

Calibrating AI Models for Wireless Communications via Conformal Prediction

KFIR M. COHEN¹ (Student Member, IEEE), SANGWOO PARK¹ (Member, IEEE), OSVALDO SIMEONE¹ (Fellow, IEEE), AND SHLOMO SHAMAI (SHITZ)² (Life Fellow, IEEE)

¹King's Communication, Learning and Information Processing (KCLIP) Laboratory, Department of Engineering, King's College London, WC2R 2LS London, U.K.

²Viterbi Faculty of Electrical and Computing Engineering, Technion—Israel Institute of Technology, Haifa 3200003, Israel

CORRESPONDING AUTHOR: S. PARK (sangwoo.park@kcl.ac.uk)

The work of Kfir M. Cohen and Sangwoo Park was supported by the European Research Council (ERC) through the European Union's Horizon 2020 Research and Innovation Programme under Grant 725731. The work of Osvaldo Simeone was supported in part by the European Research Council (ERC) through the European Union's Horizon 2020 Research and Innovation Programme under Grant 725731; in part by the Open Fellowship of the Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/W024101/1; in part by the European Union's Horizon Europe Project CENTRIC under Grant 101096379; and in part by the Realising Enabling Architectures and Solutions for Open Networks (REASON) Project, a U.K. Government funded project through the Future Open Networks Research Challenge (FONRC) sponsored by the Department of Science Innovation and Technology (DSIT). The work of Shlomo Shamai was supported in part by the European Union's Horizon 2020 Research and Innovation Programme under Grant 694630 and in part by the German Research Foundation (DFG) via the German–Israeli Project Cooperation (DIP) under Project SH 1937/1-1.

ABSTRACT When used in complex engineered systems, such as communication networks, artificial intelligence (AI) models should be not only as accurate as possible, but also well calibrated. A well-calibrated AI model is one that can reliably quantify the uncertainty of its decisions, assigning high confidence levels to decisions that are likely to be correct, and low confidence levels to decisions that are likely to be erroneous. This paper investigates the application of conformal prediction as a general framework to obtain AI models that produce decisions with formal calibration guarantees. Conformal prediction transforms probabilistic predictors into set predictors that are guaranteed to contain the correct answer with a probability chosen by the designer. Such formal calibration guarantees hold irrespective of the true, unknown, distribution underlying the generation of the variables of interest, and can be defined in terms of ensemble or time-averaged probabilities. In this paper, conformal prediction is applied for the first time to the design of AI for communication systems in conjunction with both frequentist and Bayesian learning, focusing on the key tasks of demodulation, modulation classification, and channel prediction. For demodulation and modulation classification, we apply both validation-based and cross-validation-based conformal prediction; while we investigate the use of online conformal prediction for channel prediction. For each task, we evaluate the probability that the set predictor contains the true output, validating the theoretical coverage guarantees of conformal prediction, as well as the informativeness of the predictor via the average predicted set size.

INDEX TERMS Bayesian learning, calibration, conformal prediction, cross-validation, reliability, set prediction, wireless communications.

I. INTRODUCTION

A. MOTIVATION

HOW reliable is your artificial intelligence (AI)-based model? The most common metric to design an AI model and to gauge its performance is the average *accuracy*. However, in applications in which AI decisions are used within a larger system, AI models should not only be as accurate as possible, but they should also be able to reliably quantify the uncertainty of their decisions. As an example, consider an unlicensed link that uses AI tools

to predict the best channel to access out of four possible channels. A predictor that assigns the probability vector of [90%, 2%, 5%, 3%] to the possible channels predicts the same best channel – the first – as a predictor that outputs the probability vector [30%, 20%, 25%, 25%]. However, the latter predictor is less certain of its decision, and it may be preferable for the unlicensed link to refrain from accessing the channel when acting on less confident predictions, e.g., to avoid excessive interference to licensed links [1], [2].

As in the example above, AI models typically report a confidence measure associated with each prediction, which reflects the model's *self-evaluation* of the accuracy of a decision. Notably, neural network models implement *probabilistic predictors* that produce a probability distribution across all possible values of the output variable. The self-reported model confidence, however, may not be a reliable measure of the true, unknown, accuracy of a prediction. In such situations, the AI model is said to be poorly calibrated.

As illustrated in the example in Fig. 1, accuracy and calibration are distinct criteria, with neither criterion implying the other. It is, for instance, possible to have an accurate predictor that consistently underestimates the accuracy of its decisions, and/or that is overconfident where making incorrect decisions (see fourth column in Fig. 1). Conversely, one can have inaccurate predictions that estimate correctly their uncertainty (see fifth column in Fig. 1).

Deep learning models tend to produce either overconfident decisions [3], or calibration levels that rely on strong assumptions about the ground-truth, unknown, data generation mechanism [4], [5], [6], [7], [8], [9]. This paper investigates the use of *conformal prediction* (CP) [10], [11], [12] as a framework to design provably well-calibrated AI predictors, with *distribution-free* calibration guarantees that do not require making any assumption about the ground-truth data generation mechanism.

B. CONFORMAL PREDICTION FOR AI-BASED WIRELESS SYSTEMS

CP leverages probabilistic predictors to construct well-calibrated *set predictors*. Instead of producing a probability vector, as in the examples in Fig. 1, a set predictor outputs a subset of the output space, as exemplified in Fig. 2. A set predictor is *well calibrated* if it contains the correct output with a pre-defined *coverage* probability selected by the system designer. For a well-calibrated set predictor, the size of the prediction set for a given input provides a measure of the uncertainty of the decision. Set predictors with smaller average prediction size are said to be more *efficient* [10].

This paper investigates CP as a general mechanism to obtain AI models with formal calibration guarantees for communication systems. The calibration guarantees of CP hold irrespective of the true, unknown, distribution underlying the generation of the variables of interest, and are defined either in terms of ensemble averages [10] or in terms of long-term averages [14]. CP is applied in conjunction to both frequentist and Bayesian learning, and specific applications are discussed to demodulation, modulation classification, and channel prediction.

C. RELATED WORK

Most work on AI for communications relies on conventional frequentist learning tools (see, e.g., the review papers [15], [16], [17], [18]). *Frequentist learning* is based on the minimization of the (regularized) training loss, which is

interpreted as an estimate of the ground-truth population loss. When data is scarce, this estimate is unreliable, and hence the focus on a single, optimized, model parameter vector often yields probabilistic predictors that are poorly calibrated, producing overconfident decisions [3], [19], [20], [21].

