A Supervised Learning Framework for Joint Estimation of Angles-of-Arrival and Number of Sources

N.B. When citing this work, cite the original published paper.

(article starts on next page)

RESEARCH ARTICLE

A Supervised Learning Framework for Joint Estimation of Angles-of-Arrival and Number of Sources

NOUD KANTERS [1], (Graduate Student Member, IEEE),
AND ANDRÉS ALAYÓN GLAZUNOV [1,2], (Senior Member, IEEE)
[1]Department of Electrical Engineering, University of Twente, 7500 AE Enschede, The Netherlands
[2]Department of Electrical Engineering, Chalmers University of Technology, 41296 Gothenburg, Sweden

Corresponding author: Noud Kanters (n.b.kanters@utwente.nl)

ABSTRACT Machine learning is a promising technique for angle-of-arrival (AOA) estimation of waves impinging a sensor array. However, the majority of the methods proposed so far only consider a known, fixed number of impinging waves, i.e., a fixed number of sources (NOS). This paper proposes a machine-learning-based estimator designed for the case when the NOS is variable and hence unknown a priori. The proposed estimator comprises a framework of single-label classifiers. Each classifier predicts if waves are present within certain randomly selected segments of the array's field of view (FOV), resulting from discretising the FOV with a certain (FOV) resolution. The classifiers' predictions are combined into a probabilistic angle spectrum, whereupon the NOS and the AOAs are estimated jointly by applying a probability threshold whose optimal level is learned from data. The estimator's performance is assessed using a new performance metric: the joint AOA estimation success rate. Numerical simulations show that for low SNR (-10 dB), a low FOV resolution ($2°$) yields a higher success rate than a high resolution ($1°$), whereas the opposite applies for mid (0 dB) and high (10 dB) SNRs. In nearly all simulations, except one at low SNR and a high FOV resolution, the proposed estimator outperforms the MUSIC algorithm if the maximum allowed AOA estimation error is approximately equal to (or larger than) the FOV resolution.

INDEX TERMS Angle-of-arrival estimation, number of sources detection, supervised learning, feedforward neural network.

I. INTRODUCTION

Angle-of-arrival (AOA) estimation of waves impinging a sensor array has been studied extensively as it has applications in various fields from array signal processing, e.g., wireless communications, radar and sonar [1]. In many practical applications, the number of waves impinging the array, henceforth referred to as the number of sources (NOS), is not constant, meaning it has to be estimated as well. Solutions to this problem can be categorised into separable and joint detection methods [2]. In separable detection, the NOS is estimated before estimating the AOAs, for example through model order estimators Akaike's information criterion (AIC) or the minimum description length (MDL) [3]. On the contrary, in joint detection the NOS and the AOAs are estimated simultaneously.

The associate editor coordinating the review of this manuscript and approving it for publication was Hasan S. Mir.

Conventional AOA estimators generally require an estimate of the NOS prior to the AOA estimation, and hence they correspond to the separable detection category. Beamforming algorithms, e.g., the Bartlett and the Capon beamformers [4], belong to this class of conventional estimators. Their resolution, i.e., their ability to resolve closely spaced sources, depends directly on the physical size of the array [1]. This limitation does not apply to subspace-based super-resolution algorithms like multiple signal classification (MUSIC) [5], estimation of signal parameters via rotational invariance techniques (ESPRIT) [6], and variants thereof like root-MUSIC [7]. However, these algorithms require the computationally expensive eigenvalue decomposition. Moreover, the resolution of the MUSIC algorithm deteriorates for highly correlated signals, whereas ESPRIT and root-MUSIC can only be applied in combination with particular array geometries [1]. Maximum likelihood (ML) methods, e.g., [8],

[9], do not suffer from these fundamental limitations, but their computational complexity grows exponentially with the NOS. In order to mitigate the aforementioned shortcomings, various sparsity-based approaches have been proposed, e.g., [10], [11], [12]. While these methods can handle scenarios of unknown NOSs (i.e., joint detection), spurious sources are often present in the resulting power spectra [13].

Recently, supervised-learning-based AOA estimation algorithms have been proposed to further improve the accuracy and/or the computational efficiency. These algorithms learn a mapping between array outputs and AOAs from data directly. Hence, they do not require specific assumptions regarding the array geometry or the data model. The majority of these supervised-learning-based works are only applicable if the NOS is fixed, henceforth referred to as scenario I. In other words, they can be considered part of the separable detection category, but they do not consider the NOS detection itself. For example in [14], the 2D AOA estimation (i.e., azimuth and elevation angle estimation) of a single source is performed by combining the conventional MUSIC algorithm with different learning algorithms, i.e., neural networks (NN), Gaussian processes (GP) and regression trees (RT). All of them consistently outperform the baseline MUSIC algorithm in terms of the average AOA estimation error, with improvements up to 50% for GP and RT in particular high-SNR, low-elevation situations. Similarly, [15] considers the 2D AOA estimation of a single source through an ensemble of five convolutional neural networks (CNNs), [16] investigates the 1D AOA estimation of two sources through a deep neural network (DNN) and [17] proposes to emulate a large array through a DNN, whereupon the 1D single-source AOA is estimated using the MUSIC algorithm. In [18], multi-source 1D AOA estimation is performed by using a separate support vector machine (SVM) for the estimation of each AOA. Although this implies that the NOS determines how many SVMs are required, the NOS detection itself is not considered.

In many practical applications the NOS is not constant (which we refer to as scenario II in this paper), hence it is to be estimated too. Therefore, the joint estimation of the NOS and the AOAs, i.e., the alternative to separable detection, is of great relevance. Clearly, an AOA estimator performing joint estimation comes with increased complexity, as it should be capable of estimating a variable number of parameters. Since this is not straightforward to implement using existing learning algorithms, it has received less attention. Nevertheless, a number of solutions have been proposed. For example, in [19], a single DNN is deployed for the estimation of both the NOS (restricted to be between 1 and 4 by design) and the 1D AOAs. More freedom in terms of the NOSs that can be handled is provided by the methods presented in [20] and [21]. There, the estimators (comprising multiple parallel DNNs in [20] and a single CNN in [21]) are tailored to a 1D grid of search angles (1° resolution) within the FOV of the sensor array. Hence, they formulate the AOA estimation

problem as a classification problem and aim to find those search angles which represent AOAs. In [20], the predictions for all search angles are combined into an angle spectrum, whereupon the arguments of the highest peaks are returned as the AOA estimates. However, it is not explained how the estimator deals with scenarios of unknown NOSs. On the contrary, [21] uses a user-defined (i.e., not optimized) confidence level to determine the NOS. Neither [20] nor [21] investigates how the grid resolution itself affects the predictions of the learning algorithm.

In this paper, we adopt an approach comparable to the ones presented in [20] and [21]: we discretise the array's FOV, whereupon the joint AOA estimation problem is solved through classification. The main contributions of this paper can be summarized as follows:

- A machine learning framework (MLF) is proposed to jointly estimate the NOS and the AOAs of waves impinging a sensor array of arbitrary geometry. The MLF consists of an ensemble of classifiers, trained through supervised learning, which are organized along a framework based on the ensemble method random k-labelsets (RAkEL) [22]. Consequently, the proposed MLF can, in principle, be deployed in combination with any learning algorithm capable of single-label multi-class classification. Modifications to the RAkEL method are implemented to tailor it to the AOA estimation problem.

- A peak detection algorithm is devised in order to jointly extract the NOS and the AOAs from the probabilistic angle spectrum. This algorithm comprises a probability threshold, whose level is optimized based on data. The spectrum peaks above the threshold are located, whereupon the number of peaks and their arguments are returned as the NOS estimate and the AOA estimates, respectively.

- The impact of the resolution of the FOV discretisation (FOV resolution) on the predictions of the individual classifiers as well as on the final AOA estimates is investigated through numerical simulations, using feedforward NNs as the learning algorithm. It is shown that increasing the FOV resolution does not necessarily improve the overall joint AOA estimation success rate (see next point) of the MLF, depending on the signal-to-noise ratio (SNR).

- A new performance metric, the joint AOA estimation success rate, is introduced. This metric is based on the *success rate* proposed in [15], but here we adapt it to take into account both the NOS and the AOAs, and to make it depend on a user-defined maximum allowed AOA estimation error. Its expected value in case of ideal classifiers is derived for the case of uniformly distributed random AOAs.

- The MLF is compared to the conventional MUSIC algorithm [5] combined with the MDL and the AIC NOS estimators [3]. Numerical simulations representing a variety of SNRs (−10, 0, 10 dB) and FOV resolutions (2° and 1°) in both scenarios I and II show that the proposed MLF achieves a higher rate of successful joint AOA estimation than the MUSIC algorithm if the maximum allowed AOA estimation error is of the order of (or larger than) the FOV resolution.

This applies to nearly all considered cases, except one at low SNR (-10 dB) and high FOV resolution ($1°$) in scenario I.

The following notations apply throughout the paper. The transpose operator is denoted by $(\cdot)^T$, $(\cdot)^H$ stands for complex conjugate transpose and $E[\cdot]$ is the expectation operator. Scalars are denoted as a or A (lightface), whereas **a** (boldface lowercase) denotes a column vector and **A** (boldface uppercase) is a matrix. $\Re(\cdot)$ and $\Im(\cdot)$ represent the real and imaginary part of a complex variable or function, respectively. The $n \times n$ identity matrix is denoted as $\mathbf{I}_n$ and diag(**a**) is a diagonal matrix with the elements of **a** on the diagonal.

The remainder of this paper is structured as follows. The data model and the problem statement are discussed in Section II. The proposed AOA estimator is presented in Section III, whereupon performance metrics are described in Section IV. The conducted simulations are described in Section V and their results are presented and analyzed in Section VI. Conclusions follow in Section VII.

