D AoAs evenly spaced in $[-\pi/2 \ldots \pi/2]$. And, as a design decision, we set D in such a way that the resulting number of elements (analog beamformers) in the codebook $B = \binom{D}{N_{\rm RF}}$ is larger than or equal to $N_{\rm BS}$.

A straightforward choice for $\mathbf{b}_{i,j} = \mathbf{b}_{i,j}(\theta_d)$ would be phased arrays, namely,

$$\mathbf{s}(\theta_d) = \frac{1}{\sqrt{N_{\rm BS}}}\left[1, e^{j\pi \sin(\theta_d)}, \ldots, e^{j(N_{\rm BS}-1)\pi \sin(\theta_d)}\right] \tag{5}$$

For large $N_{\rm BS}$, however, the resulting beams become too narrow (to recall, their 3 dB beamwidth is $2/N_{\rm BS}$). As we will discuss in Section IV ahead, this yields a poor classification performance and, thus, must be avoided. Instead, we let $\mathbf{b}$ (sub-indices i, j have been dropped, for brevity) be the solution to the following optimization problem:

$$\begin{aligned} &\underset{\mathbf{b}}{\text{maximize}} \int_{\theta_d - \frac{\Delta\theta}{2}}^{\theta_d + \frac{\Delta\theta}{2}} \mathbf{b}^H \mathbf{s}(\theta)\, \mathbf{s}^H(\theta)\, \mathbf{b}\, d\theta \\ &\text{subject to } \|\mathbf{b}\|^2 = 1. \end{aligned} \tag{6}$$

where $\frac{\Delta\theta}{2}$ is a design parameter related with the interval around θ_d where we want to *broaden* the beam. The score function in (6) can be re-written as $\mathbf{b}\mathbf{C}\mathbf{b}^H$ where we have defined:

$$\mathbf{C} = \int_{\theta_d - \frac{\Delta\theta}{2}}^{\theta_d + \frac{\Delta\theta}{2}} \mathbf{s}(\theta)\, \mathbf{s}^H(\theta)\, d\theta. \tag{7}$$

The (n, m) element in $\mathbf{C}$ thus reads

$$[\mathbf{C}]_{n,m} = \Delta\theta e^{j\pi(n-m)\theta_d} \operatorname{sinc}\left((n-m)\frac{\Delta\theta}{2}\right) \tag{8}$$

with $\operatorname{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$. From all the above, one concludes that $\mathbf{b}$ is given by the eigenvector associated with the largest eigenvalue of matrix $\mathbf{C}$[1].

In Fig. 2, we depict the spatial response of the first element ($\mathbf{B}_0$) in the codebook for different values of $\Delta\theta/2$. Clearly, the larger the value of $\Delta\theta/2$, the broader the beams with (i) minimum beamwidth for $\Delta\theta/2 = 0$ (i.e., phased arrays, top); (ii) directional beams with intermediate beamwidth (and lower sidelobes) for $\Delta\theta/2 = 3/D$ (middle); and (iii) non-directional behavior for $\Delta\theta/2 = 8/D$ (bottom).

Complementarily, in Fig. 3 (left) we show the histograms of analog beamformer occurences with optimal selection (i.e.,

[1]Clearly, the amplitude of each element in $\mathbf{B}_i$ is not necessarily 1. This means that the ABF network can no longer be realized through a set of phase shifters. Instead, each phase shifter must be accompanied by an amplitude control device [17]. Yet this increases the overall cost of the ABF network, for a reduced number of amplitude levels, the cost can be affordable. And, in any case, it is substantially lower than that of using digital beamforming for the whole set of $N_{\rm BS}$ antennas.

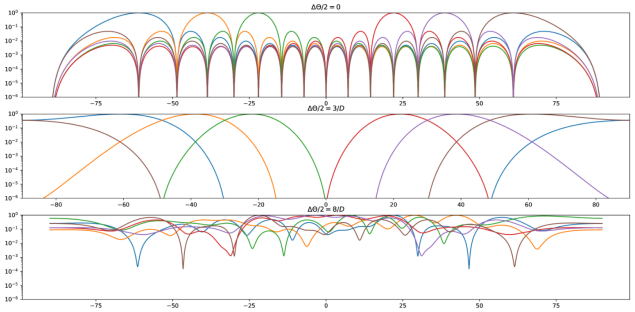

FIGURE 2. Periodogram (logscale) of the first element in the ABF codebook for different values of $\Delta\theta/2$ ($N_{\mathbf{BS}} = 16$, $N_{\mathbf{RF}} = 6$, $D = 8$, $B = 28$).

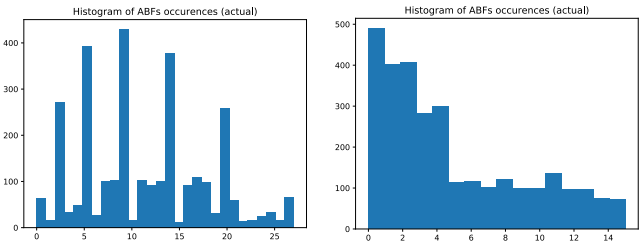

FIGURE 3. Histogram of analog beamformer occurrences with optimal selection and (i) all possible beamformers ($B = 28$, left); and (ii) a subset of most popular ABFs ($B_{\mathbf{sub}} = 16$, right) of the codebook ($N_{\mathbf{BS}} = 16$, $N_{\mathbf{RF}} = 6$, $K = 4$,$D = 8$).

finding out the beamformer with the highest sum-rate, see next section, via exhaustive search). Interestingly, the distribution is not uniform over analog beamformers. This stems from the fact that, even if the AoAs of the various users are uniformly distributed in $[-\pi/2, \ldots, \pi/2]$ the actual selection depends on the interplay of the AoAs of the user of interest *and* those of the interferers (and their number). Since the distribution is not uniform, it makes sense to drop those analog beamformers which are seldom used and restrict the search to the subset of B_{sub} most *popular* ones. By doing so, the computational complexity associated to the selection process decreases and the resulting models to be fit can be of a lower dimension (e.g., layers and/or neurons per layer in the case of deep neural networks). In Fig. 3 (right), we show the histogram of beamfomer occurrences *after* restricting the search to the $B_{\text{sub}} = N_{\text{BS}} = 16$ most popular ones. Clearly, the distribution is more homogeneous now. However, using smaller codebooks has some impact on system performance, which is scenario-dependent. This will be discussed in detail in the computer simulation results section ahead.

Finally, Figure 4 below depicts the periodogram of the *set* of analog beamformers in a codebook designed according to (6).