Bayesian learning offers a principled way to address this problem [22], [23]. This is done by producing as the output of the learning process not a single model parameter vector, but rather a distribution in the model parameter space, which quantifies the model's epistemic uncertainty caused by limited access to data. A model trained via Bayesian learning produces probabilistic predictions that are averaged over the trained model parameter distribution. This *ensembling* approach to prediction ensures that disagreements among models that fit the training data (almost) equally well are accounted for, substantially improving model calibration [24], [25].

In practice, Bayesian learning is implemented via approximations such as variational inference (VI) or Monte Carlo (MC) sampling, yielding scalable learning solutions [23]. VI methods approximate the exact Bayesian posterior distribution with a tractable variational density [26], [27], [28], [29], while MC techniques obtain approximate samples from the Bayesian posterior distribution [30], [31], [32]. Among other applications to communications systems, Bayesian learning was studied for network allocation in [33], [34], and [35], for massive MIMO detection in [36], [37], and [38], for channel estimation in [39], [40], and [41], for user identification in [42], and for multi user detection in [43] and [44]. Extensions to Bayesian meta-learning have been investigated in [20].

Exact Bayesian learning offers formal guarantees of calibration only under the assumption that the assumed model is *well specified* [4], [5]. In practice, this means that the assumed neural network models should have sufficient capacity to represent the ground-truth data generation mechanism, and that the predictive uncertainty should be unimodal for continuous outputs (since conventional likelihoods are unimodal, e.g., Gaussian) [5], [23], [24]. These assumptions are easily violated in practice, especially in communication systems in which lower-complexity models must be implemented on edge devices, and access to data for specific network configurations is limited. Specific examples are provided in [21] for applications including modulation classification [45], [46] and localization [47], [48].

Robustified versions of Bayesian learning that are based on the optimization of a modified free energy criterion were shown empirically to partly address the problem of model misspecification [4], [5], with implications for communication systems presented in [21]. However, robust Bayesian learning solutions do not have formal guarantees of calibration in the presence of misspecified models.

Another family of methods that aim at enhancing the calibration of probabilistic models implement a validation-based post-processing phase. Platt scaling [49] and temperature scaling [3] find a fixed parametric mapping of the trained

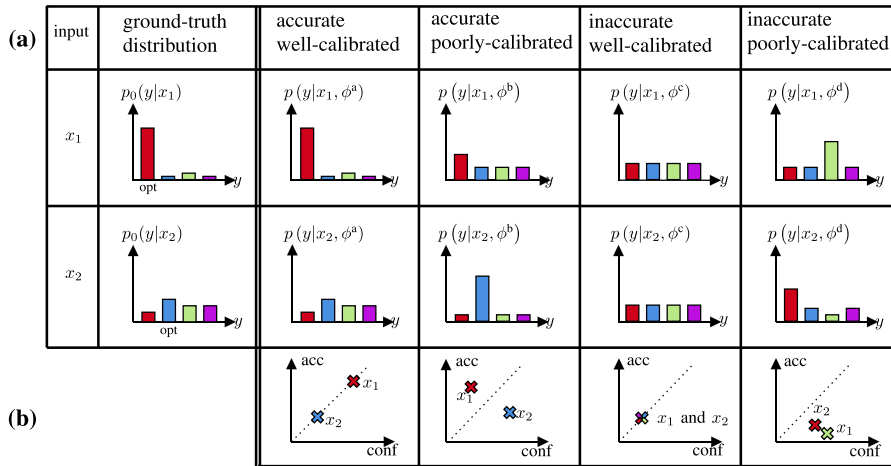


FIGURE 1. (a) Examples of probabilistic predictors for two inputs x_1 and x_2 : As compared to the ground-truth distribution in the second column, the first predictor (third column) is accurate, assigning the largest probability to the optimal decision (indicated as “opt” in the second column) and also well calibrated, reproducing the true accuracy of the decision; the second predictor (fourth column) is still accurate, but it is underconfident (for input x_1) and overconfident on the correct decision (for input x_2); the third predictor (fifth column) is not accurate, producing a uniform distribution across all output values, but is well calibrated if the data set is balanced [13]; and the last predictor (sixth column) is both inaccurate and poorly calibrated, providing overconfident decisions. (b) Confidence versus accuracy for the decisions made by the corresponding predictors.

input	ground-truth distribution	well-calibrated efficient	well-calibrated inefficient	poorly-calibrated efficient
x_1		{▲}	{▲★}	{◆}
x_2		{★◆}	{▲★◆}	{★}
x_3		{▲★}	{▲★◆}	{★}

FIGURE 2. Set predictors produce subsets of the range of the output variable (here {▲★◆}) for each input. Calibration is measured with respect to a desired coverage level $1 - \alpha$: A set predictor is well calibrated if the true label is included in the prediction set with probability at least $1 - \alpha$. A well-calibrated set predictor can be inefficient if it returns excessively large set predictions (forth column). In contrast, a poorly-calibrated set predictor (fifth column) returns set predictions that include the true value of the label with a probability smaller than $1 - \alpha$.

model output that minimizes the validation loss, while isotonic regression [50] applies a non-parametric binning approach. These recalibration-based approaches cannot guarantee calibration, as they may overfit the validation data set [51] and they are sensitive to the inaccuracy of the starting model [52].

Conformal prediction is a general framework for the design of set predictors that satisfy formal, *distribution-free*, guarantees of calibration [10], [11]. Given a desired miscoverage

probability α , CP returns set predictions that include the correct output value with probability at least $1 - \alpha$ under the only assumption that the data distribution is *exchangeable*. This condition is weaker than the standard assumption of “i.i.d.” data made in the design of most machine learning systems.

The original work on CP, [10], introduced *validation-based* CP and *full* CP. Since then, progress has been made on reducing computational complexity, minimizing the size of the prediction sets, and further alleviating the assumptions of exchangeability. *Cross-validation-based* CP was proposed in [53] to reduce the computational complexity as compared to full CP, while improving the efficiency of validation-based CP. The authors of [54], [55] proposed the optimization of a *CP-aware loss* to improve the efficiency of validation-based CP, while avoiding the larger computational cost of cross-validation. The work [56] proposed reweighting as a means to handle distribution shifts between the examples in the data set and the test point. Other research directions include improvements in the training algorithms [57], [58], and the introduction of novel calibration metrics [59], [60]. Finally, *online* CP, presented in [14] and [61], was shown to achieve long-term calibration over time without requiring statistical assumptions on the data generation.

CP provably guarantees calibration at the cost of relaxing point predictions to sets of point predictions. Another line of work on formal, distribution-free, post-hoc calibration introduces *Venn predictors* which relax probabilistic predictors to sets of probabilistic predictors [10, Sec. 6], [62], [63], [64]. Such sets contain at least one well-calibrated probabilistic

predictor, and the size of each set is no larger than the number of possible labels. In general, the sets of probabilistic predictor produced by Venn predictors, unlike those output by CP, do not appear to offer actionable information for decision processes. In contrast, Venn predictors have been explored in [65] for *explanation* purposes via perturbations of input features [66].