II. DATA MODEL AND PROBLEM STATEMENT

Let's consider Q narrowband sources (i.e., incident plane waves) in the far-field of a uniform linear array (ULA) composed of N sensors with inter-element spacing d. It is assumed that the sources and the sensors are all in the same plane, such that the direction-of-arrival (DOA) of each incident plane wave can be described by a single parameter, i.e., an angle-of-arrival (AOA). Hence, a one-dimensional (1D) AOA estimation problem is considered. The AOA of the q^{th} wave equals θ_q, with $q = 1, \ldots, Q$, and is defined with respect to the ULA's broadside. The problem addressed in this paper is the joint estimation of the number of sources (NOS) Q and the AOAs $\theta_1, \ldots, \theta_Q$ given T snapshots of the sensor array output.

The sensor array output $\mathbf{y}(t) \in \mathbb{C}^{N \times 1}$, sampled at time instance t, is represented by the signal model

$$\mathbf{y}(t) = \mathbf{A}(\theta_1, \ldots, \theta_Q)\mathbf{s}(t) + \mathbf{n}(t), \tag{1}$$

where $\mathbf{s}(t) \in \mathbb{C}^{Q \times 1}$ and $\mathbf{n}(t) \in \mathbb{C}^{N \times 1}$ represent the signal waveforms and the element noise, respectively, and $\mathbf{A}(\theta_1, \ldots, \theta_Q) \in \mathbb{C}^{N \times Q}$ is the array manifold consisting of Q steering vectors, i.e.

$$\mathbf{A}(\theta_1, \ldots, \theta_Q) = [\mathbf{a}_1(\theta_1), \ldots, \mathbf{a}_Q(\theta_Q)]. \tag{2}$$

The q^{th} steering vector $\mathbf{a}_q(\theta_q) \in \mathbb{C}^{N \times 1}$ describes the array response to the q^{th} wave and is defined as

$$\mathbf{a}_q(\theta_q) = \left[1, e^{j\frac{2\pi}{\lambda}d \sin\theta_q}, \ldots, e^{j\frac{2\pi}{\lambda}(N-1)d \sin\theta_q}\right]^T, \tag{3}$$

where λ is the wavelength of the transmitted signal.

In this paper, $\mathbf{s}(t)$ and $\mathbf{n}(t)$ are both assumed to be i.i.d. zero-mean complex Gaussian random variables. Hence, the signal covariance matrix is given by

$$\mathbf{P} = E[\mathbf{s}(t)\mathbf{s}^H(t)] = \text{diag}([\sigma_1^2, \ldots, \sigma_Q^2]^T), \tag{4}$$

where σ_q^2 denotes the variance of the q^{th} signal. It is assumed that the noise power is equal over all sensors, such that the noise covariance matrix is defined as

$$\mathbf{Q} = E[\mathbf{n}(t)\mathbf{n}^H(t)] = \nu^2 \mathbf{I}_N, \tag{5}$$

where ν^2 is the noise variance. Hence, the covariance matrix equals

$$\mathbf{R} = E[\mathbf{y}(t)\mathbf{y}^H(t)] = \mathbf{A}\mathbf{P}\mathbf{A}^H + \mathbf{Q}. \tag{6}$$

In practice, $\mathbf{R}$ has to be estimated from noisy array measurements. For an array measurement consisting of T snapshots, the maximum likelihood estimate, $\hat{\mathbf{R}}$, is computed as

$$\hat{\mathbf{R}} = \frac{1}{T} \sum_{t=1}^{T} \mathbf{y}(t)\mathbf{y}^H(t), \tag{7}$$

where it is assumed that the AOAs $\theta_1, \ldots, \theta_Q$ (and therefore Q as well) are identical for all T snapshots $\{\mathbf{y}(1), \ldots, \mathbf{y}(T)\}$.

The machine learning framework developed for the joint AOA estimation problem is presented next.

III. SUPERVISED-LEARNING-BASED JOINT AOA ESTIMATION FRAMEWORK

The proposed learning-based estimator comprises two main components: (I) an ensemble of learning-based classifiers organized along a framework, and (II) a procedure to convert the predictions of these classifiers to angle-of-arrival (AOA) estimates. We proceed by first presenting each component and the related aspects, followed by a description of the deployment procedure of the estimator as a whole.

A. AOA ESTIMATION FRAMEWORK

AOA estimation in scenarios with a variable number of sources (NOS) implies that the number of parameters to be estimated is variable too. As this is not straightforward to implement using existing supervised learning algorithms, a framework is devised to recast the problem. This framework is the core of the estimator as it defines the number of classifiers in the ensemble, what their target outputs should be during training, and how their predictions should be interpreted and converted into AOA estimates during deployment.

1) MULTI-SOURCE AOA ESTIMATION THROUGH CLASSIFICATION

Consider the array's field of view (FOV) defined by the interval $[\theta_{\min}, \theta_{\max})$. This interval is discretised into M non-overlapping segments. Although not necessary, the presented method is specialized to a regular discretisation. Therefore, each segment spans $\Delta\theta$ degrees, where

$$\Delta\theta = \frac{\theta_{\max} - \theta_{\min}}{M}. \tag{8}$$

Hence, $\Delta\theta$ denotes the angle resolution of the FOV discretisation, henceforth abbreviated as the FOV resolution. The i^{th} FOV segment is defined by the interval $[\theta_{i,\min}, \theta_{i,\max})$, where

$$\theta_{i,\min} = \theta_{\min} + (i-1)\Delta\theta \tag{9a}$$

$$\theta_{i,\max} = \theta_{\min} + i\Delta\theta \tag{9b}$$

and $i = 1, \ldots, M$. Using the discretised FOV, we recast the AOA estimation problem as a classification problem: the proposed estimator aims to find those, and only those, FOV segments which include at least one of the AOAs $\theta_1, \ldots, \theta_Q$. This is a so-called multi-label multi-class (or simply multi-label) classification problem [23]: M distinct labels (here, non-overlapping FOV segments) exist, of which at most[1] Q should be assigned to a single instance (here, a collection of T snapshots of the array output).

Multi-label classification problems have been addressed successfully by transforming them into multiple single-label classification problems through the random k-labelsets (RAkEL) method [22]. This method is the basis for the AOA estimation framework, hence we present its main principles below.

2) RAkEL FOR MULTI-LABEL CLASSIFICATION [22]

RAkEL transforms a multi-label problem of M labels, $\{\lambda_1, \ldots, \lambda_M\}$, into m single-label problems of 2^k labels (where $k < M$) such that it can be solved by m single-label classifiers $h_1, \ldots, h_m$. This is achieved in two steps. First, the multi-label problem is divided in m smaller (but still multi-label) problems by generating m subsets of k labels (called k-labelsets). The second step is the transformation of the smaller multi-label problems into single-label problems via a method called label powerset (LP). The LP of k-labelset R_j ($j = 1, \ldots, M$), denoted as $\mathcal{P}(R_j)$, is the set containing all 2^k possible subsets of R_j as its elements. For example, if $k = 2$ and $R_j = \{\lambda_a, \lambda_b\}$, then $\mathcal{P}(R_j) = \{\{\}, \{\lambda_a\}, \{\lambda_b\}, \{\lambda_a, \lambda_b\}\}$. Hence, by defining 2^k new labels, each of them representing a different element of $\mathcal{P}(R_j)$, the j^{th} multi-label problem can be solved indirectly by single-label classifier h_j.

The k-labelsets can be generated via random sampling either with or without replacement, referred to as RAkEL$_\text{o}$ and RAkEL$_\text{d}$, respectively. Here, the subscript 'o' stands for overlapping and the 'd' for disjoint. Contrary to RAkEL$_\text{d}$, a label could be included in multiple k-labelsets in RAkEL$_\text{o}$ in which case the final prediction on whether to assign this label is obtained by a majority voting procedure. Averaged over 8 datasets from different fields, RAkEL$_\text{o}$ outperforms RAkEL$_\text{d}$ in terms of the F_1-score, a measure for predictive performance [22]. The authors of [22] recommend to use RAkEL$_\text{o}$ with a small k (e.g., $k = 3$) and $M < m < 2M$, as it is more efficient to use a large m than a large k in terms of the computational burden.

3) COMBINING RAkEL$_\text{d}$ AND RAkEL$_\text{o}$ FOR AOA ESTIMATION

In this paper, RAkEL is applied for the sake of joint AOA and NOS estimation. Hence, the labels $\lambda_1, \ldots, \lambda_M$ represent the FOV segments, where the i^{th} segment is defined by the interval $[\theta_{i,\min}, \theta_{i,\max})$ (9). However, rather than using either RAkEL$_\text{o}$ or RAkEL$_\text{d}$, we propose to combine both variants, because of the following: when generating the

TABLE 1. Example RAkEL-based [22] AOA estimation framework with $k = 2$, $M = 4$, $L = 2$.