B. BASEBAND BEAMFORMING (BBF)

For baseband beamforming, we adopt a zero-forcing ZF scheme, as e.g., in [7]. By defining the effective channel matrix as the product of the actual channel and the analog

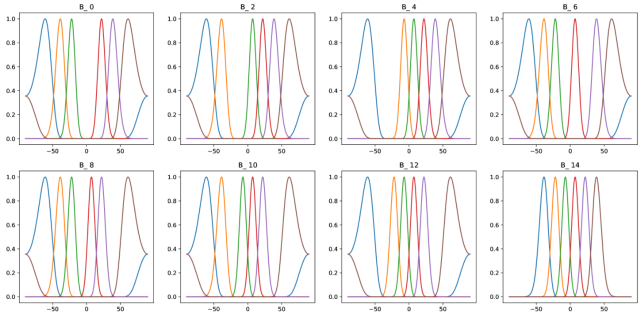

FIGURE 4. Periodograms of selected elements in the ABF codebook ($N_{\mathbf{BS}} = 16$, $N_{\mathbf{RF}} = 6$, $\Delta\theta/ = 3/D$, $D = 8$).

beamforming matrix, namely, $\mathbf{H}_{\text{eff}} = \mathbf{B}^H\mathbf{H}$, we have

$$\mathbf{W} = \left(\mathbf{H}_{\text{eff}}^H\mathbf{H}_{\text{eff}}\right)^{-1}\mathbf{H}_{\text{eff}}^H. \tag{9}$$

IV. PROBLEM STATEMENT

Our goal is to select the analog beamformer from the B-element codebook $\mathcal{B}$ such that the sum rate at the output of the BBF block is maximized. From the system model presented in the preceding section, the sum rate can be expressed as:

$$\begin{aligned} R(b) &= \sum_{k=1}^{K} \log_2\left(1+\gamma_k(b)\right) \\ &= \sum_{k=1}^{K} \log_2\left(1+\frac{P^{\text{t}}/\rho_k}{\sigma^2\left[\left(\mathbf{H}_{\text{eff}}(b)^H\mathbf{H}_{\text{eff}}(b)\right)^{-1}\right]_{kk}}\right) \end{aligned} \tag{10}$$

where $[\mathbf{A}]_{kk}$ denotes the k-the element in the diagonal of matrix A, and where the dependency of $\mathbf{H}_{\text{eff}}$ on the codebook index b has been made explicit.

In order to select the optimal beamformer (and avoid exhaustive searches over the codebook elements), one needs to leverage on H as channel state information (CSI). However, under the assumption of a hybrid beamforming architecture, H cannot be readily estimated in the digital domain (since the number of RF chains, $N_{\mathbf{RF}}$, is smaller than the number of antennas, $N_{\mathbf{BS}}$). Alternatively, we propose to use some sufficient statistic derived from $\mathbf{H}$. In particular, we investigate whether the AoA and received power for each user terminal can be exploited. The rationale behind can be found in Fig. 5 which depicts which is the optimal analog beamformer (found via exhaustive search) for each tuple of AoAs for the $K = 2$ user case. Interestingly, for a suitable value of $\Delta\theta/2 = 3$ (middle plot), the analog beamformers that maximize the sum-rate can be clustered into a small number of *connected* regions (one per beamformer). Hence, the selection process can be modeled as a classification problem that, in turn, can be effectively solved with machine and deep learning tools (see Section V ahead).

The appearance of the scatter plot above these lines critically depends on the beamwidth of the pre-defined set of analog beamformers (see discussions in the preceding

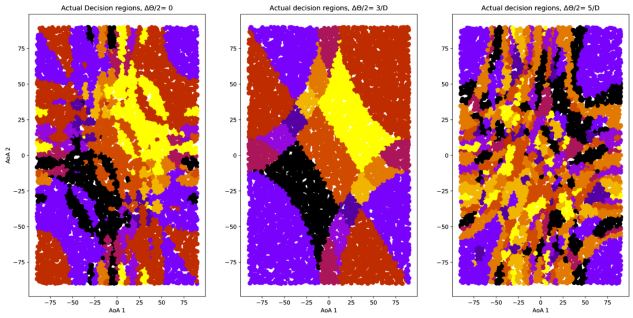

FIGURE 5. Actual decision regions with identical received powers and different values of $\Delta\theta/2$. The different colors are associated with different elements in the codebook ($N_{\mathrm{BS}} = 10$, $N_{\mathrm{RF}} = 3$, $K = 2$, SNR = 20dB, $B = 10$, $D = 5$).

section). For instance, a codebook built on phased arrays (i.e., $\Delta\theta = 0$, left plot in Fig. 5) generates many more connected regions, actually more than one per beamfomer. This stems from the fact that, for very narrow beams, multiple non adjacent beamformers yield comparable sum rates. Determining the boundaries between beamformer regions, the task that the classification algorithm is confronted with, becomes much more difficult. And, hence, performance experiences a very severe degradation. This is why, in Section III-A, we designed an ad-hoc codebook of beamformers with larger beamwidths. Likewise, for *non-directional* analog beamformers
($\Delta\theta/2 = 5/D$, right plot), the number of connected regions explodes. In fact, for non-directional beamformers, the directional information provided by AoAs is not a sufficient statistic anymore.

A. AoA AND RECEIVED POWER ESTIMATION: CAPON VS. MUSIC

By resorting to the Capon (i.e., minimum variance) power spectral estimator [15], we can effectively estimate the set of AoAs and received powers. To that aim, we use the $N_{\mathrm{RF}} \times 1$ baseband signal $\mathbf{y} = \mathbf{B}^H\mathbf{H}\mathbf{x} + \mathbf{B}^H\mathbf{n}$, that is, a *filtered* version of the received signal *after* ABF and analog to digital conversion (to recall, the received signal itself lies in the RF domain). Hence, the power spectral density reads

$$\mathbf{P}_b(\theta) = \frac{1}{\mathbf{s}^H(\theta)\mathbf{B}_{\mathrm{b}}\left(\mathbf{R}_{\mathrm{yy}}\right)^{-1}\mathbf{B}_{\mathrm{b}}^H\mathbf{s}(\theta)} \tag{11}$$

$$= \frac{1}{\mathbf{s}^H(\theta)\mathbf{B}_{\mathrm{b}}\left(\mathbf{B}_{\mathrm{b}}^H\mathbf{R}_{\mathrm{rr}}\mathbf{B}_{\mathrm{b}}\right)^{-1}\mathbf{B}_{\mathrm{b}}^H\mathbf{s}(\theta)} \tag{12}$$

with $\mathbf{R}_{\mathrm{yy}}$ and $\mathbf{R}_{\mathrm{rr}}$ standing for the covariance matrices of the filtered and received signals, respectively. It is worth noting that a *directional* analog beamformer $\mathbf{B}$, such as those designed in Section III-A, might cancel out (part of or all) the received signals. To avoid that, a set of non-directional beamformers $\mathbf{B}_{\mathrm{b}}$ must be used for AoA estimation. Specifically, we let $\mathbf{B}_{\mathrm{b}} = \mathbf{Q}\left(\mathbf{Q}^H\mathbf{Q}\right)^{\frac{1}{2}}$, where $\mathbf{Q}$ is a $N_{\mathrm{BS}} \times N_{\mathrm{RF}}$ matrix of complex-valued random i.i.d. Gaussian entries. Bearing all the above in mind, the estimated AoAs and received powers ($\hat{\theta}_k$ and $\hat{P}_k$ for $k = 1 \ldots K$) are given, respectively, by the arguments and function values associated to the maxima of $\mathbf{P}_b(\theta)$.