D. MAIN CONTRIBUTIONS

To the best of our knowledge, with the exception of the conference version [67] of this paper, this is the first work to investigate the application of CP to the design of AI models for communication systems. The main contributions of this paper are as follows.

- We provide a self-contained introduction to CP by focusing on validation-based CP [10], cross-validation-based CP [53], and online conformal prediction [61]. The presentation details connections to conventional probabilistic predictors, as well as the performance metrics used to assess calibration and efficiency.
- We propose the application of offline CP to the problems of symbol demodulation and modulation classification. The experimental results validate the theoretical property of CP methods of providing well-calibrated decisions. Furthermore, they demonstrate that naïve predictors that only rely on the output of either frequentist or Bayesian learning tools often result in poor calibration.
- Finally, we study the application of online CP to the problem of predicting received signal strength for over-the-air measured signals [68]. We demonstrate that online CP can obtain the predefined target long-term coverage rate at the cost of negligible increase in the prediction interval as compared to naïve predictors.

The conference version [67] of this work presented results only for symbol demodulation, while not providing background material on CP and not considering online CP. In contrast, this work is self-contained, presenting CP from first principles and including also online CP. Furthermore, this work investigates applications of CP to modulation classification and to channel prediction by leveraging real-world data sets [68], [69]. For reproducibility purposes, we have made our code publicly available.¹

The rest of this paper is organized as follows. In Sec. II, we define set predictors, and introduce the relevant performance metrics. Then, in Sec. III, naïve set predictors are introduced that do not provide guarantees in terms of calibration. Sec. IV describes conformal prediction, a general methodology to obtain well-calibrated set predictors. Sec. V details online conformal prediction, which is well suited for time-varying data. Applications to wireless communications are investigated in the following sections: Symbol demodulation is studied in Sec. VI; modulation classification in Sec. VII; and channel prediction in Sec. VIII. Sec. IX concludes the paper.

¹<https://github.com/kclip/cp4wireless>

II. PROBLEM DEFINITION

This section introduces set predictors, along with key performance metrics of coverage and inefficiency. To this end, we start by describing the data-generation model and reviewing probabilistic predictors.

A. DATA-GENERATION MODEL

We consider the standard supervised learning setting in which the learner is given a data set $\mathcal{D} = \{z[i]\}_{i=1}^N$ of N examples of input-output pairs $z[i] = (x[i], y[i])$ for $i = 1, \dots, N$, and is tasked with producing a prediction on a test input x with unknown output y . Writing $z = (x, y)$ for the test pair, data set \mathcal{D} and test point z follow the unknown *ground-truth*, or *population*, *distribution* $p_0(\mathcal{D}, z)$. Apart from Sec. V, we further assume throughout that the population distribution $p_0(\mathcal{D}, z)$ is *exchangeable* – a condition that includes as a special case the traditional independent and identically distributed (i.i.d.) data-generation setting. Note that we will not make explicit the distinction between random variables and their realizations, which will be clear from the context.

Mathematically, exchangeability requires that the joint distribution $p_0(\mathcal{D}, z)$ does not depend on the ordering of the $N + 1$ variables $\{z[1], \dots, z[N], z\}$. Equivalently, by de Finetti's theorem [70], there exists a latent random vector c with distribution $p_0(c)$ such that, conditioned on c , the variables $\{z[1], \dots, z[N], z\}$ are i.i.d. Writing the conditional i.i.d. distribution as

$$p_0(\mathcal{D}, z|c) = p_0(z|c) \prod_{i=1}^N p_0(z[i]|c) \quad (1)$$

for some ground-truth sampling distribution $p_0(z|c)$ given the variable c , under the exchangeability assumption, the joint distribution can be expressed as

$$p_0(\mathcal{D}, z) = \mathbb{E}_{p_0(c)}[p_0(\mathcal{D}, z|c)], \quad (2)$$

where $\mathbb{E}_{p(x)}[\cdot]$ denotes the expectation with respect to distribution $p(x)$.

The vector c in (2) can be interpreted as including context variables that determine the specific learning task. For instance, in a wireless communication setting, the vector c may encode information about channel conditions. In Sec. V, we will consider a more general setting in which no assumptions are made on the distribution of the data.

B. PROBABILISTIC PREDICTORS

Before introducing set predictors, we briefly review conventional probabilistic predictors. *Probabilistic predictors* implement a parametric conditional distribution model $p(y|x, \phi)$ on the output $y \in \mathcal{Y}$ given the input $x \in \mathcal{X}$, where $\phi \in \Phi$ is a vector of model parameters. Given the training data set \mathcal{D} , frequentist learning produces an optimized single vector $\phi_{\mathcal{D}}^*$, while Bayesian learning returns a distribution $q^*(\phi|\mathcal{D})$ on the model parameter space Φ [23], [24]. In either case, we will denote as $p(y|x, \mathcal{D})$ the resulting optimized

predictive distribution

$$p(y|x, \mathcal{D}) = \begin{cases} p(y|x, \phi_{\mathcal{D}}^*) & \text{for frequentist learning} \\ \mathbb{E}_{q^*(\phi|\mathcal{D})}[p(y|x, \phi)] & \text{for Bayesian learning.} \end{cases} \quad (3)$$

Note that the predictive distribution for Bayesian learning is obtained by averaging, or ensembling, over the optimized distribution $q^*(\phi|\mathcal{D})$. We refer to Appendix IX for basic background on frequentist and Bayesian learning.

From (3), one can obtain a point prediction \hat{y} for output y given input x as the probability-maximizing output as

$$\hat{y}(x|\mathcal{D}) = \operatorname{argmax}_{y' \in \mathcal{Y}} p(y'|x, \mathcal{D}). \quad (4)$$

In the case of a discrete set \mathcal{Y} , the hard predictor (4) minimizes the probability of detection error under the model $p(y|x, \mathcal{D})$. The probabilistic prediction $p(y|x, \mathcal{D})$ also provides a measure of predictive uncertainty for all possible outputs $y \in \mathcal{Y}$. In particular, for the point prediction $\hat{y}(x|\mathcal{D})$ in (4), we have the predictive, self-reported, *confidence level*

$$\operatorname{conf}(x|\mathcal{D}) = \max_{y' \in \mathcal{Y}} p(y'|x, \mathcal{D}) = p(\hat{y}(x|\mathcal{D})|x, \mathcal{D}). \quad (5)$$

As illustrated in Fig. 1, the performance of a probabilistic predictor can be evaluated in terms of both accuracy and *calibration*, with the latter quantifying the quality of uncertainty quantification via the confidence level (5) [3]. Specifically, a probabilistic predictor $p(y|x, \mathcal{D})$ is said to be *well calibrated* [3] if the probability that the hard predictor $\hat{y} = \hat{y}(x|\mathcal{D})$ equals the true label matches its confidence level π for all possible values of probability $\pi \in [0, 1]$. Mathematically, calibration is defined by the condition

$$\mathbb{P}(y = \hat{y} | p(\hat{y}|x, \mathcal{D}) = \pi) = \pi, \quad \text{for all } \pi \in [0, 1] \quad (6)$$

where the probability $\mathbb{P}(\cdot)$ follows the ground-truth distribution $p_0(x, y)$. Stronger definitions, like that introduced in [71], require the predictive distribution to match the ground-truth distribution also for values of y that are distinct from (4).