classif.	k-labelset	label-subset	predict.	$\bar{P}_{i,j,\tilde{k}}$ (11)			
h_j	R_j	$\tilde{R}_{j,\tilde{k}}$	$\tilde{P}_{j,\tilde{k}}$	λ_1	λ_2	λ_3	λ_4
Layer 1							
h_1	$\{\lambda_1, \lambda_3\}$	$\tilde{R}_{1,1} = \{\}$	$\tilde{P}_{1,1}$	0	0	0	0
		$\tilde{R}_{1,2} = \{\lambda_1\}$	$\tilde{P}_{1,2}$	$\tilde{P}_{1,2}$	⋮	0	⋮
		$\tilde{R}_{1,3} = \{\lambda_3\}$	$\tilde{P}_{1,3}$	0		$\tilde{P}_{1,3}$	
		$\tilde{R}_{1,4} = \{\lambda_1, \lambda_3\}$	$\tilde{P}_{1,4}$	$\tilde{P}_{1,4}$	0	$\tilde{P}_{1,4}$	0
h_2	$\{\lambda_2, \lambda_4\}$	$\tilde{R}_{2,1} = \{\}$	$\tilde{P}_{2,1}$	0	0	0	0
		$\tilde{R}_{2,2} = \{\lambda_2\}$	$\tilde{P}_{2,2}$	⋮	$\tilde{P}_{2,2}$	⋮	0
		$\tilde{R}_{2,3} = \{\lambda_4\}$	$\tilde{P}_{2,3}$		0		$\tilde{P}_{2,3}$
		$\tilde{R}_{2,4} = \{\lambda_2, \lambda_4\}$	$\tilde{P}_{2,4}$	0	$\tilde{P}_{2,4}$	0	$\tilde{P}_{2,4}$
Layer 2							
h_3	$\{\lambda_1, \lambda_4\}$	⋮	⋮	⋮	0	0	⋮
h_4	$\{\lambda_2, \lambda_3\}$	⋮	⋮	0	⋮	⋮	0

k-labelsets via random sampling with replacement (RAkEL$_\text{o}$), one cannot control the number of k-labelsets in which a particular label is included. More specifically, as each label is selected with equal probability, there is a probability of $((M-k)/M)^m$ for a label not to be included in any k-labelset. For the application addressed in this work, this implies that certain segments of the FOV might not be considered by the AOA estimator. This is clearly problematic as the estimator would not be able to 'see' waves with AOAs within those segments. To circumvent this problem without having to increase m and/or k (which increases the computational burden), it is proposed to approximate RAkEL$_\text{o}$ by using L independent 'layers' of RAkEL$_\text{d}$. Consequently, each label is included in exactly L k-labelsets and the majority voting procedure of RAkEL$_\text{o}$ can be applied for all labels $\lambda_1, \ldots, \lambda_M$. The total number of classifiers in this layered framework equals

$$m = L \lceil M/k \rceil, \tag{10}$$

where $\lceil \cdot \rceil$ rounds up the argument to the nearest integer.[2] It is worthwhile to note that both increasing the FOV resolution (i.e., decreasing $\Delta\theta$) and increasing the number of layers L results in a larger number of classifiers in the framework. An example of the proposed layered framework is presented in the first 2 columns of Table 1.

B. CONVERTING CLASSIFIER PREDICTIONS TO JOINT AOA ESTIMATES

Section III-A described how the AOA estimation problem is decomposed into multiple single-label classification problems. Here, we present how the classifiers' predictions are converted to NOS and AOA estimates when the estimator is deployed. To be as generic as possible regarding the single-label learning algorithm, it is assumed that the its

---

[1] The number of labels to be assigned is smaller than Q if multiple AOAs belong to the same FOV segment.

[2] By proper choice of M, the existence of a labelset consisting of less than k labels can be prevented and rounding can be discarded.

prediction comprises a set of probabilities, rather than a single index.

1) CLASSIFIER PREDICTIONS

Let's denote the elements of $\mathcal{P}(R_j)$ (i.e., the label subsets of R_j) as $\tilde{R}_{j,1}, \ldots, \tilde{R}_{j,2^k}$. Hence, by definition it holds that $\tilde{R}_{j,\tilde{k}} \subseteq R_j \subset \{\lambda_1, \ldots, \lambda_M\}$, where $\tilde{k} = 1, \ldots, 2^k$. Furthermore, we denote any prediction of classifier h_j as the set $\{\tilde{P}_{j,1}, \ldots, \tilde{P}_{j,2^k}\}$, where $\tilde{P}_{j,\tilde{k}}$ represents the probability that there is at least one AOA within every FOV segment represented by the labels in $\tilde{R}_{j,\tilde{k}}$, according to classifier h_j. Hence, $\tilde{P}_{j,\tilde{k}}$ is directly related to the label subset $\tilde{R}_{j,\tilde{k}}$, as visualised by columns 3 and 4 of Table 1. We assume that $0 \leq \tilde{P}_{j,\tilde{k}} \leq 1$ and that $\sum_{\tilde{k}=1}^{2^k} \tilde{P}_{j,\tilde{k}} = 1$.

Rather than converting the probabilistic predictions $\tilde{P}_{j,1}, \ldots, \tilde{P}_{j,2^k}$ to Boolean variables (i.e., 1 for the highest probability and 0 for all the others) and subsequently applying the majority voting procedure of RAkEL$_o$, we adopt another approach to estimate the AOAs. This approach is presented next.

2) COMPUTING NOS AND AOA ESTIMATES

To prevent the loss of information in this stage of the estimation process, all probabilistic predictions $\tilde{P}_{j,\tilde{k}}$ (with $j = 1, \ldots, m$ and $\tilde{k} = 1, \ldots, 2^k$) are converted to per-label-predictions $\bar{P}_{i,j,\tilde{k}}$ ($i = 1, \ldots, M$) according to

$$\bar{P}_{i,j,\tilde{k}} = \begin{cases} \tilde{P}_{j,\tilde{k}}, & \lambda_i \in \tilde{R}_{j,\tilde{k}}, \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

An example is presented in the 4 rightmost columns of Table 1. Then, the per-label-predictions $\bar{P}_{i,j,\tilde{k}}$ are combined into segment probabilities $P_1, \ldots, P_M$ as

$$P_i = \frac{1}{L} \sum_{j=1}^{m} \sum_{\tilde{k}=1}^{2^k} \bar{P}_{i,j,\tilde{k}}. \quad (12)$$

The division by L in (12) guarantees that $0 \leq P_i \leq 1$, as each label is included in exactly L k-labelsets. Hence, P_i represents the probability that there is at least one AOA within the i^{th} FOV segment, according to the L classifiers evaluating it.

Finally, we interpret the sequence of probabilities $P_1, \ldots, P_M$ as an angle spectrum, similar to the work presented in [20]. In order to jointly extract the NOS and the AOAs from this spectrum, we propose to use a straightforward peak detection algorithm. This algorithm locates all spectrum peaks above a threshold and returns the number of peaks as the NOS estimate, $\hat{Q}$, and their arguments as the AOA estimates $\hat{\theta}, \ldots, \hat{\theta}_{\hat{Q}}$. This approach prevents the existence of spurious AOA estimates at labels neighboring a desired label.[3] However, if two or more AOAs correspond to neighboring FOV segments, they cannot be distinguished.

---

[3]Also, the proposed method accommodates the use of different FOV discretisations (e.g., random non-uniform discretisations) for the different framework layers. A first step in this direction is presented in [24].

FIGURE 1. Example spectrum and resulting AOA estimates.

Since each peak has a plateau width of $\Delta\theta$ (8) degrees (the resolution of the discretised FOV), the centre of the plateau is taken as the estimate. The possible AOA estimates are therefore defined by the centres of the FOV segments $c_1, \ldots, c_M$, where

$$c_i = \frac{1}{2}(\theta_{i,\min} + \theta_{i,\max}) = \theta_{\min} + (i - 1/2)\Delta\theta. \quad (13)$$

An example spectrum and its corresponding AOA estimates are presented in Fig. 1.

C. DEPLOYMENT PROCESS FLOW

The complete AOA estimation procedure is visualised in Fig. 2. Three stages can be identified: (I) the preparation stage, (II) the training stage and (III) testing/estimation stage. Details regarding each of these stages are explained next.

1) PREPARATION STAGE

The core of the preparation stage is the construction of the RAkEL-based framework, i.e., the generation of the k-labelsets $R_1, \ldots R_m$, as described in Section III-A. For this, the array's FOV and the framework's topology need to be defined through the parameters $\theta_{\min}$, $\theta_{\min}$, and L, k, M, respectively, whereupon the FOV resolution $\Delta\theta$ (8) and the number of classifiers m (10) follow automatically.

Besides the framework construction, a number of settings regarding the classifier training (e.g., the learning algorithm and its corresponding design parameters) and the threshold optimization need to be defined during the preparation stage as well. Details are clarified below.

2) TRAINING STAGE

In the training stage, the AOA estimator is optimized based on training data. We assume a training set of D_{trn} instances is available, where an 'instance' is a collection of T snapshots of the noisy array output (1), paired with the corresponding AOAs (i.e., the AOAs for which these array outputs were computed). As the training stage is composed of two branches, (I) the classifier training branch and (II) the threshold optimization branch, the training set must be split in two (not necessarily equally large) parts.

The details of the classifier training branch depend on the employed learning algorithm. However, in general, the

FIGURE 2. Flowchart of AOA estimator deployment procedure.

procedure contains the following steps. First, the training data need to be prepared such that they can be used for supervised learning, meaning input-output pairs need to be composed. The input component of an input-output pair, the so-called feature vector, contains the available information based on which the learning algorithm computes its prediction. Hence, in the present work, the feature vector is derived from the array data. It is worthwhile to note that every instance from the training set (the part used for classifier training) yields m input-output pairs, i.e., one for each classifier, all sharing the same feature vector. After computing the feature vectors for all instances, element-wise feature normalization is applied, since some learning algorithms are sensitive to scale [25]. The output components of input-output pairs represent the prediction targets. Contrary to the inputs, they need to be computed for classifier (and each instance, clearly) individually, as each classifier is associated with its own k-labelset. Since it is assumed that each prediction of a single-label classifier comprises 2^k probabilities (Section III-B1), this must also apply to the prediction targets. Hence, for one particular training instance, the targets for classifier h_j ($j = 1, \ldots, m$), denoted as $\{\tilde{P}^{(t)}_{j,1}, \ldots, \tilde{P}^{(t)}_{j,2^k}\}$, are computed as

$$\tilde{P}^{(t)}_{j,\tilde{k}} = \begin{cases} 1, & \text{if } \tilde{R}_{j,\tilde{k}} = (R_j \cap \bar{\Lambda}) \\ 0, & \text{otherwise,} \end{cases} \quad (14)$$

where $\tilde{k} = 1, \ldots, 2^k$ and R_j is the k-labelset associated with h_j, with label-subsets $\tilde{R}_{j,1}, \ldots, \tilde{R}_{j,2^k}$, and where $\bar{\Lambda}$ is the set containing exactly those labels representing FOV segments which include at least one of the instance's AOAs. Hence, $\bar{\Lambda}$ is defined as

$$\bar{\Lambda} = \{\lambda_i \mid i \in \{1, \ldots, M\}$$
$$\wedge (\exists \theta_q)[\theta_q \in \Theta \wedge \theta_{i,\min} \le \theta_q < \theta_{i,\max}]\}, \quad (15)$$

where $\Theta = \{\theta_1, \ldots, \theta_Q\}$ with $\theta_1, \ldots, \theta_Q$ and Q being the true AOAs and the true NOS of the instance under consideration, respectively. After composing the input-output pairs for all instances and all classifiers, the actual training is carried out. As each classifier learns its own mapping, it is proposed to track the learning progress of each classifier individually by means of a validation set in order to determine when to stop training.