However, the spatial resolution of the Capon method is rather limited. To alleviate this, one can resort to super-resolution methods like MUSIC (MUltiple SIgnal Classification) [15]. The main idea behind MUSIC (and other subspace-based methods) consists in finding $\hat{\theta}_k$ as the values of θ such that the *filtered* column vector $\mathbf{B}_{\mathrm{b}}^H\mathbf{s}(\theta)$ lies on the signal subspace of $\mathbf{R}_{\mathrm{yy}}$. We can express the true covariance matrix of the filtered signal as

$$\mathbf{R}_{\mathrm{yy}} = \mathbf{E}_S\mathbf{\Lambda}_S\mathbf{E}_S^H + \sigma\mathbf{E}_N\mathbf{E}_N^H \tag{13}$$

where Λ_S is a $K \times K$ diagonal matrix containing the largest eigenvalues $\lambda_1 \geq \lambda_2 \geq \ldots \lambda_K$ of $\mathbf{R}_{\mathrm{yy}}$; $\mathbf{E}_S$ is an $N_{\mathrm{RF}} \times K$ matrix that contains the (signal) eigenvectors corresponding to the K largest eigenvalues; and $\mathbf{E}_N$ is an $N_{\mathrm{RF}} \times (N_{\mathrm{RF}} - K)$ matrix that contains the (noise) eigenvectors associated with the smallest (noise) eigenvalue σ^2. Hence, $\mathbf{B}_{\mathrm{b}}^H\mathbf{s}(\theta)$ lies on the signal subspace of $\mathbf{R}_{\mathrm{yy}}$ for the values of θ such that

$$\mathbf{s}^H(\theta)\mathbf{B}_{\mathrm{b}}\mathbf{E}_N\mathbf{E}_N^H\mathbf{B}_{\mathrm{b}}^H\mathbf{s}(\theta) = 0 \tag{14}$$

The estimated AoAs can thus be computed as the arguments associated to the maxima of the function:

$$\Pi_b(\theta) = \frac{1}{\mathbf{s}^H(\theta)\mathbf{B}_{\mathrm{b}}\mathbf{E}_N\mathbf{E}_N^H\mathbf{B}_{\mathrm{b}}^H\mathbf{s}(\theta)} \tag{15}$$

This function, however, is not a proper power spectral density and, consequently, the received powers cannot be directly estimated from its maxima. The following expression [18] can be used instead:

$$\hat{P}_k = \frac{1}{\mathbf{s}^H(\theta_k)\mathbf{B}_{\mathrm{b}}\hat{\mathbf{E}}_S\left(\hat{\Lambda}_S - \hat{\sigma}^2\mathbf{I}_K\right)^{-1}\hat{\mathbf{E}}_S^H\mathbf{B}_{\mathrm{b}}^H\mathbf{s}(\theta_k)} \tag{16}$$

where $\hat{\Lambda}_S$ and $\hat{\mathbf{E}}_S$ denote the eigenvalues and eigenvectors of the *sample* covariance matrix, respectively; and where the quantity $\hat{\sigma}^2 = \frac{1}{N_{\mathrm{RF}}-K}\sum_{k=K+1}^{N_{\mathrm{RF}}}\hat{\lambda}_k$ is an estimator of the (filtered) noise power.

Figure 6 depicts the spectral densities obtained by the periodogram, Capon, and MUSIC methods with the received and filtered signals[2]. Unsurprisingly, the spatial resolution of the Capon method is higher than that of the periodogram. However, that of MUSIC is even higher, as evidenced by the set of much narrower peaks. As a result, MUSIC succeeds in separating the two signals impinging on the array from $-80.3°$ and $-75.8°$ whereas the Capon method merges them together. This, in turn, causes a gross error in the AoA estimate since one of those AoAs is attributed to some other (noisy) maximum in the power spectral density. We also observe that the impact of using a filtered version of the signal (MUSIC filtered vs. MUSIC) is negligible: the spectral

[2]For filtered signals, we use the *first* analog beamformer ($\mathbf{B}_0$) in the non-directional codebook. However, similar results can be obtained with any other codebook elements.

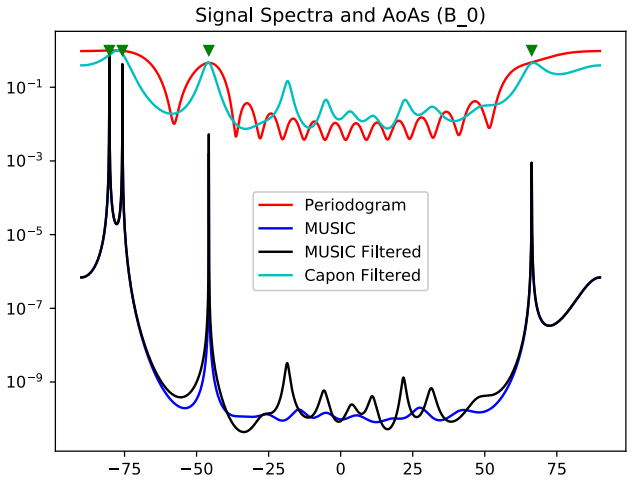

FIGURE 6. Normalized power spectral densities for the periodogram, Capon, and MUSIC methods with the received and filtered signals (non-directional random analog beamformers). Green markers denote the actual AoAs ($N_{BS} = 16$, $N_{RF} = 6$, $K = 4$, SNR = 20dB).

density itself is a bit distorted but, still, the resolution and the location of peaks associated to the angles of arrival (mostly) remain unchanged.

V. DATA-DRIVEN ANALOG BEAM SELECTION

The task of selecting the optimal analog beamformer can be modeled as a supervised learning problem. Specifically it can be cast into a multi-class classification task where:

- The input is given by the set of (estimated) AoAs and received powers for all active users, namely, the row vector $\mathbf{t} = \left[\theta_1, \ldots, \theta_K, P_1^{\mathrm{r}}, \ldots, P_K^{\mathrm{r}}\right]$. The number of active users K is assumed to be known.
- The output is an index to the element (beamformer) in the codebook, $b \in [1, \ldots, B]$, yielding the highest sum rate after ZF. Such an output avoids performing the costly matrix product and inversion in (10) for *every* analog beamformer in the codebook that an exhaustive search would entail.

The next subsections describe (i) the training set to be used; and (ii) the machine and deep learning schemes considered in this work.

A. TRAINING SET DESCRIPTION

The training set comprises a total of M examples $\mathbf{t}_m$ stacked in a training matrix $\mathbf{T} = \left[\mathbf{t}_1^T, \ldots, \mathbf{t}_M^T\right]^T \in \mathbb{R}^{M\times 2K}$ and a class label vector $\mathbf{c} = [c_1, \ldots, c_M]^T \in \mathbb{N}^{M\times 1}$ with the corresponding outputs. The generation of the training set comprises the following steps:

1) Generation the received signal for a total of M realizations of the communication scenario (user locations, transmit powers, path-loss).
2) Collection /estimation of the AoA and receive power information as per Section IV-A.
3) Performing ABF and BBF on the received signal for each example with all the possible analog beamformers in the codebook, as in (1).
4) Computation of the corresponding sum rates as per (10);
5) Letting the label of each example be the index of the analog beamformer yielding the highest sum rate, i.e., $\mathbf{c}_m \in \{1, 2, \ldots, B\}$.

To avoid significant bias in the training, features are normalized prior to their use by the learning schemes, namely, $t_{ij} \leftarrow \left(t_{ij} - E_i t_{ij}\right) / \left(\max_i \left[t_{ij}\right] - \min_i \left[t_{ij}\right]\right)$

B. MACHINE LEARNING SCHEMES

The so-called k-Nearest Neighbors (kNN) scheme is one of the simplest ML-based classification methods. For a new example $\mathbf{t}$, the kNN classifier finds the k nearest[3] entries in the training set $\mathbf{T}$. Then kNN declares the class label for the new example $\mathbf{t}$ based on a (possibly weighted) majority of the nearest neighbors' labels. The only hyper-parameter of the algorithm is G, the number of nearest neighbors to find. Typically, G is set to be an odd number, to avoid ambiguity.