C. SET PREDICTORS

A *set predictor* is defined as a set-valued function $\Gamma(\cdot|\mathcal{D}) : \mathcal{X} \rightarrow 2^{\mathcal{Y}}$ that maps an input x to a subset of the output domain \mathcal{Y} based on data set \mathcal{D} . We denote the size of the set predictor for input x as $|\Gamma(x|\mathcal{D})|$. As illustrated in the example of Fig. 2, the set size $|\Gamma(x|\mathcal{D})|$ generally depends on input x , and it can be taken as a measure of the uncertainty of the set predictor.

The performance of a set predictor is evaluated in terms of calibration, or coverage, as well as of inefficiency. *Coverage* refers to the probability that the true label is included in the predicted set; while *inefficiency* refers to the average size $|\Gamma(x|\mathcal{D})|$ of the predicted set. There is clearly a trade-off between two metrics. A conservative set predictor that always produces the entire output space, i.e., $\Gamma(x|\mathcal{D}) = \mathcal{Y}$, would trivially yield a coverage probability equal to 1, but at the cost of exhibiting the worst possible inefficiency of $|\mathcal{Y}|$.




FIGURE 3. (a) The validity condition (7) assumed in offline CP is relevant if one is interested in the average performance with respect to realizations $(\mathcal{D}, \mathbf{z}) \sim p_0(\mathcal{D}, \mathbf{z})$ of training set \mathcal{D} and test variable $\mathbf{z} = (x, y)$. Input variable x is not explicitly shown in the figure, and the horizontal axis runs over the training examples in \mathcal{D} and the test example \mathbf{z} . (b) In online CP, the set predictor Γ_i uses its input $x[i]$ and all previously observed pairs $\mathbf{z}[1], \dots, \mathbf{z}[i-1]$ with $\mathbf{z}[i] = (x[i], y[i])$ to produce a prediction set. The long-term validity (21) assumed by online CP is defined as the empirical time-average rate at which the predictor Γ_i includes the true target variable $y[i]$.

Conversely, a set predictor that always produces an empty set, i.e., $\Gamma(x|\mathcal{D}) = \emptyset$, would achieve the best possible inefficiency, equal to zero, while also presenting the worst possible coverage probability equal to zero.

Let us denote a set predictor $\Gamma(\cdot|\cdot)$ for short as Γ . Formally, the *coverage* level of set predictor Γ is the probability that the true output y is included in the prediction set $\Gamma(x|\mathcal{D})$ for a test pair $\mathbf{z} = (x, y)$. This can be expressed as $\operatorname{coverage}(\Gamma) = \mathbb{P}(y \in \Gamma(x|\mathcal{D}))$, where the probability $\mathbb{P}(\cdot)$ is taken over the ground-truth joint distribution $p_0(\mathcal{D}, (x, y))$ in (2). The set predictor Γ is said to be $(1 - \alpha)$ -*valid* if it satisfies the inequality

$$\operatorname{coverage}(\Gamma) = \mathbb{P}(y \in \Gamma(x|\mathcal{D})) \geq 1 - \alpha. \quad (7)$$

When the desired coverage level $1 - \alpha$ is fixed by the predetermined target *miscoverage level* $\alpha \in [0, 1]$, we will also refer to set predictors satisfying (7) as being well calibrated.

Following the discussion in the previous paragraph, it is straightforward to design a valid, or well-calibrated, set predictor, even for the restrictive case of miscoverage level $\alpha = 0$. This can be, in fact, achieved by producing the full set $\Gamma(x|\mathcal{D}) = \mathcal{Y}$ for all inputs x . One should, therefore, also consider the inefficiency of predictor Γ . The *inefficiency* of set predictor Γ is defined as the average prediction set size

$$\operatorname{inefficiency}(\Gamma) = \mathbb{E}\left[|\Gamma(x|\mathcal{D})|\right], \quad (8)$$

where the average is taken over the data set \mathcal{D} and the test pair (x, y) following their exchangeable joint distribution $p_0(\mathcal{D}, (x, y))$.

In practice, the coverage condition (7) is relevant if the learner produces multiple predictions using independent data set \mathcal{D} , and is tested on multiple pairs (x, y) . In fact, in this case, the probability in (7) can be interpreted as the fraction of predictions for which the set predictor $\Gamma(x|\mathcal{D})$ includes the correct output. This situation, illustrated in Fig. 3(a), is quite common in communication systems, particularly at




FIGURE 4. A naïve probabilistic-based (NPB) set predictor uses a pre-trained probabilistic predictor to include all output values to which the probabilistic predictor assigns the largest probabilities that reach the coverage target $1 - \alpha$. This naïve scheme has no formal guarantee of calibration, i.e., it does not guarantee the coverage condition (7), unless the original probabilistic predictor is well calibrated.

the lower layers of the protocol stack. For instance, the data \mathcal{D} may correspond to pilots received in a frame, and the test point z to a symbol within the payload part of the frame (see Sec. VI). While the coverage condition (7) is defined under the assumption of a fixed ground-truth distribution $p_0(\mathcal{D}, z)$, in Sec. V we will allow for temporal distributional shifts and we will focus on validity metrics defined as long-term time averages (see Fig. 3(b)).

III. NAÏVE SET PREDICTORS

Before describing CP in the next section, in this section we review two naïve, but natural and commonly used, approaches to produce set predictors, that fail to satisfy the coverage condition (7).