Once the training has been terminated for all classifiers, the threshold optimization branch is initiated. This branch aims to optimize the threshold level (i.e., probability level) employed in the peak detection algorithm (Fig. 1). The process is as follows. First, feature vectors are computed for all threshold optimization training instances. This is done in the same way as in the classifier training branch, except that the feature-wise normalization is done using the normalization constants (feature-wise means and variances) derived from the classifier training data. In this way, we emulate the

estimation stage, in which one can only normalize based on training data as well. The feature vectors are fed through the ensemble of trained classifiers, whereupon the resulting predictions $\tilde{P}_{j,\tilde{k}}$ are converted to angle spectra according to the procedure described in Section III-B. As these spectra (of which there are as many as there are threshold optimization training instances) only contain values between 0 and 1 by definition, the optimal threshold must be between these values as well. The actual threshold optimization is a matter of computing the AOA and NOS estimates for all spectra for a user-defined set of threshold values (to be defined in the preparation stage). The threshold level that maximizes the number of spectra for which the estimated NOS $\hat{Q}$ equals the true NOS Q is considered optimal and is used within the estimation stage.

3) TESTING/ESTIMATION STAGE
After finishing the training stage, the estimator can be applied for AOA estimation. For each instance, the estimation procedure is identical to the one described by the threshold optimization branch, except that the optimal threshold level is now known and can be applied directly. To assess performance of the estimator, a test set of D_{tst} instances is used.

IV. PERFORMANCE METRICS
In the present work, the accuracy of the estimates obtained from the proposed joint angle-of-arrival (AOA) estimator depends on (I) the framework topology (defined by framework parameters $\theta_{\min}$, $\theta_{\max}$ and L, k, M), and (II) the predictive performance of the single-label classifiers used within the framework. The metrics employed to study the impact of the above on the AOA estimates are defined below.

A. $P(\hat{Q}=Q)$ AND RMSE
Assuming joint AOA estimation is performed for P instances, the probability that the number of sources (NOS) is estimated correctly is defined as

$$P(\hat{Q}=Q) = \frac{P'}{P} \times 100\%, \qquad (16)$$

where

$$P' = \text{num}(\hat{Q}_p = Q_p), \qquad (17)$$

where Q_p and $\hat{Q}_p$ are the true and estimated NOS for instance p ($p = 1, \ldots, P$), respectively, and $\text{num}(x)$ denotes the operation of counting the number of statements for which x is true.

For the $P' \leq P$ instances for which the NOS is estimated correctly, the AOA estimates are evaluated by means of the root-mean-square error (RMSE), which is computed as

$$\text{RMSE} = \sqrt{\frac{1}{P'}\sum_{p'=1}^{P'}\left[\frac{1}{Q_{p'}}\sum_{q=1}^{Q_{p'}}(\theta_{p',q} - \hat{\theta}_{p',q})^2\right]}, \qquad (18)$$

where $\theta_{p',q}$ and $\hat{\theta}_{p',q}$ are the q^{th} true AOA and the q^{th} AOA estimate in instance p' ($p' = 1, \ldots, P'$), respectively, and $Q_{p'}$ represents the instance's NOS. For each instance, the AOAs and AOA estimates are sorted in the same order before computing the RMSE.

B. JOINT AOA ESTIMATION SUCCESS RATE
Given that the proposed AOA estimator jointly performs NOS detection and AOA estimation, a metric is devised which takes into account both these aspects. It is based on the *success rate* proposed in [15] and expressed as

$$f_{\text{sr}}(\tilde{\theta}) = \frac{\text{num}\left(\left[\hat{Q}_p = Q_p\right] \cap \left[|\hat{\theta}_{p,q} - \theta_{p,q}| \leq \tilde{\theta}\right]\right)}{P} \times 100\%, \qquad (19)$$

where $q = 1, \ldots, Q_p$ and $p = 1, \ldots, P$, with Q_p and P as defined above. Hence, (19) implies that the joint AOA estimate for the p^{th} instance is successful only if the NOS is estimated correctly, i.e., $\hat{Q}_p = Q_p$, and all AOA estimation errors $|\hat{\theta}_{p,q} - \theta_{p,q}|$ (computed after sorting) are smaller than or equal to the maximum allowed AOA estimation error $\tilde{\theta}$. It is worthwhile to note that estimation errors up to $\Delta\theta/2$ are expected due to the finite FOV resolution.

As a reference for the success rate $f_{\text{sr}}(\tilde{\theta})$, we introduce $f_{\text{sr,exp}}(\tilde{\theta})$, which represents the success rate that would be expected if all classifiers in the framework were ideal, i.e., if their predictions $\tilde{P}_{j,\tilde{k}}$ equal the prediction targets $\tilde{P}_{j,\tilde{k}}^{(t)}$ (14) for all considered instances. Hence, $f_{\text{sr,exp}}(\tilde{\theta})$ is a measure for success rate limitations imposed by the framework's topology. In the case of a regular FOV discretisation and uniformly distributed random AOAs sharing the interval $[\theta_{\min}, \theta_{\max})$, $f_{\text{sr,exp}}(\tilde{\theta})$ is computed as

$$f_{\text{sr,exp}}(\tilde{\theta}) = \begin{cases} \left(\frac{\tilde{\theta}}{\Delta\theta/2}\right)^Q f_{\text{sr,exp,max}}, & \text{if } \tilde{\theta} < \Delta\theta/2, \\ f_{\text{sr,exp,max}}, & \text{otherwise}, \end{cases} \qquad (20)$$

where

$$f_{\text{sr,exp,max}} = \binom{M+1-Q}{Q}\frac{Q!}{M^Q} \times 100\%. \qquad (21)$$

Derivations of (20) and (21) are presented in Appendix A and B, respectively. It is worthwhile to note that (20) assumes that the NOS Q is equal in all evaluated instances. If not, $f_{\text{sr,exp}}(\tilde{\theta})$ is computed for all possible values of Q individually and a (weighted) average is applied afterwards.

C. F_1-SCORE
Besides evaluating the NOS and AOA estimates directly, the predictions of the single-label classifiers are evaluated as well. This is done by means of the F_1-score (see, e.g., [26]). As the F_1-score is computed per label and per classifier, the notation $F_1(j, \tilde{k})$ is used from here, where the index $j = 1, \ldots, m$ refers to the classifier and the index $\tilde{k} = 1, \ldots, 2^k$ to the label. The F_1-score is defined as the harmonic mean of two other metrics, precision and recall, with the subscript

1 indicating that precision and recall both contribute with equal weights to the mean, i.e.,

$$F_1(j, \tilde{k}) = 2 \times \frac{\text{precision}(j, \tilde{k}) \times \text{recall}(j, \tilde{k})}{\text{precision}(j, \tilde{k}) + \text{recall}(j, \tilde{k})}. \qquad (22)$$

Here, precision$(j, \tilde{k})$ is defined as the ratio

$$\text{precision}(j, \tilde{k}) = \frac{\text{tp}_{j,\tilde{k}}}{\text{tp}_{j,\tilde{k}} + \text{fp}_{j,\tilde{k}}}, \qquad (23)$$

where $\text{tp}_{j,\tilde{k}}$ and $\text{fp}_{j,\tilde{k}}$ denote the number of true and false positives for label $\tilde{k}$ and classifier h_j, respectively.[4] Hence, precision is a measure for a classifier's exactness. Furthermore, recall$(j, \tilde{k})$ is defined as the ratio

$$\text{recall}(j, \tilde{k}) = \frac{\text{tp}_{j,\tilde{k}}}{\text{tp}_{j,\tilde{k}} + \text{fn}_{j,\tilde{k}}}, \qquad (24)$$

where $\text{fn}_{j,\tilde{k}}$ denotes the number of false negatives for label $\tilde{k}$ and classifier h_j. Hence, recall represents the fraction of all instances of label $\tilde{k}$ that are actually classified as such and is therefore a measure for a classifier's completeness. Consequently, it holds that $0 \le F_1(j, \tilde{k}) \le 1$, with a higher value indicating a higher predictive performance.

In this work, the assessment of all classifiers yields $m \times 2^k$ F_1-scores. We compute averages $\bar{F}_1(Q_h)$, $Q_h = 1, \ldots, k$, by averaging all F_1-scores (22) which correspond to labels representing subsets $\tilde{R}_{j,\tilde{k}}$ with the same subset cardinality $|\tilde{R}_{j,\tilde{k}}|$. Hence, $\bar{F}_1(Q_h)$ is defined as[5]

$$\bar{F}_1(Q_h) = \frac{1}{m} \sum_{j=1}^{m} \frac{1}{|S_{j,Q_h}|} \sum_{\kappa \in S_{j,Q_h}} F_1(j, \kappa), \qquad (25)$$

where S_{j,Q_h} is the set containing those indices $\tilde{k}$ that refer to the elements of $\mathcal{P}(R_j)$ (the label powerset of the k-labelset of classifier h_j) whose cardinality equals Q_h, i.e.,

$$S_{j,Q_h} = \{\tilde{k} \mid \tilde{k} \in \{1, \ldots 2^k\} \wedge |\tilde{R}_{j,\tilde{k}}| = Q_h\}. \qquad (26)$$

For example, if $R_j = \{\lambda_a, \lambda_b\}$ ($k = 2$) and we denote its subsets $\{\}, \{\lambda_a\}, \{\lambda_b\}, \{\lambda_a, \lambda_b\}$ as $\tilde{R}_{j,1}, \ldots, \tilde{R}_{j,4}$, respectively, then $S_{j,0} = \{1\}$, $S_{j,1} = \{2, 3\}$ and $S_{j,2} = \{4\}$.