We will also consider a Support Vector Classifier (SVC). For a multi-class classification problems (like ours), SVCs need to solve B binary classification problems. Each of them, identifies one category vs. the other categories (i.e. one-vs.-rest approach). Specifically, during the training phase, binary SVCs fit a separating hyperplane with optimal geometric margin (i.e. maximum distance to the closest elements in each of the two classes). Next, in the test phase, those hyperplanes are used to decide in which region (i.e., class, beamformer index) the new example t lies[4]. Since our problem is not linearly separable (see Fig. 5), we introduce a Gaussian (radial-based) kernel function to transform the input data. This allows to transform hyperplanes into more general boundaries. For this kernel in particular, the hyperparameters here are σ_k, the variance of the Gaussian kernel; and, C, a penalty parameter aimed to balance bias and overrfitting (see [12] for details).

C. DEEP LEARNING SCHEME

Feedforward deep neural networks (DNN) like the multi-layer perceptron (MLP) can be regarded as universal function approximators. As shown in Fig. 7, our L-layer MLP consists of: one input layer with $2K$ elements (AoAs, received powers), one output layer with B neurons (the number of elements in the codebook), and $L-2$ hidden layers with N_l; $l = 2, \ldots, L-1$ neurons each. This MLP defines a mapping $f(\mathbf{t}, \Phi) : \mathbb{R}^{2K} \mapsto \mathbb{R}^B$ of an input vector $\mathbf{t}$ onto an output vector $\mathbf{r}_L$ through L iterative processing

[3] In this work, we restrict ourselves to the most popular distance measure: the Euclidean distance, namely, $d(\mathbf{t}_m, \mathbf{t}) = \|\mathbf{t}_m - \mathbf{t}\|^2$

[4] If binary SVCs predict multiple labels, the one with the highest confidence score is finally selected.

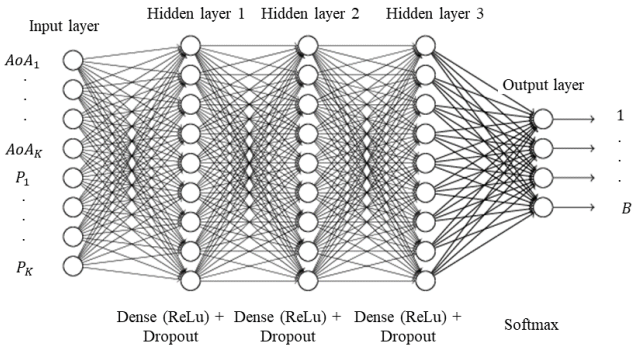

FIGURE 7. Multi-layer Perceptron (MLP), a feedforward deep neural network.

steps:

$$\mathbf{r}_l = f_l\left(\mathbf{r}_{l-1}; \phi_l\right); \quad l = 1 \ldots L \tag{17}$$

where $f_l(\mathbf{r}_{l-1}, \phi_l) : \mathbb{R}^{N_{l-1}} \mapsto \mathbb{R}^{N_l}$ is the mapping associated to the l-th layer. Such mapping depends, on the one hand, on the output from the previous layer and, on the other, on the parameter set of that layer, ϕ_l, with $\Phi = [\phi_1, \ldots, \phi_L]$ denoting the set of all parameters in the model (see below for details). The l-th layer is said to be *dense* if $f_l\left(\mathbf{r}_{l-1}; \phi_l\right)$ is of the form

$$f_l\left(\mathbf{r}_{l-1}; \phi_l\right) = \sigma\left(\mathbf{G}_l \mathbf{r}_{l-1} + \mathbf{u}_l\right) \tag{18}$$

with $\mathbf{G}_l \in \mathbb{R}^{N_l \times N_{l-1}}$ and $\mathbf{u}_l \in \mathbb{R}^{N_l}$ denoting a matrix of weights and a bias vector, respectively; and $\sigma(\cdot)$ is a (typically non-linear) activation function. In this work, Rectified Linear Units (ReLU) are used as activation functions in all hidden layers. On the contrary, a Softmax activation function is used at the output layer, so that the probability for each beamformer can be effectively computed. Further, each hidden layer is followed by a dropout layer in order to avoid overfitting when using the training set. In the training phase, the parameter set in each layer, i.e., $\phi_l = [\mathbf{G}_l, \mathbf{u}_l]$, is adjusted via the so-called stochastic adaptive gradient descent algorithm (Adagrad, [19]) and backpropagation [12]. The goal is to minimize the categorical cross entropy (i.e., maximize classification accuracy).

VI. PERFORMANCE EVALUATION

In this section, we assess the performance of the proposed ML- and DL-based beam selection schemes. The system scenario comprises one BS equipped with $N_{\text{BS}} = 16$ receive antennas and $N_{\text{RF}} = 6$ RF chains. We consider uplink channels with a single scattering path (i.e., $L_p = 1$), same transmit power $P_k = P$ for all k and (possibly) different path-losses for the various users. The angles-of-arrival were uniformly distributed in $[-\pi/2, \ldots, \pi/2]$.

For the implementation, training and evaluation of the proposed schemes, we resorted to the Scikit-learn [20] and Tensorflow [21] libraries, for ML (kNN, SVC) and DL (MLP) ones, respectively. Besides, we also used Keras [22], which runs on top of TensorFlow (and also other DL libraries like Theano or CNTK) and provides a high-level neural network API (Application Programming Interface) to it. Unless otherwise stated, the number of examples in the training and test sets was set to $M = 3000$ and $M_{\text{test}} = 3000$, respectively (details on the training set can be found in Section V-A). As for tuning, the hyper-parameters in each model were optimized via grid search: $G = [3, 5, 10, 20, 50, 80, 100, 200]$, $\sigma_k = \left[10^{-4}, 10^{-3}, 0.01, 0.1, 0.2, 0.5, 1, 5\right]$, and $C = [1, 10, 100, 1000, 2000, 5000]$. To that aim, a cross-validation factor of 5 (i.e., 20% of the examples were reserved for testing) was considered, both for the kNN and SVC schemes. To recall, for the SVC we used a Radial Basis Function as a Kernel, and a 'one-over-the-rest' decision function (see Section V-B) A different model was fit for each SNR. As for the MLP, the number of layers and neurons in the hidden layers was $L = 5$ and $N_l = 125$, respectively. During training, the dropout rate was set to 20% which sufficed to prevent overfitting. The total number of training iterations was set to 200. For the stochastic adaptive gradient descent algorithm (Adagrad), the learning rate was set to 0.01, and a batch size of 32 samples was used for the computation of the gradient (for further details on the configuration of the MLP scheme, please, see Section V-C). All the aforementioned parameters were manually optimized.

Computer simulation results are given in terms of classification accuracy and sum-rate (see next subsections). As benchmarks, we include (i) a *random* beam selection (uniform distribution in $\{1, \ldots, B\}$); and (ii) an *optimal* beam selection, which is accomplished via exhaustive search.