A. NAÏVE SET PREDICTORS FROM PROBABILISTIC PREDICTORS

Given a probabilistic predictor $p(y|x, \mathcal{D})$ as in (3), one could construct a set predictor by relying on the confidence levels reported by the model. Specifically, aiming at satisfying the coverage condition (7), given an input x , one could construct the smallest subset of the output domain \mathcal{Y} that covers a fraction $1 - \alpha$ of the probability designed by model $p(y|x, \mathcal{D})$. Mathematically, the resulting *naïve probabilistic-based* (NPB) set predictor is defined as

$$\Gamma^{\text{NPB}}(x|\mathcal{D}) = \underset{\Gamma \subseteq 2^{\mathcal{Y}}}{\operatorname{argmin}} |\Gamma| \quad (9)$$

$$\text{s.t. } \sum_{y' \in \Gamma} p(y'|x, \mathcal{D}) \geq 1 - \alpha$$

for the case of a discrete set, and an analogous definition applies in the case of a continuous domain \mathcal{Y} . Fig. 4 illustrates the NPB for a prediction problem with output domain size $|\mathcal{Y}| = 4$. Given that, as mentioned in Sec. I, probabilistic predictors are typically poorly calibrated, the naïve set predictor (9) does not satisfy condition (7) for the given desired miscoverage level α , and hence it is not well calibrated. For example, in the typical case in which the probabilistic predictor is overconfident [3], the predicted sets (9) tend to be too small to satisfy the coverage condition (7).

B. NAÏVE SET PREDICTORS FROM QUANTILE PREDICTORS

While the naïve probabilistic-based set predictor (9) applies to both discrete and continuous target variables, we now focus on the important special case in which \mathcal{Y} is a real number, i.e., $\mathcal{Y} = \mathbb{R}$. This corresponds to scalar regression problems, such as for channel prediction (see Sec. VIII). Under this assumption, one can construct a naïve set predictor based on estimates of the $\alpha/2$ - and $(1 - \alpha/2)$ -quantiles $y_{\alpha/2}(x)$ and $y_{1-\alpha/2}(x)$ of the ground-truth distribution $p_0(y|x)$ (obtained from the joint distribution $p_0(\mathcal{D}, z)$). In fact, writing as

$$y_q(x) = \inf \left\{ y \in \mathbb{R} : \int_{-\infty}^y p_0(y'|x) dy' \leq q \right\} \quad (10)$$

the q -quantile, with $q \in [0, 1]$, of the ground-truth distribution $p_0(y|x)$, the interval $[y_{\alpha/2}(x), y_{1-\alpha/2}(x)]$ contains the true value y with probability $1 - \alpha$.

Defining the pinball loss as [72]

$$\ell_q(y, \hat{y}) = \max \left\{ -(1 - q)(y - \hat{y}), q(y - \hat{y}) \right\} \quad (11)$$

for $q \in [0, 1]$, the quantile $y_q(x)$ in (10) can be obtained as [73]

$$y_q(x) = \underset{\hat{y} \in \mathbb{R}}{\operatorname{argmin}} \mathbb{E}_{p_0(y|x)} [\ell_q(y, \hat{y})]. \quad (12)$$

Therefore, given a parametrized predictive model $\hat{y}(x|\phi)$, the quantile $y_q(x)$ can be estimated as $\hat{y}(x|\phi_{\mathcal{D}, q})$ with optimized parameter vector

$$\phi_{\mathcal{D}, q} = \underset{\phi}{\operatorname{argmin}} \left\{ \frac{1}{N} \sum_{(x, y) \in \mathcal{D}} \ell_q(y, \hat{y}(x|\phi)) \right\}. \quad (13)$$

With the estimate $\hat{y}(x|\phi_{\mathcal{D}, \alpha/2})$ of quantile $y_{\alpha/2}(x)$ and estimate $\hat{y}(x|\phi_{\mathcal{D}, 1-\alpha/2})$ of quantile $y_{1-\alpha/2}(x)$, we finally obtain the *naïve quantile-based* (NQB) predictor

$$\Gamma^{\text{NQB}}(x|\mathcal{D}) = \left[\hat{y}(x|\phi_{\mathcal{D}, \alpha/2}), \hat{y}(x|\phi_{\mathcal{D}, 1-\alpha/2}) \right]. \quad (14)$$

The naïve set prediction in (14) fails to satisfy the condition (7), since the empirical quantiles $\hat{y}_q(x)$ generally differ from the ground-truth quantiles $y_q(x)$.

IV. CONFORMAL PREDICTION

In this section, we review CP-based set predictors, which have the key property of guaranteeing the $(1 - \alpha)$ -validity condition (7) for any predetermined miscoverage level α , irrespective of the ground-truth distribution $p_0(\mathcal{D}, z)$ of the data. We specifically focus on validation-based CP [10] and cross-validation-based CP [53], which are more practical variants of full CP [10], [74]. In Sec. V, we cover online CP [14], [61].

A. VALIDATION-BASED CP (VB-CP)

In this subsection, we describe *validation-based* CP (VB-CP), which partitions the available set $\mathcal{D} = \mathcal{D}^{\text{tr}} \cup \mathcal{D}^{\text{val}}$ into a training set \mathcal{D}^{tr} with N^{tr} samples and a validation set \mathcal{D}^{val} with $N^{\text{val}} = N - N^{\text{tr}}$ samples (Fig. 5(a)). This class of methods is also known as inductive CP [10] or split CP [53].

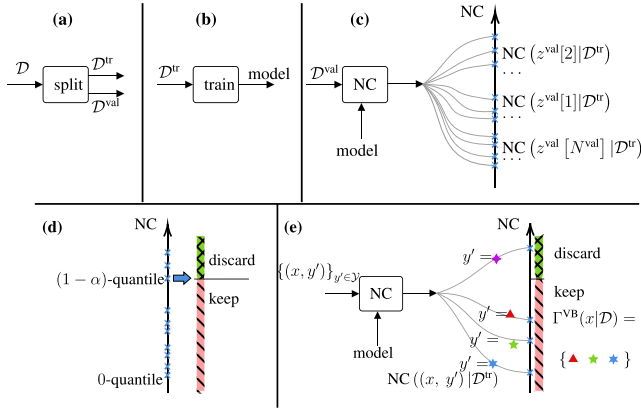


FIGURE 5. Validation-based conformal prediction (VB-CP): (a) The data set is split into training and validation set; (b) A single model is trained over the training data set; (c)-(d) Post-hoc calibration is done by evaluating the NC scores on the validation set (c) and by identifying the $(1 - \alpha)$ -quantile of the validation NC scores. This divides the axis of NC scores into a “keep” region of NC scores smaller than the threshold, and into a complementary “discard” region (d). (e) For each test input x , VB-CP includes in the prediction set all labels $y' \in \mathcal{Y}$ for which the NC score of the pair (x, y') is within the “keep” region.

VB-CP operates on any pre-trained probabilistic model $p(y|x, \mathcal{D}^{\text{tr}})$ obtained using the training set \mathcal{D}^{tr} as per (3). At test time, given an input x , VB-CP relies on a validation set to determine which labels $y' \in \mathcal{Y}$ should be included in the predicted set. Specifically, for any given test input x , a label $y' \in \mathcal{Y}$ is included in set $\Gamma^{\text{VB}}(x|\mathcal{D})$ depending on the extent to which the candidate pair (x, y') “conforms” with the examples in the validation set.