V. NUMERICAL SIMULATIONS

In this section, we describe the simulations that were conducted to assess the performance of the proposed angle-of-arrival (AOA) estimator. A summary of the simulation parameters is presented in Table 2. Details are given below.

---

[4]Since the classifiers' predictions are assumed to be probabilities rather than boolean variables, true/false positives/negatives are ill-defined. For the sake of F_1-score computation, we therefore assign boolean 1 to the label corresponding to the highest probability and boolean 0 to all the others.

[5]Note that if $\text{tp}_{j,\tilde{k}} = 0$, then $F_1(j, \tilde{k})$ is not defined. In this case, this particular $F_1(j, \tilde{k})$ is excluded from (25) and the average is taken over all remaining valid F_1-scores.

TABLE 2. Simulation parameters.

Parameter	Value	
	Scenario I	Scenario II
Sources and Signals		
Number of sources	$Q = 2$	$Q \sim U(1, 4)$
AOAs	$\theta_1, \ldots, \theta_Q \sim U(-60°, 60°)$	
SNR	$\text{SNR} \in \{-10, 0, 10\}$ dB	
AOA Estimation Framework		
FOV	$[\theta_{\min}, \theta_{\max}) = [-60°, 60°)$	
# framework layers	$L \in \{1, 3, 5\}$	
FOV resolution	$\Delta\theta \in \{2°, 1°\}$	
Labelsets	$k = 3$	
Sensor array		
Configuration	ULA	
# sensors	$N = 8$	
Inter-element spacing	$d = \lambda/2$	
Single-Label Classifiers		
Learning algorithm	Feedforward neural networks	
Input layer, # neurons	$N^2 = 64$	
Hidden layers, #	2	5
Hidden layers, # neurons	64, 36	100, 100, 100, 100, 50
Hidden layers, activ. funct.	ReLU	
Output layer, # neurons	$2^k = 8$	
Output layer, activ. func.	Softmax	
Optimizer	Adam	
Learning rates	$\alpha = 0.001, \beta_1 = 0.9, \beta_2 = 0.999$	
Loss function	Categorical cross entropy	
Mini-batch, # instances	32	
Threshold Optimization		
Evaluated thresholds	$0.01, 0.02, 0.03, \ldots, 1$	
Datasets		
# snapshots per instance	$T = 100$	
# instances training set	$D_{\text{trn}} = 80\,000$	$D_{\text{trn}} = 320\,000$
Fraction classifier training	80%	
Fraction classifier validation	10%	
Fraction threshold optim.	10%	
# instances test set	$D_{\text{tst}} = 50\,000$	
Benchmark Algorithms		
AOA estimator	MUSIC	
Angle spectrum resolution	Low: $\Delta\theta$, High: $0.1°$	
Number of sources estimators	MDL, AIC	

A. SIMULATION CONDITIONS

The data for training and testing the proposed estimator are generated synthetically using the data model presented in Section II. Similar to [21], the simulations address two scenarios regarding the number of sources (NOS):

(I) the NOS Q is assumed to be constant over all instances, i.e. $Q = 2$, and

(II) the NOS Q varies over the different instances, i.e., Q is assumed to be a random variable drawn from the discrete uniform distribution $Q \sim U(1, 4)$, meaning up to 4 impinging waves are considered.

The following have been assumed for both scenarios. A uniform linear array (ULA) of $N = 8$ sensors with $\lambda/2$ inter-element spacing is considered, where λ is the wavelength of the considered plane waves. The sources transmit uncorrelated signals of equal power, i.e., $\mathbf{P} = \sigma^2 \mathbf{I}_Q$ (4). The waves' AOAs are assumed to be random variables following the continuous uniform distribution, i.e., $\theta_1, \ldots, \theta_Q \sim U(-60°, 60°)$. The array's field of view (FOV) is defined by the interval $[\theta_{\min}, \theta_{\max}) = [-60°, 60°)$. The number of

FOV segments evaluated by each classifier, i.e., the number of labels in a k-labelset, is set to $k = 3$, as suggested for RAkEL$_o$ in [22].

For both scenarios, simulations are performed to investigate the impact of the signal-to-noise ratio (SNR) σ^2/ν^2, the FOV resolution (represented by $\Delta\theta$) and the number of layers in the framework, L. Specifically, the following values are considered: SNR $\in \{-10, 0, 10\}$ dB, $\Delta\theta \in \{2°, 1°\}$ (meaning $M = 60$ and $M = 120$, respectively (8)) and $L \in \{1, 3, 5\}$. Hence, $2 \times 3 \times 2 \times 3 = 36$ (scenarios × SNRs × resolutions × framework layers) simulations are performed. Here, a 'simulation' comprises all three deployment stages presented in Section III-C. All random variables (NOS Q, AOAs $\theta_1, \ldots, \theta_Q$, waveforms $\mathbf{s}(t)$ and element noise $\mathbf{n}(t)$) follow the same distributions for all instances (an instance being a collection of T snapshots of the array output) within a simulation, whether they are training or testing instances. New realizations are generated for each instance (NOS and AOAs) and for each snapshot (waveforms and element noise) individually.

B. LEARNING-PARAMETERS AND DATA SETS
In this work, the feedforward neural network (FFNN) (see, e.g., [27]) is employed as the single-label learning algorithm. The FFNN is one of the simplest type of neural networks (NNs) that exist, but still allows for sufficient design freedom to fit in the proposed AOA estimation framework. FFNNs are composed of an input layer, one or multiple hidden layers and an output layer. Each layer consists of a number of neurons. The number of neurons in the input layer is imposed by the dimension of the feature vectors. In the present work, the feature vector of an instance is composed as

$$\mathbf{r} = [\hat{R}_{1,1}, \ldots, \hat{R}_{N,N}, \Re(\hat{R}_{1,2}), \Im(\hat{R}_{1,2}), \\ \Re(\hat{R}_{1,3}), \Im(\hat{R}_{1,3}), \ldots]^T, \quad (27)$$

where $\hat{R}_{i,j}$ is the element at row i and column j of $\hat{\mathbf{R}}$ and $\hat{\mathbf{R}}$ is computed according to (7) using $T = 100$ snapshots of the array output (1). Hence, the number of neurons in the input layer equals N^2, with N being the number of sensors in the array.[6] Since the array data follow the Gaussian distribution, element-wise standardization is applied as the normalization algorithm, meaning all element-wise means and variances equal 0 and 1, respectively [25]. The number of hidden layers and the number of neurons in these layers were determined using the method presented in [27]. They are different for the different simulation scenarios, as can be seen in Table 2, with the sequence of numbers representing the number of neurons in the hidden layers from input-side to output-side. All hidden layers are fully connected (i.e., each neuron is connected to all neurons in both the previous and the next layer) and use the ReLU activation function [28]. The number of neurons in the

output layer is imposed by the RAkEL parameter k and equals 2^k. The output layer uses the Softmax activation function (see, e.g., [28]), meaning that all 2^k outputs are between 0 and 1 and add up to 1. Hence, they represent the probabilities $\tilde{P}_{j,\tilde{k}}$ ($j = 1, \ldots, m$ and $\tilde{k} = 1, \ldots, 2^k$) which are converted to AOA estimates according to the procedure described in Section III-B.

The NN training, i.e., the optimization of the NNs' weights, is performed using the Adam optimizer [29] in combination with the categorical cross entropy loss function (see, e.g., [27]). The default [29] learning rates of $\alpha = 0.001$, $\beta_1 = 0.9$ and $\beta_2 = 0.999$ are used and each weight update is based on a mini-batch of 32 training instances. The training of a particular NN is terminated if the loss on the validation set did not decrease for 3 consecutive epochs (iterations over the training set). All simulations are implemented in Python using the TensorFlow machine learning library [30].

The evaluated threshold levels in the threshold optimization branch (Section III-C) are 0, 0.01, 0.02, ..., 1.

The training set contains $D_{\text{trn}} = 80000$ instances for scenario I and $D_{\text{trn}} = 320000$ instance for scenario II. From all training instances, 80% is used for training the classifiers, 10% for validating them (i.e., determining when to stop training), and 10% for optimizing the threshold level. In all simulations, the estimator is tested using $D_{\text{tst}} = 50000$ test instances.

C. BENCHMARK ALGORITHMS
The joint AOA estimates obtained from the proposed estimator are compared (using the performance metrics presented in Section IV) to those obtained from the well-known MUSIC algorithm [5]. Since the MUSIC algorithm belongs to the separable detection category, a NOS estimate is required prior to estimating the AOAs. Two NOS estimators are considered: the minimum description length (MDL) and the Akaike information criterion (AIC) [3]. For each simulation, the MUSIC angle spectrum is evaluated at two angle resolutions: (I) a lower resolution, equal to the FOV resolution $\Delta\theta$ of the proposed AOA estimation framework, and (II) a higher resolution of $0.1°$.

VI. RESULTS AND ANALYSIS
In this section, simulation results are presented and analysed. Results pertaining the fixed number of sources (NOS) scenario are discussed first.

A. RESULTS SIMULATION SCENARIO I: FIXED NOS
1) NUMBER OF FRAMEWORK LAYERS AND FOV RESOLUTION
Fig. 3 shows the joint AOA estimation success rate $f_{\text{sr}}(\tilde{\theta})$ (19) for various values of the maximum allowed AOA estimation error $\tilde{\theta}$, for all 18 simulations conducted within this scenario (3 SNRs × 2 values for $\Delta\theta$ × 3 values for L). The expected success rate in the case of ideal classifiers $f_{\text{sr,exp}}(\tilde{\theta})$ (20), which depends on $\Delta\theta$ but not on the SNR nor on L, is shown

---

[6]Since $\hat{\mathbf{R}}$ is Hermitian, only the diagonal elements and the elements on the upper right half of (7) are used. In fact, in case of isotropic sensors as considered here, the diagonal elements do not contain any information. Still, we include them in the feature vector, such that the impact of physically more realistic arrays can be easily investigated in the future.