A. CLASSIFICATION PERFORMANCE WITH CAPON SPECTRAL ESTIMATION

Here, we investigate the classification accuracy that can be achieved when the AoAs and received powers are found via Capon Spectral Estimation. First, we focus on the decision regions generated by the kNN scheme (Fig. 8). Clearly, the regions are very similar to the actual ones, which are displayed in Fig.5 since the classification accuracy for this toy example was 85%. Circles with two different colors indicate classification errors (the inner circle is associated to the optimal beam selection decision). Interestingly, most of the errors occur between adjacent regions (e.g. red-yellow, yellow-orange); and, also, the scatterplot is symmetrical with respect to the secondary diagonal since an exchange of AoA among the two active users yields exactly the same solution (beamformer). Similar results were obtained with the SVC and MLP schemes.

Next, we turn our attention to the impact of the receive SNR at the RF antennas on classification accuracy (Fig. 9). The two upper plots illustrate the performance with the *actual* AoAs and received powers, in order to isolate the impact of CSI estimation. First, we observe that classification accuracy clearly depends on algorithmic complexity: MLP outperforms SVC, and SVC outperforms kNN. For identical receive powers

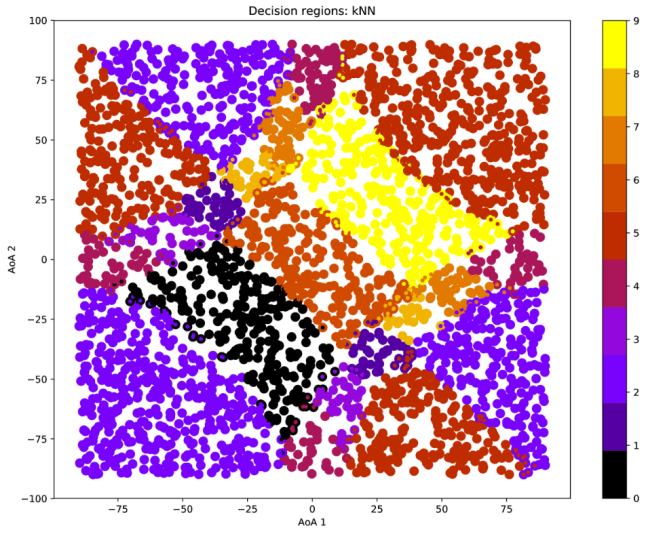

FIGURE 8. Decision regions for the kNN scheme. Circles with two colors indicate classification errors. Received powers are identical for both users ($N_{BS} = 10$, $N_{RF} = 3$, $K = 2$, SNR = 20dB, $B = 10$, $D = 5$, $\Delta\theta/2 = 3/D$).

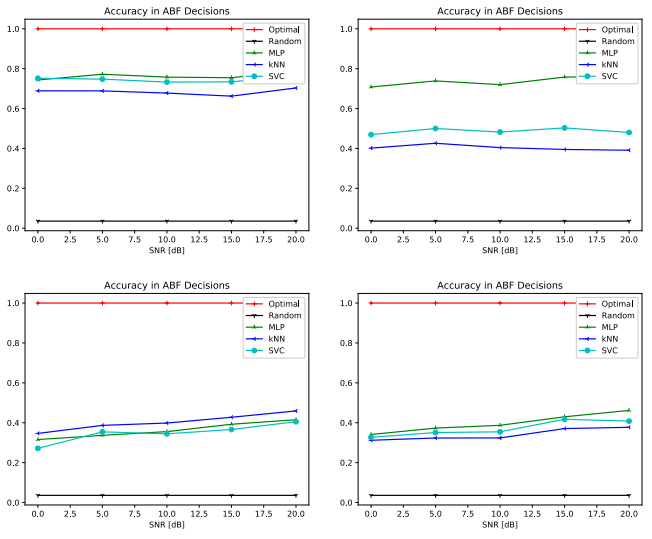

FIGURE 9. Classification accuracy with (i) actual (top) and estimated (bottom) AoAs and receive powers; and (ii) identical (left) or different (right) receive powers ($N_{BS} = 16$, $N_{RF} = 6$, $K = 4$, $B = 28$, $D = 8$, $\Delta\theta/2 = 3/D$, Capon estimator).

(top, left), the performance loss with respect to the optimal classifier is comparable for the three approaches and in the range of 20% to 30%. Still, their performance is radically higher than that of random beam selection (classification accuracy of $1/B$, where B is the cardinality of the codebook). In the presence of fluctuations in the received power of ± 2 dB (e.g., due to imperfect power control and/or different user-to-BS distances), the classification accuracy exhibits some degradation (top, right). This holds, in particular, for the ML approaches. Apparently, the more complex map of decision regions that such fluctuations entail are better handled by the MLP since, in general, DNNs turn out to be more sophisticated classifiers than *classical* ML ones. Besides, the impact of the SNR on performance is negligible. Apparently, the fact that ML/DL models are trained for specific SNRs largely mitigates such impact.

However, performance is notably different with *estimated* AoAs and received powers (via Capon power spectra). Even with identical received powers (bottom, left), the classification accuracy substantially degrades: from $70\% - 80\%$ to $30\% - 40\%$. The main reason for that is two-fold. First of all, we have the limited spatial resolution of the Capon estimator that, in addition, must operate with a filtered version of the received signal (see Section IV-A). In other words, when the AoAs of two users are too close, they are merged together which implies that (i) their received powers are added up; and (ii) the second AoA (and received power) is attributed to any of the *noisy* spectrum peaks shown in Fig. 6. When this occurs, the resulting AoA/received power tuples can be quite different and, thus, the probability of selecting a beamformer which is radically different from the optimal one increases too. Besides, even if users continue to be separable in the angular domain, the additive noise has a direct impact on the quality of AoA and received power estimates. Hence, and unlike the case with actual CSI (top row), classification accuracy improves when SNR increases. Finally, now we observe that MLP experiences difficulties in handling the combined effects of CSI estimation (bottom plots): its performance is now comparable to that of kNN and SVC.

B. PERFORMANCE WITH MUSIC VS. CAPON SPECTRAL ESTIMATION

Next, we analyze the impact of the power spectral estimation strategy on performance. As discussed in Section IV-A, using a high-resolution spectral estimation technique like MUSIC allows to estimate both AoA and received powers more accurately. This, in turn, results into an enhanced classification accuracy for all the proposed ML/DL approaches, as Figure 10 reveals. The gain is particularly high for the MLP scheme, with classification accuracy boosting from 40% to 75% in the high-SNR regime. Besides, the resulting classification accuracy gets very close to that achievable with *actual* CSI. Consequently, hereinafter we will restrict out analysis to a system with (i) MUSIC-based AoA and received power estimation (since it outperforms Capon-based ones); and (ii) different received powers (since this is the most difficult scenario).