This “conformity” test for a candidate pair is based on a *nonconformity (NC) score*. An NC score for VB-CP can be obtained as the log-loss

$$\text{NC}(z = (x, y)|\mathcal{D}^{\text{tr}}) = -\log p(y|x, \mathcal{D}^{\text{tr}}) \quad (15)$$

or as any other score function that measures the loss of the probabilistic predictor $p(y|x, \mathcal{D}^{\text{tr}})$ on example (x, y) . It is also possible to define NC scores for quantile-based predictors as in (14), and we refer to [61] for details.

VB-CP consists of a training phase (Fig. 5(a)-(d)) and of a test phase (Fig. 5(e)). During *training*, the data set \mathcal{D}^{tr} is used to obtain a probabilistic predictor $p(y|x, \mathcal{D}^{\text{tr}})$ as in (3) (Fig. 5(b)). Then, NC scores $\text{NC}(z^{\text{val}}[i]|\mathcal{D}^{\text{tr}})$, as in (15), are evaluated on all points $z^{\text{val}}[i]$, $i = 1, \dots, N^{\text{val}}$ in the validation set \mathcal{D}^{val} (Fig. 5(c)). Finally, the real line of NC scores is partitioned into a “keep” region and a “discard” region (Fig. 5(d)), choosing as a threshold the $(1 - \alpha)$ -empirical quantile of the N^{val} NC scores $\{\text{NC}(z^{\text{val}}[i]|\mathcal{D}^{\text{tr}})\}_{i=1}^{N^{\text{val}}}$. Accordingly, we “keep” the labels y' with NC scores that are smaller than the $(1 - \alpha)$ -empirical quantile of the validation NC scores, and “discard” larger NC scores.

During *testing* (Fig. 5(e)), given a test input x , $|\mathcal{Y}|$ NC scores are evaluated, one for each of the candidate labels $y' \in \mathcal{Y}$, using the same trained model $p(y|x, \mathcal{D}^{\text{tr}})$. All candidate

labels y' for which the NC score $\text{NC}((x, y')|\mathcal{D}^{\text{tr}})$ falls within the “keep” region are included in the predicted set of VB-CP. Mathematically, the VB-CP set predictor is obtained as

$$\Gamma^{\text{VB}}(x|\mathcal{D}) = \left\{ y' \in \mathcal{Y} \mid \text{NC}((x, y')|\mathcal{D}^{\text{tr}}) \leq Q_{\alpha}(\{\text{NC}(z^{\text{val}}[i]|\mathcal{D}^{\text{tr}})\}_{i=1}^{N^{\text{val}}}) \right\}, \quad (16)$$

where the *empirical quantile from the top* for a set of N real values $\{r[i]\}_{i=1}^N$ is defined as

$$Q_{\alpha}(\{r[i]\}_{i=1}^N) = \lceil (1 - \alpha)(N + 1) \rceil \text{th smallest value of the set } \{r[i]\}_{i=1}^N \cup \{+\infty\}. \quad (17)$$

B. CROSS-VALIDATION-BASED CP (CV-CP)

VB-CP has the computational advantage of requiring the training of a single model, but the split into training and validation data causes the available data to be used in an inefficient way. This data inefficiency generally yields set predictors with a large average size (8). Unlike VB-CP, *cross-validation-based CP (CV-CP)* [53] trains multiple models, each using a subset of the available data set \mathcal{D} . As detailed next and summarized in Fig. 6, during the training phase, each data point $z[i]$ in the validation set is assigned an NC score based on a model trained using a subset of the data set \mathcal{D} that excludes $z[i]$, with $i \in \{1, \dots, N\}$. Then, for testing, the inclusion of a label y' in the prediction set for an input x is based on a comparison of NC scores evaluated for the pair (x, y') with all the N validation NC scores.

Specifically, as illustrated in Fig. 6, K -fold CV-CP [53], referred here as K -CV-CP, first partitions the data set \mathcal{D} into K disjoint folds $\{\mathcal{S}_k\}_{k=1}^K$, each with N/K points, i.e., $\cup_{k=1}^K \mathcal{S}_k = \mathcal{D}$ (Fig. 6(a)), for a predefined integer $K \in \{2, \dots, N\}$ such that the ratio N/K is an integer.

During *training*, the K subsets $\mathcal{D} \setminus \mathcal{S}_k$ are used to train K probabilistic predictors $p(y|x, \mathcal{D} \setminus \mathcal{S}_k)$ defined as in (3) (Fig. 6(b)). Each trained model $p(y|x, \mathcal{D} \setminus \mathcal{S}_k)$ is used to evaluate the $|\mathcal{S}_k| = N/K$ NC scores $\text{NC}(z_k|\mathcal{D} \setminus \mathcal{S}_k)$ for all validation data points $z_k \in \mathcal{S}_k$ that were not used for training the model (Fig. 6(c)). Unlike VB-CP, K -CV-CP requires keeping in memory all the N validation scores for testing. These points are illustrated as crosses in Fig. 6(c).

During *testing*, for a given test input x and for any candidate label $y' \in \mathcal{Y}$, CV-CP evaluates K NC scores, one for each of the K trained models. Each such NC score $\text{NC}((x, y')|\mathcal{D} \setminus \mathcal{S}_k)$ is compared with the N/K validation scores obtained on fold \mathcal{S}_k . We then count how many of the N/K validation scores are larger than $\text{NC}((x, y')|\mathcal{D} \setminus \mathcal{S}_k)$. If the sum of all such counts, across the K folds $\{\mathcal{S}_k\}_{k=1}^K$, is larger than a fraction α of all N data points, then the candidate label y' is included in the prediction set (Fig. 6(d)). This criterion follows the same principle of VB-CP of including all candidate labels y' that “conform” well with a sufficiently large fraction of validation points.

Mathematically, K -CV-CP is defined as

$$\Gamma^{K\text{-CV}}(x|\mathcal{D}) = \left\{ y' \in \mathcal{Y} \mid \sum_{k=1}^K \sum_{z_k \in \mathcal{S}_k} \mathbb{1}(\text{NC}((x, y')|\mathcal{D} \setminus \mathcal{S}_k)) \leq \text{NC}(z_k|\mathcal{D} \setminus \mathcal{S}_k) \geq \lfloor \alpha(N+1) \rfloor \right\}, \quad (18)$$

where $\mathbb{1}(\cdot)$ is the indicator function ($\mathbb{1}(\text{true}) = 1$ and $\mathbb{1}(\text{false}) = 0$). The left-hand side of the inequality in (18) implements the sums, shown in Fig. 6(d), over counts of validation NC scores that are larger than the corresponding NC score for the candidate pair (x, y') .