FIGURE 3. Joint AOA estimation success rate $f_{sr}(\tilde{\theta})$ of the proposed MLF vs. maximum allowed AOA estimation error $\tilde{\theta}$ for all simulations in scenario I. See Table 2.

FIGURE 4. Average F_1-scores $\bar{F}_1(Q_h)$ of the proposed MLF vs. label subset cardinality Q_h, for all simulations with $L = 5$ in Scenario I. See Table 2.

as a reference. As can be seen from Fig. 3, increasing the number of framework layers L increases the success rate $f_{sr}(\tilde{\theta})$ for all the six considered {SNR, $\Delta\theta$}-couples, although the improvements are limited, especially when comparing $L = 3$ and $L = 5$. Hence, we conclude that the general recommendation of using RAkEL$_o$ with $M < m < 2M$ and a small k [22], equivalent to using $3 < L < 6$ for $k = 3$ (10) in the layered framework proposed here, can be loosened for the present AOA application.

Fig. 3 also shows that for the two highest SNRs, the rate of successful AOA estimation is increased by using the higher FOV resolution ($\Delta\theta = 1°$) rather than the lower one ($\Delta\theta = 2°$), especially if $L \geq 3$. For example, the success rate at $\tilde{\theta} = 1°$ (i.e., assuming AOA estimation errors up to $1°$ are allowed) increases from 70.6% to 89.1% (0 dB SNR, $L = 5$) and from 83.1% to 96.2% (10 dB SNR, $L = 5$). Considering the tightness of the reference $f_{sr,exp}(\tilde{\theta})$ to the success rates achieved by the MLF at these SNRs for $\Delta\theta = 2°$ and the fact that the success rates increase (in absolute sense) when going to $\Delta\theta = 1°$, it is concluded that the performance of the MLF is limited by the FOV resolution when using $\Delta\theta = 2°$. On the contrary, when looking at the -10 dB SNR cases, it is observed that the success rate actually decreases when increasing the FOV resolution, e.g., from 21.4% to 12.3% for $L = 5$ and $\tilde{\theta} = 1°$. It is worthwhile to note that the resolution increase is obtained at the expense of an increased computational cost. That is because the number of classifiers (here, NNs) to be trained is inversely proportional with $\Delta\theta$, see (10) and (8). Hence, for the -10 dB SNR case, using the lower resolution is clearly the better option, both from

the AOA estimation accuracy perspective as well as from the resource perspective.

To get a better insight in the impact of the FOV resolution, we proceed by evaluating the predictive performance of the NNs by means of the averaged F_1-scores (Section IV-C).

2) NEURAL NETWORK PREDICTIVE PERFORMANCE

Fig. 4 shows the F_1-scores $\bar{F}_1(Q_h)$ (25) for the various subset cardinalities $Q_h = 0, \ldots, 2$. Note that even though $k = 3$, $\bar{F}_1(Q_h > 2)$ is not defined because $Q = 2$ in all simulations considered here. The results presented in Fig. 4 are based on the simulations with $L = 5$. Thus, $\bar{F}_1(Q_h)$ is computed by averaging the F_1-scores (22) of 100 and 200 NNs for the $\Delta\theta = 2°$ and $\Delta\theta = 1°$ frameworks, respectively (10). As can be seen from Fig. 4, $\bar{F}_1(Q_h)$ decreases when increasing the FOV resolution (i.e., decreasing $\Delta\theta$) for all SNRs and for all values of Q_h except $Q_h = 0$. This can be explained by a phenomenon called class imbalance [31]. Although a detailed discussion is outside the scope of this work, it is worthwhile to note that this effect is expected to get stronger when further increasing the FOV resolution, as more and more instances from the training set will correspond to $Q_h = 0$. Consequently, learning an accurate mapping for instances corresponding to other values of Q_h becomes more difficult.

While Fig. 4 shows that increasing the FOV resolution decreases the predictive performance at *all* considered SNRs, Fig. 3 shows that the joint AOA estimation success rate $f_{sr}(\tilde{\theta})$ only decreases at *low* SNR. This might sound paradoxical, but it is not: when increasing the FOV resolution while keeping the maximum allowed AOA estimation error $\tilde{\theta}$ fixed, one might (if $\tilde{\theta} > \Delta\theta/2$) obtain a successful AOA estimate also using non-perfect predictions. That this is indeed the case can be understood by evaluating the success rates relative to the references $f_{sr,exp}(\tilde{\theta})$. As can be seen from Fig. 3 for $L = 5$ and at mid and high SNR, the success rates are further apart from the references $f_{sr,exp}(\tilde{\theta})$ for the higher FOV resolution ($\Delta\theta = 1°$) than for the lower FOV resolution ($\Delta\theta = 2°$). As the references $f_{sr,exp}(\tilde{\theta})$ assume ideal classifiers, this indicates that indeed the NNs' predictions are further from ideal for the higher FOV resolution, as confirmed by results shown in Fig. 4. Interestingly, it is found that the threshold level, which is optimized during the training stage, increased from 0.05 to

TABLE 3. $P(\hat{Q} = Q)$ and RMSE for Scenario I. See Table 2.

	SNR [dB]	-10		0		10	
	$\Delta\theta$ [deg]	2	1	2	1	2	1
$P(\hat{Q} = Q)$ [%]	MLF, $L = 5$	**83.5**	50.8	94.5	**97.1**	94.7	**98.1**
	MDL		0.5		91.9		97.4
	AIC		63.7		86.8		89.8
RMSE [deg]	MLF, $L = 5$	**2.2**	13.1	0.7	**0.5**	0.6	**0.4**
	MUSIC $(\Delta\theta)$ + MDL	3.7	3.6	7.5	7.2	8.5	8.1
	MUSIC $(0.1°)$ + MDL		3.6		7.2		8.0
	MUSIC $(\Delta\theta)$ + AIC	9.5	9.4	10.3	10.2	9.4	9.1
	MUSIC $(0.1°)$ + AIC		9.4		10.1		9.1

FIGURE 5. $P(\hat{Q} = Q)$ and RMSE vs. AOA interval. 10 dB SNR, $\Delta\theta = 2°$ and $L = 5$, Scenario I. Results are based on a test set in which the sources' AOAs are symmetric with respect to broadside. All other parameters are as in Table 2, Scenario I.

0.22 (0 dB SNR) and from 0.04 to 0.23 (10 dB SNR) when increasing the FOV resolution from $\Delta\theta = 2°$ to $\Delta\theta = 1°$. This indicates that at the higher resolution, there are peaks in the probabilistic angle spectra at angles other than the AOAs that need to be filtered out. This is a direct consequence of incorrect classifier predictions. At low SNR (-10 dB), the situation is different, as the decreased predictive performance resulting from an increased FOV resolution caused the success rate $f_{sr}(\tilde{\theta})$ to decrease in absolute sense as well.

3) BENCHMARK COMPARISON

In this section, the joint AOA estimates of the proposed MLF are compared to those attained from the reference algorithms MDL, AIC (NOS estimates) and MUSIC (AOA estimates). Again, the results presented for the MLF are based on the $L = 5$ simulations.

Table 3 presents the NOS estimation accuracy $P(\hat{Q} = Q)$ and the root-mean-square error (RMSE) for all considered {SNR, $\Delta\theta$}-couples. The best performing algorithm (i.e., the one achieving the highest $P(\hat{Q} = Q)$ and the lowest RMSE) is highlighted in bold for each SNR. For both metrics, the MLF outperforms the benchmark algorithms at all considered SNRs, although this requires different FOV resolutions: again, the low resolution ($\Delta\theta = 2°$) is preferred for the -10 dB SNR case, whereas the high resolution ($\Delta\theta = 1°$) achieves better results at the mid and high range SNRs.

To clarify the relatively high RMSEs for the MUSIC algorithm, we plot both $P(\hat{Q} = Q)$ and the RMSE against the

AOA interval $|\theta_2 - \theta_1|$ in Fig. 5. To this end, we synthesized additional test sets (12000 instances) in which the AOAs of the two sources are symmetric with respect to the array's broadside, i.e., $\theta_q = 90\pm\delta$ degree. All other parameters are as in Table 2. The AOA interval $|\theta_2 - \theta_1| = |90+\delta-(90-\delta)| = 2\delta$ is assumed to be a random variable following a continuous log-uniform probability distribution between $1°$ and $120°$. We grouped the instances in these test sets based on their AOA interval and computed $P(\hat{Q} = Q)$ and the RMSE for each group separately, as indicated by the vertical grid and the stair-wise graphs in Fig. 5. In this way, we 'average out' (especially at large AOA intervals) the impact of the finite resolution which is inherent to both the MLF and the MUSIC algorithm. As an additional reference, the Cramér-Rao lower bound (CRLB), see, e.g., [2], [8], is shown as well. For the sake of conciseness, we only present results for the {10 dB SNR, $\Delta\theta = 2°$}-couple, but similar observations were made in the other considered cases as well. As can be seen from Fig. 5, MDL and AIC outperform the MLF at small AOA intervals. This is because in this specific symmetric scenario, an AOA interval of at least $2\Delta\theta = 4°$ is required for the MLF to be able to resolve both sources (Section III-B2). Hence, at these small AOA intervals, the MLF never estimates the NOS correctly and therefore, the RMSE cannot be computed. Contrarily, the RMSE for the MUSIC algorithm *does* exist at small intervals, although it is nearly 2 orders of magnitude larger than the CRLB (worst case). This can be understood as follows. Since the MUSIC algorithm belongs to the separable detection category, it aims to return as many AOA estimates as required according to the NOS detection method, here MDL/AIC. In case MDL/AIC manages to estimate the correct NOS, while at the same time the MUSIC angle spectrum does not contain distinct peaks at all AOAs (which might happen for small AOA intervals [1]), the argument of another peak in the spectrum is returned. This results in large AOA estimation errors, which dominate the RMSE values presented in Table 3. Although not shown for the sake of conciseness, it is observed that at lower SNRs, MDL generally fails to correctly estimate the NOS in case of small AOA intervals. Consequently, these instances are excluded from the RMSE computation, explaining why the RMSE for the MUSIC+MDL combination is lower at the lower SNRs. This discrepancy between the NOS estimators and the MUSIC algorithm at small AOA intervals complicates the comparison with the MLF, and hence emphasizes the advantage of the proposed joint AOA estimation success rate $f_{sr}(\tilde{\theta})$ (19), as this metric considers both NOS and AOA estimates. Hence, next we compare the proposed MLF and the MUSIC algorithm (combined with MDL/AIC) in terms of the joint AOA estimation success rate.