Yet classification accuracy is a relevant performance measure, from a communications point of view we are often more interested in sum rates. This is illustrated in Fig. 11. Very interestingly, the sum-rate of the MLP approach (in combination with MUSIC) is virtually identical to that of optimal beam selection. And this happens despite of the fact that its classification accuracy is approximately 80% (see Fig. 10 above). As discussed earlier, this is due to the fact that most classification errors occur among adjacent regions. These regions, by codebook construction (see Section III-A)

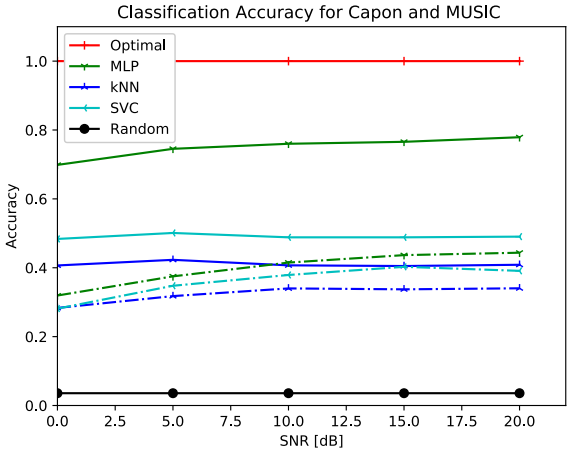

FIGURE 10. Classification accuracy with the MUSIC (solid) and Capon (dash-dotted) methods. Scenario with estimated CSI (AoAs and received powers) and different received powers ($N_{BS} = 16$, $N_{RF} = 6$, $K = 4$, $B = 28$, $D = 8$, $\Delta\theta/2 = 3/D$).

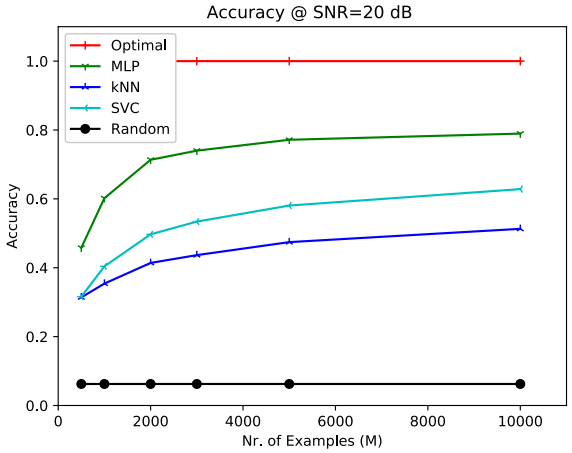

FIGURE 12. Classification accuracy vs. number of examples in the training set ($N_{BS} = 16$, $N_{RF} = 6$, $K = 4$, SNR = 20dB, $B = 28$, $B_{sub} = 16$, $\Delta\theta/2 = 3/D$, $D = 8$, MUSIC estimator).

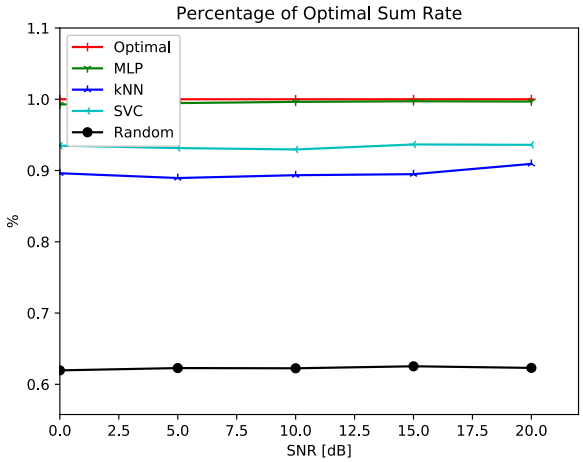

FIGURE 11. Percentage of the sum-rate vs. signal-to-noise ratio ($N_{BS} = 16$, $N_{RF} = 6$, $K = 4$, SNR = 20dB, $B = 28$, $B_{sub} = 16$, $\Delta\theta/2 = 3/D$, $D = 8$, MUSIC estimator).

are often associated to beamformers with a similar spatial response. And, for a given set of AoAs, they result into similar sum-rates. Only when classification accuracy degrades further, for instance for the kNN and SVC approaches, does the sum-rate loss become noticeable. Nevertheless, for this particular scenario, these two ML schemes still manage to retain some 90% of the achievable sum-rate with optimal beam selection. The fact that the DL approach (MLP) outperforms its ML counterparts (kNN, SVC) comes at a price of a higher computational complexity of the classification tool. A more detailed complexity analysis is left for subsequent works.

C. IMPACT OF TRAINING DATASET AND CODEBOOK SIZE

In learning techniques, the size of the training dataset is a parameter of critical importance: it has a direct impact in terms of complexity vs. performance trade-offs, and over/underfitting vs. generalization capabilities. Figure 12 below depicts classification performance as a function of the number of examples in the training dataset (M). For small M, the performance exhibited by all the ML/DL approaches is low due to underfitting. In this region, however, classification accuracy increases rapidly with M and, to a large extent, it stabilizes for $M > 3000$ examples. This holds in particular for the MLP approach. No overfitting problems were observed for the range of dataset sizes of interest (thanks to the use of dropout layers).

One more important aspect to investigate is the impact of the number of elements in the ABF codebook (cf. Section III-A). Restricting the search to the B_{sub} most *popular* elements[5] in the entire codebook, has the beneficial effects of rendering (i) the selection process less complex (fewer choices); and (ii) the dimension of the resulting models lower (and, thus, easier to fit). However, such a strategy has an impact on system performance that we investigate next.

In terms of classification accuracy (Fig. 13, left), the percentage of correct (optimal) decisions tends to increase for a reduced number of codebook elements: the number of beamformers to choose from is lower and so is the probability of making a wrong decision[6].

As for sum-rate (Fig. 13, right), the analysis is more involved. First, from the optimal selection curve, one realizes that reducing the cardinality of the codebook can only be detrimental: the optimal selection in a smaller codebook is potentially different from that of a larger one. And the selection from a larger codebook is likely to mitigate interference more effectively (the more beamformers, the more diverse interference scenarios can be effectively handled). The sum-rate degradation is particularly noticeable

[5]To perform the selection, a sufficiently high number of examples of the scenario under consideration (number of users, with/without fluctuations, etc) are generated. Specifically, the SNR was set to an intermediate value of 10 dB, although the selection barely depends on this parameter.

[6]This holds true in particular for random selection (black curve).

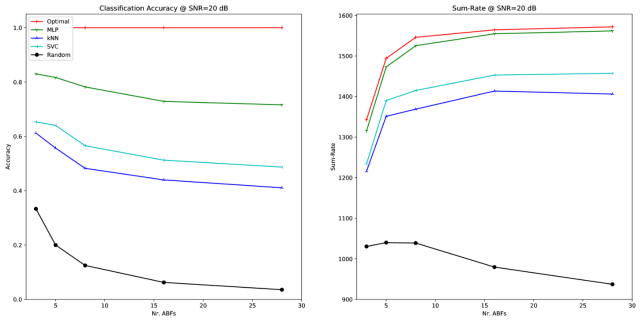

FIGURE 13. Performance vs. number of analog beamformers (B_{sub}) in the codebook: classification accuracy (left) and sum-rate (right) ($N_{\text{BS}} = 16$, $N_{\text{RF}} = 6$, $K = 4$, SNR = 20dB, $B = 28$, $\Delta\theta/2 = 3/D$, $D = 8$, MUSIC estimator).