K -CV-CP increases the computational complexity K -fold as compared to VB-CP, while generally reducing the inefficiency [53]. The special case of $K = N$, known as jackknife+ [53], is referred here as CV-CP. In this case, each of the N folds $\mathcal{S}_k, k = 1, \dots, N$ uses a single cross validation point. In general, CV-CP is the most efficient form of K -CV-CP, but it may be impractical for large data set sizes due to need to train N models. The number of folds K should strike a balance between computational complexity, as K models are trained, and inefficiency.

C. CALIBRATION GUARANTEES

VB-CP (16) satisfies the coverage condition (7) [10] under the only assumption of exchangeability (see Sec. II-A).

The validity of CV-CP requires a technical assumption on the NC score. While in VB-CP the NC score is an arbitrary score function evaluated based on any pre-trained probabilistic model, for CV-CP, the NC score $\text{NC}(z|\mathcal{D}')$ must satisfy the additional property of being invariant to permutations of the data set \mathcal{D}' used to train the underlying probabilistic model. Consider the log-loss $\text{NC}(z = (x, y)|\mathcal{D}') = -\log p(y|x, \mathcal{D}')$ (15), or any other score function based on the trained model $p(y|x, \mathcal{D}')$, as the NC score. CV-CP requires that the training algorithm used to produce model $p(y|x, \mathcal{D}')$ provides outputs that are invariant to permutations of the training set \mathcal{D}' .

Specifically, for frequentist learning, the optimization algorithm producing the parameter vector $\phi_{\mathcal{D}'}^*$ in (3) must be permutation-invariant. This is the case for standard methods such as full-batch gradient descent (GD), or for non-parametric techniques such as Gaussian processes. For Bayesian learning, the distribution $q^*(\phi|\mathcal{D}')$ in (3) must also be permutation-invariant, which is true for the exact posterior distribution [23], as well as for approximations obtained via MC methods such as Langevin MC [23], [31].

The requirement on permutation-invariance can be alleviated by allowing for probabilistic training algorithms such as stochastic gradient descent (SGD) [75]. With probabilistic training algorithms, the only requirement is that the *distribution* of the (random) output models is permutation-invariant. This is, for instance, the case if SGD is implemented by taking mini-batches uniformly at random within the training set \mathcal{D}' [75], [76], [77]. With probabilistic training algorithms, however, the validity condition (7) of CV-CP is only guar-

anteed on average with respect to the random outputs of the algorithms.

Specifically, under the discussed assumption of permutation-invariance of the NC scores, by [53, Theorems 1 and 4], CV-CP satisfies the inequality

$$\mathbb{P}(y \in \Gamma^{\text{CV}}(x|\mathcal{D})) \geq 1 - 2\alpha, \quad (19)$$

while K -CV-CP satisfies the inequality

$$\begin{aligned} \mathbb{P}(y \in \Gamma^{K\text{-CV}}(x|\mathcal{D})) &\geq 1 - 2\alpha - \min \left\{ \frac{2(1-1/K)}{N/K+1}, \frac{1-K/N}{K+1} \right\} \\ &\geq 1 - 2\alpha - \sqrt{2/N}. \end{aligned} \quad (20)$$

Therefore, validity for both cross-validation schemes is guaranteed for the larger miscoverage level of 2α . Accordingly, one can achieve miscoverage level of α , satisfying (7), by considering the CV-CP set predictor $\Gamma^{\text{CV}}(x|\mathcal{D})$ with $\alpha/2$ in lieu of α in (18). That said, in the experiments, we will follow the recommendation in [53] and [76] to use α in (18).

V. ONLINE CONFORMAL PREDICTION

In this section, we turn to online CP. Unlike the CP schemes presented in the previous section, online CP makes no assumptions about the probabilistic model underlying data generation [14], [61]. Rather, it models the observations as a deterministic stream of input-output pairs $z[i] = (x[i], y[i])$ over time index $i = 1, 2, \dots$; and it targets a coverage condition defined in terms of the empirical rate at which the prediction set Γ_i at time i covers the correct output $y[i]$.

In the offline version of CP reviewed in the previous section, all N samples of the data set \mathcal{D} are assumed to be available upfront (see Fig. 3(a)). In contrast, in online CP, a set predictor Γ_i for time index i is produced for each new input $x[i]$ over time $i = 1, 2, \dots$. Specifically, given the past observations $\{z[j]\}_{j=1}^{i-1}$, the set predictor $\Gamma_i(x[i]|\{z[j]\}_{j=1}^{i-1})$ outputs a subset of the output space \mathcal{Y} . Given a target miscoverage level $\alpha \in [0, 1]$, an online set predictor is said to be $(1 - \alpha)$ -long-term valid if the following limit holds

$$\lim_{I \rightarrow \infty} \frac{1}{I} \sum_{i=1}^I \mathbb{1}(y[i] \in \Gamma_i(x[i]|\{z[j]\}_{j=1}^{i-1})) = 1 - \alpha \quad (21)$$

for all possible sequences $z[i]$ with $i = 1, 2, \dots$. Note that the condition (21), unlike (7), does not involve any ensemble averaging with respect to the data distribution. We will take (21) as the relevant definition of calibration for online learning.

Rolling conformal inference (RCI) [61] adapts in an online fashion a *calibration parameter* $\theta[i]$ across the time index i as a function of the instantaneous error variable

$$\text{err}[i] = \mathbb{1}(y[i] \notin \Gamma_i(x[i])), \quad (22)$$

which equals 1 if the correct output value is not included in the prediction set $\Gamma_i(x[i])$, and 0 otherwise. This is done using the update rule

$$\theta[i+1] \leftarrow \theta[i] + \gamma(\text{err}[i] - \alpha), \quad (23)$$