As can be seen in in Fig. 6, the proposed MLF outperforms the MUSIC algorithm if $\tilde{\theta} \gtrsim \Delta\theta$, i.e., if the maximum allowed AOA estimation error is approximately of the same order as (or larger than) the size of the FOV segments. This applies to all variants of the MUSIC algorithm considered (low/high angle spectrum resolution, see Table 2,

FIGURE 6. Joint AOA estimation success rate $f_{sr}(\tilde{\theta})$ vs. maximum allowed AOA estimation error $\tilde{\theta}$, proposed MLF (with $L = 5$) vs. MUSIC for simulations in Scenario I. See Table 2.


FIGURE 7. Joint AOA estimation success rate $f_{sr}(\tilde{\theta})$ vs. maximum allowed AOA estimation error $\tilde{\theta}$, proposed MLF (with $L = 5$) vs. MUSIC for simulations in Scenario II. See Table 2.

and MDL/AIC NOS detection) and to almost all {SNR, $\Delta\theta$}-couples. Only for the {−10 dB SNR, $\Delta\theta = 1°$}-couple, the MUSIC+AIC combination attains a higher success rate than the MLF. Contrarily, if $\tilde{\theta} < \Delta\theta$, the high resolution MUSIC algorithm outperforms the MLF for the mid and high SNRs. This is a direct consequence of the finite FOV resolution of the MLF, because of which errors up to $\Delta\theta/2$ are to be expected, as already illustrated by $f_{sr,exp}(\tilde{\theta})$ (20) in Fig. 3.

B. RESULTS SIMULATION SCENARIO II: VARIABLE NOS

Next, we present an analysis of the simulation results pertaining the variable NOS scenario. For the sake of conciseness, we limit ourselves to the benchmark comparison, as the phenomena observed in Section VI-A, e.g., limited improvements for $L > 3$ and a decreasing predictive performance for increasing Q_h due to class imbalance, apply here as well.

1) BENCHMARK COMPARISON

Fig. 7 shows the joint AOA estimation success rate $f_{sr}(\tilde{\theta})$, plotted against maximum allowed AOA estimation error $\tilde{\theta}$, for all considered {SNR, $\Delta\theta$}-couples. Again, the results shown for the MLF were obtained using a framework with $L = 5$ layers. As can be seen, the success rates for the MLF and for the MUSIC algorithm follow the same trends as in scenario I (Fig. 6), although they have decreased in absolute sense for all values of $\tilde{\theta}$ for both algorithms. Contrary to scenario I, the MLF now outperforms the MUSIC-AIC combination in the {−10 dB SNR, $\Delta\theta = 1°$}-case as well. Nevertheless, still the $\Delta\theta = 2°$ MLF achieves higher success rates than the $\Delta\theta = 1°$ MLF at this low SNR.

To get more insight into the impact of the various NOSs on the estimator performance, we group all test instances based on the NOS Q and evaluate the success rate $f_{sr}(\tilde{\theta} = 1°)$

FIGURE 8. Success rate $f_{sr}(1°)$ vs. number of sources Q for MLF and MUSIC. Scenario II.

for each of them separately (Fig. 8). Hence, AOA estimation errors up to 1° are considered acceptable. It is worthwhile to note that the relative maximum allowed AOA estimation error $\tilde{\theta}/\Delta\theta$ is larger for the high resolution framework ($\Delta\theta = 1°$) than for the low resolution framework ($\Delta\theta = 2°$). Consequently, a higher success rate can be achieved by the high resolution MLF, even though the NNs have a lower predictive performance (not shown for the sake of conciseness) than those in the low resolution MLF. This was also observed in scenario I, for SNRs equal to 0 and 10 dB (see Fig. 6 at $\tilde{\theta} = 1°$ and Fig. 4). As can be seen from Fig. 8, the success

rate decreases for increasing NOSs Q, both for the MLF and the MUSIC algorithm. We conclude that in this simulation scenario and for this particular maximum allowed AOA estimation error ($\tilde{\theta} = 1°$), only the MLF with low FOV resolution ($\Delta\theta = 2°$) is outperformed by the (high resolution) MUSIC algorithm, and only at SNRs of 0 and 10 dB. In all other cases, the MLF achieves the highest success rate for all considered values of Q.

VII. CONCLUSION
In this paper, we proposed a machine learning framework (MLF) which jointly estimates the number of sources (NOS) and the angles-of-arrival (AOAs) of plane waves impinging a sensor array. The MLF is tailored to the array's segmented field of view (FOV) such that it can solve the joint AOA estimation problem through supervised-learning-based classification. The proposed approach is general in the sense that the MLF can, in principle, be implemented in combination with any single-label multi-class classification algorithm. Moreover, a new performance metric, the joint AOA estimation success rate, is introduced to assess the performance of the proposed MLF. Particularly, this metric depends on the user-defined maximum allowed AOA estimation error. Numerical simulations are conducted using feedforward neural networks as the learning algorithm. In scenarios representing both fixed and variable NOSs, results show that in terms of the joint AOA estimation success rate, the optimal FOV segmentation strongly depends on the signal-to-noise ratio (SNR). When increasing the FOV resolution from 2° to 1° while keeping the learning settings the same, the achieved success rate deteriorates at low SNR (-10 dB), for all considered values of the maximum allowed AOA estimation error. On the contrary, at mid (0 dB) and high (10 dB) SNRs, the success rate increases when increasing the FOV resolution. The FOV resolution is inversely proportional to the number of classifiers in the MLF. Hence, at mid and high range SNRs, an important trade-off between the estimation performance and the computational burden is to be considered. In nearly all considered cases, the MLF outperforms the multiple signal classification (MUSIC) algorithm, implemented in conjunction with the NOS estimator Akaike's information criterion (AIC) or the minimum description length (MDL). Only in case of a fixed NOS and at low SNR, the MLF with high FOV resolution (1°) is outperformed by the MUSIC-AIC combination. We conclude that the proposed MLF offers a higher rate of successful joint AOA estimation for all SNRs if the maximum allowed AOA estimation error is of the order of (or larger than) the size of the FOV segments and if the FOV resolution is selected with care.

Further research is required to determine the optimal FOV resolution for a given SNR (distribution). Moreover, further investigation into the use of alternative learning algorithms and the use of different FOV discretisations for the different framework layers is recommended. Finally, studying the use of physically realistic sensor arrays, the impact of their layout, and the impact of correlated sources on the estimation

accuracy of the proposed estimator in comparison to conventional estimators like the MUSIC algorithm is of great interest for practical applications.

APPENDIX A
EXPECTED SUCCESS RATE IDEAL CLASSIFIERS
Consider the independent and identically distributed random variables $\theta_1, \ldots, \theta_Q$, drawn from the continuous uniform distribution $U(\theta_{\min}, \theta_{\max})$. Assume the interval $[\theta_{\min}, \theta_{\max})$ is segmented in M intervals $[\theta_{i,\min}, \theta_{i,\max})$ ($i = 1, \ldots, M$) in a regular manner, meaning each interval has size $\Delta\theta = (\theta_{\max} - \theta_{\min})/M$. Let's denote the center of the i^{th} interval $c_i = (\theta_{i,\min} + \theta_{i,\max})/2$. Then, the probability P that all $\theta_1, \ldots, \theta_Q$ are at most $\tilde{\theta}$ removed from one of the interval centers $c_1, \ldots, c_M$ (for $\tilde{\theta} < \Delta\theta/2$) is computed as

$$
\begin{aligned}
P(\tilde{\theta})\big|_{\tilde{\theta}<\Delta\theta/2} &= \prod_{q=1}^{Q} M \int_{-\tilde{\theta}}^{\tilde{\theta}} \frac{1}{\theta_{\max} - \theta_{\min}} d\theta_q \\
&= \left[\frac{\theta_{\max} - \theta_{\min}}{\Delta\theta} \int_{-\tilde{\theta}}^{\tilde{\theta}} \frac{1}{\theta_{\max} - \theta_{\min}} d\theta \right]^{Q} \\
&= \left[\frac{\tilde{\theta}}{\Delta\theta/2} \right]^{Q}.
\end{aligned}
\tag{28}
$$

Clearly, if $\tilde{\theta} \geq \Delta\theta/2$, $P(\tilde{\theta}) = 1$, since the closest c_i is at a distance of at most $\Delta\theta/2$ from any point in the interval $[\theta_{\min}, \theta_{\max})$. Hence, it follows that

$$
P(\tilde{\theta}) = \begin{cases} \left(\frac{\tilde{\theta}}{\Delta\theta/2} \right)^{Q} & \text{if } \tilde{\theta} < \Delta\theta/2 \\ 1 & \text{otherwise.} \end{cases}
\tag{29}
$$

APPENDIX B
MAXIMUM EXPECTED SUCCESS RATE
Consider performing random sampling with replacement from the set $\{\lambda_1, \ldots, \lambda_M\}$, where the likelihood of selecting a particular λ_i ($i = 1, \ldots, M$) is equal for all of them. Hence, when sampling Q times, M^Q possible outcomes (permutations) exist. Assume we want to compute the percentage p of these M^Q permutations which fulfill the requirements that (I) none of the λ_i is selected multiple times, and (II) no neighbouring λ_i are selected, i.e., when λ_i is selected, λ_{i-1} and λ_{i+1} are not. Here, the latter requirement reduces to either λ_{i+1} or λ_{i-1} if $i = 1$ or $i = M$, respectively. This can be interpreted as random sampling without replacement Q times from a set of $M - (Q - 1)$ elements, for which the number of combinations equals $\binom{M-(Q-1)}{Q}$. Multiplying this by $Q!$ converts the combinations to permutations, meaning that the percentage of permutations fulfilling requirements (I) and (II) is computed as

$$
p = \binom{M - (Q - 1)}{Q} \frac{Q!}{M^Q} \times 100\%.
\tag{30}
$$

REFERENCES
[1] H. Krim and M. Viberg, "Two decades of array signal processing research: The parametric approach," *IEEE Signal Process. Mag.*, vol. 13, no. 4, pp. 67–94, Jul. 1996.