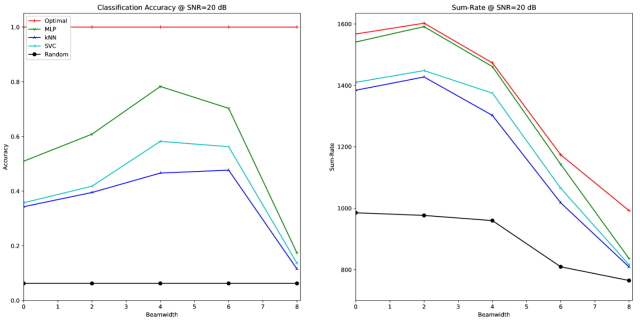

FIGURE 14. Performance vs. beamwidth $\Delta\theta/2$: classification accuracy (left) and sum-rate (right) ($N_{\text{BS}} = 16$, $N_{\text{RF}} = 6$, $K = 4$, SNR = 20dB, $B = 28$, $D = 8$, MUSIC estimator).

for codebooks with $B_{\text{sub}} < 8$ beamformers since such elements are used very often (see histogram of Fig. 3 in Section III-A).

The impact of errors in the selection of ABFs is better illustrated by the random selection curve. For small B_{sub}, the sum-rate associated to each element in the codebook is lower but the probability of making a wrong selection is lower too. For large B_{sub}, it is just the opposite. This suggests the existence of some value of B_{sub} attaining the best trade-off between the two effects (approximately, $B_{\text{sub}} = 8$ in this case).

The proposed ML/DL schemes, of course, are capable of harnessing the benefits of working with larger codebooks in a more effective manner than a mere random selection. This results in a set of curves which are monotonically increasing with the number of available ABFs (as with optimal selection decisions) but with some performance gap, accounting for the sum-rate penalty associated to incorrect selection decisions. It is worth noting that, again, the gap for the whole range of B_{sub} values is negligible for the MLP scheme. And, even more interestingly, the sum-rate degradation for $B_{\text{sub}} \geq N_{\text{BS}}$ (with respect to the $B_{\text{sub}} = B$ case) is barely noticeable. Hence, in the sequel, we let $B_{\text{sub}} = N_{\text{BS}}$.

D. IMPACT OF BEAMWIDTH AND NUMBER OF ACTIVE USERS

In this last part, we investigate the effect of beamwidth on performance. As discussed in previous sections, the classification accuracy associated to phased arrays ($\Delta\theta/2 = 0$) is poor: for narrow beams, the clustering of beamformers into disjoint regions vanishes. On the contrary, using non-directional beams (large $\Delta\theta$) is detrimental too since, in this case, AoA information is not anymore a sufficient statistic to perform beam selection. Both effects can be clearly observed in Fig. 14 left below. Hence, the optimal classification performance is attained for some intermediate beamwidth values (in this particular scenario, at $\Delta\theta/2 = 4/D$ for MLP and SVC, and $\Delta\theta/2 = 6/D$ for kNN).

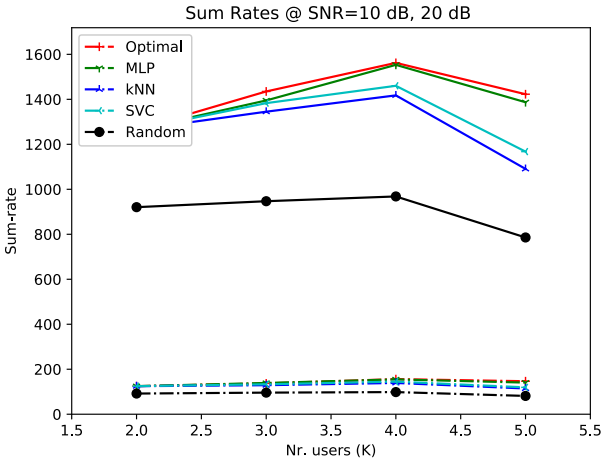

FIGURE 15. Sum-rate performance vs. number of users, for SNR = 10dB (dashed lines) and SNR = 20dB (solid lines) ($N_{\text{BS}} = 16$, $N_{\text{RF}} = 6$, $K = 4$, $B = 28$, $B_{\text{sub}} = 16$, $\Delta\theta/2 = 3/D$, $D = 8$, MUSIC estimator).

Complementarily, Figure 14 right illustrates the achievable performance in terms of sum-rate. Here, the impact of beamwidth is two-fold: not only in terms of separability of the decision regions, classification performance and, ultimately, their impact on sum-rate; but also on its interference mitigation capabilities. In other words, the larger the beamwidth, the higher the impact of interferers after the analog (ABF) and digital (BBF) beamforming blocks and, thus, the lower the sum-rate. Due to this, the optimal value of $\Delta\theta/2$ in terms of sum-rate differs from (and is smaller than) that observed for classification accuracy. For the ML and DL schemes, the degradation with respect to optimal beam selection (caused by classification errors) is, in general, small.

Finally, we focus our attention on the number of active users K (Fig. 15). For small K, the communications system is noise-limited but the sum-rate is necessarily low due to the limited number of contributors (users) to the aggregated rate. For an increasing number of users, the number of contributors increases but, to start with, the classification accuracy deteriorates. This stems from the fact that the number of inputs to the classification scheme grows linearly in the

number of users, this rendering the decision regions and the model to be fit more complex. Further, the system becomes interference-limited, this bringing in an additional source of performance degradation.

VII. CONCLUSIONS

In this paper, we have investigated how angle-of-arrival information can be exploited by deep-/machine-learning approaches to perform beam selection in a hybrid beamforming setup. Specifically, we have posed the beam selection task as a multi-class classification problem and solved it via two machine-learning (ML) approaches: k-nearest neighbors, support vector classifiers; and one deep learning approach: the multilayer perceptron). We have also designed a new beamformer set with larger and configurable beamwidth which is tailored to the beam selection task. Computer simulation results reveal that classification accuracy exhibits some degradation in the presence of received power fluctuations which is larger for ML approaches (actual CSI). With estimated CSI (Capon method), the degradation is more severe, both for ML and DL schemes. This is mostly due to the limited spatial resolution of the AoA estimation block operating on a filtered received signal, and the impact of additive noise on CSI estimates (yet the impact of the latter is moderate). In general, MLP outperforms both kNN and SVC methods in terms of classification accuracy. MUSIC allows to estimate both AoA and received powers more accurately, this resulting into an enhanced classification in particular for the MLP scheme (close to 80%). As a result, its sum-rate is virtually identical to that of optimal beam selection since most classification errors occur among adjacent regions. As for the ML schemes, they retain some 90% of the sum-rate with optimal beam selection. As for the size of the trainig dataset, we found out that classification performance saturates when the number of examples reaches 3000. Concerning the number of analog beamformers, we considered a scenario with a total of $B = 28$ initial elements in the codebook. We learnt that, if the selection is restricted to the subset of $B_{\text{subset}} = 16$ most popular beamfomers (or more), the sum-rate remains unchanged. On the contrary, for subsets of $B_{\text{subset}} \geq 8$ elements, performance degrades very rapidly (despite that classification accuracy improves, due to the reduced number of elements in the codebook). Besides, we realized that there exists an optimal value for the beamwidth of analog beamformers. This makes it possible to strike the optimal balance between too narrow beams (i.e., in phased arrays) exhibiting poor classification accuracy/sum-rate; and too broad/non-directional ones for which the angles or arrival are not a sufficient statistic to perform classification. Concerning the number of active users, a similar behaviour (convex, with some optimal value) was observed. For a small number of users, the (aggregated) sum-rate is necessarily low. For an increasing number of users, the classification accuracy deteriorates and the system becomes interference-limited, this driving sum-rate again towards lower values.