[2] H. L. Van Trees, *Optimum Array Processing: Part IV Detection, Estimation, Modulation Theory*. Hoboken, NJ, USA: Wiley, 2004.

[3] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-33, no. 2, pp. 387–392, Apr. 1985.

[4] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *Proc. IEEE*, vol. 57, no. 8, pp. 1408–1418, Aug. 1969.

[5] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. AP-34, no. 3, pp. 276–280, Mar. 1986.

[6] R. Roy and T. Kailath, "Esprit-estimation of signal parameters via rotational invariance techniques," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 7, pp. 984–995, Jul. 1989.

[7] A. Barabell, "Improving the resolution performance of eigenstructure-based direction-finding algorithms," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, vol. 8, Sep. 1983, pp. 336–339.

[8] P. Stoica and A. Nehorai, "MUSIC, maximum likelihood, and cramer-rao bound," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 5, pp. 720–741, Mar. 1989.

[9] I. Ziskind and M. Wax, "Maximum likelihood localization of multiple sources by alternating projection," *IEEE Trans. Acoust., Speech Signal Process.*, vol. ASSP-36, no. 10, pp. 1553–1560, Oct. 1988.

[10] Z. Yang, L. Xie, and C. Zhang, "Off-grid direction of arrival estimation using sparse Bayesian inference," *IEEE Trans. Signal Process.*, vol. 61, no. 1, pp. 38–43, Jan. 2013.

[11] P. Chen, Z. Cao, Z. Chen, and X. Wang, "Off-grid DOA estimation using sparse Bayesian learning in MIMO radar with unknown mutual coupling," *IEEE Trans. Signal Process.*, vol. 67, no. 1, pp. 208–220, Jan. 2019.

[12] Z.-M. Liu, Z.-T. Huang, and Y.-Y. Zhou, "An efficient maximum likelihood method for direction-of-arrival estimation via sparse Bayesian learning," *IEEE Trans. Wireless Commun.*, vol. 11, no. 10, pp. 1–11, Oct. 2012.

[13] Z. Yang, J. Li, P. Stoica, and L. Xie, "Sparse methods for direction-of-arrival estimation," in *Academic Press Library in Signal Processing*, vol. 7. Amsterdam, The Netherlands: Elsevier, 2018, pp. 509–581.

[14] A. Khan, S. Wang, and Z. Zhu, "Angle-of-arrival estimation using an adaptive machine learning framework," *IEEE Commun. Lett.*, vol. 23, no. 2, pp. 294–297, Feb. 2019.

[15] W. Zhu, M. Zhang, P. Li, and C. Wu, "Two-dimensional DOA estimation via deep ensemble learning," *IEEE Access*, vol. 8, pp. 124544–124552, 2020.

[16] Y. Kase, T. Nishimura, T. Ohgane, Y. Ogawa, D. Kitayama, and Y. Kishiyama, "DOA estimation of two targets with deep learning," in *Proc. 15th Workshop Positioning, Navigat. Commun. (WPNC)*, Oct. 2018, pp. 1–5.

[17] A. M. Ahmed, U. S. K. M. Thanthrige, A. E. Gamal, and A. Sezgin, "Deep learning for DOA estimation in MIMO radar systems via emulation of large antenna arrays," *IEEE Commun. Lett.*, vol. 25, no. 5, pp. 1559–1563, May 2021.

[18] M. Pastorino and A. Randazzo, "A smart antenna system for direction of arrival estimation based on a support vector regression," *IEEE Trans. Antennas Propag.*, vol. 53, no. 7, pp. 2161–2168, Jul. 2005.

[19] O. Bialer, N. Garnett, and T. Tirer, "Performance advantages of deep neural networks for angle of arrival estimation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 3907–3911.

[20] Z.-M. Liu, C. Zhang, and P. S. Yu, "Direction-of-arrival estimation based on deep neural networks with robustness to array imperfections," *IEEE Trans. Antennas Propag.*, vol. 66, no. 12, pp. 7315–7327, Dec. 2018.

[21] G. Papageorgiou, M. Sellathurai, and Y. Eldar, "Deep networks for direction-of-arrival estimation in low SNR," *IEEE Trans. Signal Process.*, vol. 69, pp. 3714–3729, 2021.

[22] G. Tsoumakas, I. Katakis, and L. Vlahavas, "Random K-labelsets for multilabel classification," *IEEE Trans. Knowl. Data Eng.*, vol. 23, no. 7, pp. 1079–1089, Jul. 2011.

[23] M.-L. Zhang and Z.-H. Zhou, "A review on multi-label learning algorithms," *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 8, pp. 1819–1837, Aug. 2014.

[24] N. B. Kanters, "Direction-of-arrival estimation of an unknown number of signals using a machine learning framework," M.S. thesis, Dept. Elect. Eng., Univ. Twente, Enschede, The Netherlands, 2020.

[25] A. Zheng and A. Casari, *Feature Engineering for Machine Learning: Principles and Techniques for Data Scientists*. Sebastopol, CA, USA: O'Reilly Media, 2018.

[26] G. Tsoumakas and I. Vlahavas, "Random K-labelsets: An ensemble method for multilabel classification," in *Proc. Eur. Conf. Mach. Learn.* Berlin, Germany: Springer, 2007, pp. 406–417.

[27] F. Chollet, *Deep Learning With Python*. Shelter Island, NY, USA: Manning Publications Co., 2017.

[28] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep Learning*, vol. 1. Cambridge, MA, USA: MIT Press, 2016.

[29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.

[30] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, and S. Ghemawat, "TensorFlow: Large-scale machine learning on heterogeneous distributed systems," 2016, *arXiv:1603.04467*.

[31] M. Galar, A. Fernandez, E. Barrenechea, H. Bustince, and F. Herrera, "A review on ensembles for the class imbalance problem: Bagging-, boosting-, and hybrid-based approaches," *IEEE Trans. Syst., Man, C, Appl. Rev.*, vol. 42, no. 4, pp. 463–484, Jul. 2011.

NOUD KANTERS (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the University of Twente, Enschede, The Netherlands, in 2016 and 2020, respectively, where he is currently pursuing the Ph.D. degree with the Radio Systems Group. His current research interests include angle-of-arrival estimation, fundamental limitations on antenna-channel interactions, and machine learning.

ANDRÉS ALAYÓN GLAZUNOV (Senior Member, IEEE) was born in Havana, Cuba. He received the M.Sc. (Engineer-Researcher) degree in physical engineering from Peter the Great St. Petersburg Polytechnic University (Polytech), St. Petersburg, Russia, in 1994, the Ph.D. degree in electrical engineering from Lund University, Lund Sweden, in 2009, and the Docent (Habilitation) degree in antenna systems from the Chalmers University of Technology, Gothenburg, Sweden, in 2017.

From 1996 to 2005, he held various research and specialist positions at the Telecom industry, e.g., Ericsson Research, Telia Research, and TeliaSonera, in Stockholm, Sweden. From 2001 to 2005, he was the Swedish Delegate to the European Cost Action 273, and from 2018 to 2020, he was the Dutch Delegate to the European Cost Action IRACON. He is one of the pioneers in producing the first standardized OTA measurement techniques for 3GPP, and devising novel OTA techniques, e.g., the random-LOS and the hybrid antenna characterization techiques. He has contributed to, or initiated various European research projects, e.g., more recently, the is3DMIMO, the WAVECOMBE, the 5VC and the Build-Wise projects under the auspices of the H2020 European Research and Innovation Program. He has also contributed to the international 3GPP and the ITU standardization bodies. From 2009 to 2010, he held the Marie Curie Senior Research Fellowship at the Centre for Wireless Network Design, University of Bedfordshire, Luton, U.K. From 2010 to 2014, he held a postdoctoral position with the Electromagnetic Engineering Laboratory, KTH-Royal Institute of Technology, Stockholm. From 2014 to 2018, he was an Assistant Professor at the Chalmers University of Technology. He is currently an Associate Professor with the Department of Electrical Engineering, University of Twente, Enschede, the Netherlands, where he is leading the antenna systems, propagation and OTA research. He is also an Affiliate Associate Professor with the Chalmers University of Technology, where he is leading the OTA characterization of antenna systems research area. He is the author of more than 150 scientific and technical publications. He is the coauthor and Co-Editor of the text book *LTE-Advanced and Next Generation Wireless Networks—Channel Modelling and Propagation* (Wiley, 2012). His current research interests include mmWave sensor array design, MIMO antenna systems, electromagnetic theory, fundamental limitations on antenna-channel interactions, radio propagation channel measurements, modeling and simulations, wireless performance in the built environment, and the OTA characterization of antenna systems and wireless devices.

• • •