Future work in this area, includes the generalization to scenarios with multi-path propagation. In some cases, this might entail the use of alternative sufficient statistics (e.g., received signal impinging on the array or related transformation) as an input to machine and deep learning methods.

REFERENCES

[1] C. Jiang, H. Zhang, Y. Ren, Z. Han, K.-C. Chen, and L. Hanzo, "Machine learning paradigms for next-generation wireless networks," *IEEE Wireless Commun.*, vol. 24, no. 2, pp. 98–105, Apr. 2017.

[2] *IMT Vision—Framework and Overall Objectives of the Future Development of IMT for 2020 and Beyond*, document Rec. ITU-R M.2083-0, International Telecommunication Union, Sep. 2015.

[3] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Trans. Cogn. Commun. Netw.*, vol. 3, no. 4, pp. 563–575, Dec. 2017.

[4] H. Ye, G. Y. Li, and B.-H. Juang, "Power of deep learning for channel estimation and signal detection in OFDM systems," *IEEE Wireless Commun. Lett.*, vol. 7, no. 1, pp. 114–117, Feb. 2018.

[5] C.-K. Wen, W.-T. Shih, and S. Jin, "Deep learning for massive MIMO CSI feedback," *IEEE Wireless Commun. Lett.*, vol. 7, no. 5, pp. 748–751, Oct. 2018.

[6] J. Joung, "Machine learning-based antenna selection in wireless communications," *IEEE Commun. Lett.*, vol. 20, no. 11, pp. 2241–2244, Nov. 2016.

[7] Y. Long, Z. Chen, J. Fang, and C. Tellambura, "Data-driven-based analog beam selection for hybrid beamforming under mm-wave channels," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 2, pp. 340–352, May 2018.

[8] R. Pal, K. V. Srinivas, and A. K. Chaitanya, "A beam selection algorithm for millimeter-wave multi-user MIMO systems," *IEEE Commun. Lett.*, vol. 22, no. 4, pp. 852–855, Apr. 2018.

[9] Y. Niu, Z. Feng, Y. Li, Z. Zhong, and D. Wu, "Low complexity and near-optimal beam selection for millimeter wave MIMO systems," in *Proc. 3th Int. Wireless Commun. Mobile Comput. Conf. (IWCMC)*, Jun. 2017, pp. 634–639.

[10] S. Qiu, K. Luo, and T. Jiang, "Beam selection for mmWave massive MIMO systems under hybrid transceiver architecture," *IEEE Commun. Lett.*, vol. 22, no. 7, pp. 1498–1501, Jul. 2018.

[11] A. Klautau, P. Batista, N. González-Prelcic, Y. Wang, and R. W. Heath, Jr., "5G MIMO data for machine learning: Application to beam selection using deep learning," in *Proc. Inf. Theory Appl. Workshop (ITA)*, Feb. 2018, pp. 1–9.

[12] A. Smola and S. Vishwanathan, *Introduction to Machine Learning*. Cambridge, U.K.: Cambridge Univ. Press, 2008.

[13] D. He, C. Liu, T. Q. S. Quek, and H. Wang, "Transmit antenna selection in MIMO wiretap channels: A machine learning approach," *IEEE Wireless Commun. Lett.*, vol. 7, no. 4, pp. 634–637, Aug. 2018.

[14] A. Mukherjee and A. Hottinen, "Learning algorithms for energy-efficient MIMO antenna subset selection: Multi-armed bandit framework," in *Proc. 20th Eur. Signal Process. Conf. (EUSIPCO)*, Aug. 2012, pp. 659–663.

[15] H. L. Van Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*. Hoboken, NJ, USA: Wiley, 2002.

[16] C. Antón-Haro and X. Mestre, "On the applicability of machine and deep learning techniques for beam selection in hybrid beamforming architectures with partial CSI," in *Proc. IEEE Int. Conf. Commun. Workshop Mach. Learn. Wireless Commun. (MLCOM)*, Jan. 2019.

[17] F. Gholam, J. Via, and I. Santamaria, "Beamforming design for simplified analog antenna combining architectures," *IEEE Trans. Veh. Technol.*, vol. 60, no. 5, pp. 2373–2378, Jun. 2011.

[18] M. L. McCloud and L. L. Scharf, "A new subspace identification algorithm for high-resolution DOA estimation," *IEEE Trans. Antennas Propag.*, vol. 50, no. 10, pp. 1382–1390, Oct. 2002.

[19] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *J. Mach. Learn. Res.*, vol. 12, p. 2121–2159, Jul. 2011.

[20] *Scikit-Learn: Machine Learning in Python*. Accessed: Sep. 20, 2018. [Online]. Available: https://scikit-learn.org

[21] *Tensorflow: An Open Source Machine Learning Library for Research and Production*. Accessed: Sep. 20, 2018. https://www.tensorflow.org

[22] *Keras: The Python Deep Learning Library*. Accessed: Sep. 20, 2018. [Online]. Available: https://keras.io/

CARLES ANTÓN-HARO received the Ph.D. degree in telecommunications engineering from the Technical University of Catalonia. In 1999, he joined Ericsson, Spain, where he participated in rollout projects of 2G and 3G mobile networks. He is currently a Senior Researcher and the Director of R&D Programs and with the Centre Tecnòlogic de Telecomunicacions de Catalunya. As a Senior Researcher, his research interests include signal processing and estimation theory for communications, this including machine learning, the sensor and IoT networks, M2M communications, array signal processing, MIMO, energy harvesting, and smart grids. He has published more than 110 technical papers in IEEE journals and conferences (two Best Paper Awards) and books/book chapters. He has supervised five Ph.D. Theses (one in progress). He is a Senior Member of the IEEE. In recent years, he has been actively involved in the organization of major conferences such as the IEEE Wireless Communications and Networking Conference 2018 (General Chair), Workshop on Integrating Communications, Control, and Computing Technologies for Smart Grid ICC17 (TPC Chair), or European Signal Processing Conference 2011 (General Vice-Chair). He is an Associate Editor of *EURASIP Journal on Wireless Communications and Networks*.

XAVIER MESTRE received the M.S. and Ph.D. degrees in electrical engineering from the Technical University of Catalonia (UPC), in 1997 and 2002, respectively, and the Licenciate Degree in mathematics, in 2011. From 1998 to 2002, he was with Communications Signal Processing Group, UPC, where he was a Research Assistant. In 2003, he joined the Telecommunications Technological Center of Catalonia, where he currently holds a position as a Senior Research Associate in the area of radio communications. He is currently the Head of the Advanced Signal and Information Processing Department. During this time, he has actively participated in multiple European projects and several contracts with the local industry, some of which have led to commercialized products. He has authored three granted patents. He has authored nine book chapters, 41 international journal papers, and more than 90 articles in international conferences. He has participated in the organization of multiple conferences: IEEE Wireless Communications and Networking Conference 2018 (General Vice-Chair) and the IEEE International Conference on Acoustics, Speech and Signal Processing 2020 (General Chair). He is a Senior Member of the IEEE and an elected member of the IEEE Sensor Array and Multichannel Signal Processing Technical Committee. During the pursuit of his Ph.D., he was a recipient of the 1998-2001 Ph.D. Scholarship (granted by the Catalan Government) and received the 2002 Rosina Ribalta Second Prize for the Best Doctoral Thesis Project within the areas of Information Technologies and Communications by the Epson Iberica Foundation. He has been an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, (2007–2011), (since 2015).