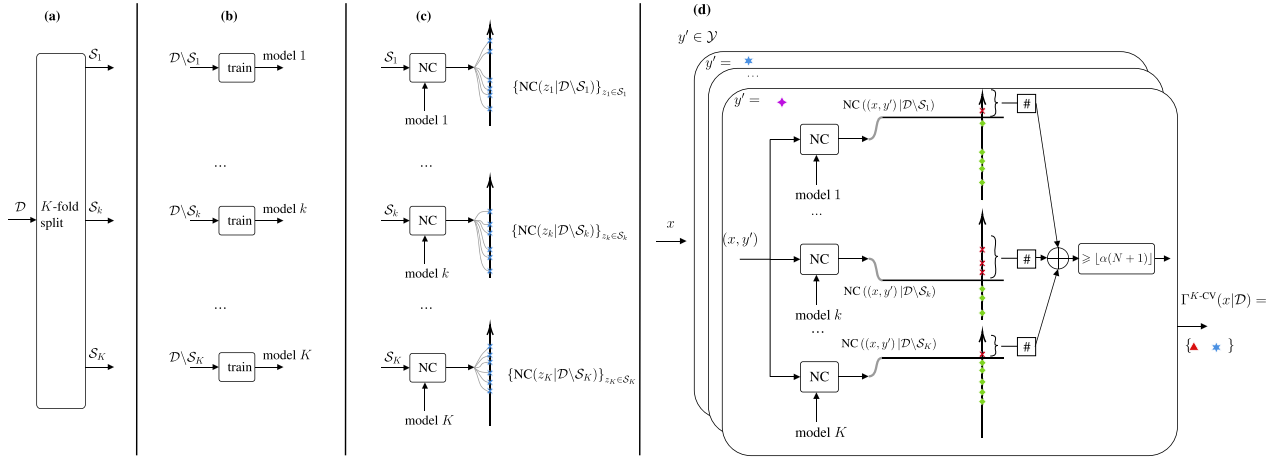


FIGURE 6. K -fold cross-validation-based conformal prediction (K -CV-CP): (a) The N data pairs of data set \mathcal{D} are split into K -folds each with $|\mathcal{S}_k| = N/K$ samples; (b) K models are trained, each using a leave-fold-out data set of $|\mathcal{D} \setminus \mathcal{S}_k| = N - N/K$ pairs; (c) NC scores are computed on the N/K holdout data points for each fold \mathcal{S}_k ; (d) For each test input x , all labels $y' \in \mathcal{Y}$ for which the number of “higher-NC” validation points exceeds a fraction α of the total N points are considered in the prediction set. CV-CP is the special case with $K = N$.

where $\gamma > 0$ is a learning rate. Accordingly, the parameter θ is increased by $\gamma(1 - \alpha)$ if an error occurs at time i , and is decreased by $\gamma\alpha$ otherwise. Intuitively, a large positive parameter $\theta[i]$ indicates that the set predictor should be more inclusive in order to meet the validity constraint (21); and vice versa, a large negative value of $\theta[i]$ suggests that the set predictor can reduce the size of the prediction sets without affecting the long-term validity constraint (21).

Following [61], we elaborate on the use of the calibration parameter $\theta[i]$ in order to ensure condition (21) for an online version of the naïve quantile-based predictor (14) for scalar regression. A similar approach applies more broadly (see [14], [78], and [79]). Denote the data set $\mathcal{D}[i] = \{z[j]\}_{j=1}^{i-1}$ as having all previously observed labeled data set up till time $i-1$. The key idea behind RCI is to extend the naïve prediction interval (14) depending on the calibration parameter $\theta[i]$ as

$$\Gamma_i^{\text{RCI}}(x[i]|\mathcal{D}[i]) = \left[\hat{y}(x|\phi_{\mathcal{D}[i], \alpha/2}) - \varphi(\theta[i]), \hat{y}(x|\phi_{\mathcal{D}[i], 1-\alpha/2}) + \varphi(\theta[i]) \right], \quad (24)$$

where

$$\varphi(\theta) = \text{sign}(\theta)(\exp^{|\theta|} - 1) \quad (25)$$

is the so-called stretching function, a fixed monotonically increasing mapping.

The set predictor RCI (24) “corrects” the NQB set predictor (14), via the additive stretching function $\varphi(\theta[i])$ based on the calibration parameter $\theta[i]$. As the time index i rolls, the calibration parameter $\theta[i]$ adaptively inflates and deflates according to (23). Upon each observation of new label $y[i]$, the quantile predictor model parameters $\phi_{\mathcal{D}[i], \alpha/2}$ and $\phi_{\mathcal{D}[i], 1-\alpha/2}$ can also be updated, without affecting the long-term validity condition (21) [61, Theorem 1]. We refer to Appendix IX for further details on online CP.

VI. SYMBOL DEMODULATION

In this section, we focus on the application of offline CP, as described in Sec. IV, to the problem of symbol demodulation in the presence of transmitter hardware imperfections. This problem was also considered in [20] and [80] by focusing on frequentist and Bayesian learning. Unlike [20], [80], we investigate the use of CP as a means to obtain set predictors satisfying the validity condition (7).

A. PROBLEM FORMULATION

The problem of interest consists of the demodulation of symbols from a discrete constellation based on received baseband signals subject to hardware imperfections, noise, and fading. The goal is to design *set demodulators* that output a subset of all possible constellation points with the guarantee that the subset includes the true transmitted signal with the desired target probability $1 - \alpha$. Set-valued demodulation can be useful when combined with channel decoders that can concurrently explore several paths [81]. Furthermore, while we do not investigate this direction here, CP could be applied directly to decoding, yielding a form of list decoding (see, e.g., [82]) with formal reliability guarantees.

To keep the notation consistent with the previous sections, we write as $y[i]$ the i -th transmitted symbols, and as $x[i]$ the corresponding received signal. Each transmitted symbol $y[i]$ is drawn uniformly at random from a given constellation \mathcal{Y} . We model I/Q imbalance at the transmitter and phase fading as in [67]. Accordingly, the ground-truth channel law connecting symbols $y[i]$ into received samples $x[i]$ is described by the equality

$$x[i] = e^{j\psi} f_{\text{IQ}}(y[i]) + v[i], \quad (26)$$

for a random phase $\psi \sim \text{U}[0, 2\pi)$, where the additive noise is $v[i] \sim \mathcal{CN}(0, \text{SNR}^{-1})$ for signal-to-noise ratio level SNR.

Furthermore, the I/Q imbalance function [83] is defined as

$$f_{IQ}(y[i]) = \bar{y}_I[i] + j\bar{y}_Q[i], \quad (27)$$

where

$$\begin{bmatrix} \bar{y}_I[i] \\ \bar{y}_Q[i] \end{bmatrix} = \begin{bmatrix} 1 + \epsilon & 0 \\ 0 & 1 - \epsilon \end{bmatrix} \begin{bmatrix} \cos \delta & -\sin \delta \\ -\sin \delta & \cos \delta \end{bmatrix} \begin{bmatrix} y_I[i] \\ y_Q[i] \end{bmatrix}, \quad (28)$$

with $y_I[i]$ and $y_Q[i]$ being the real and imaginary parts of the modulated symbol $y[i]$; and $\bar{y}_I[i]$ and $\bar{y}_Q[i]$ standing for the real and imaginary parts of the transmitted symbol $f_{IQ}(y[i])$. In (28), the channel state c consists of the tuple $c = (\psi, \epsilon, \delta)$ encompassing the complex phase ψ and the I/Q imbalance parameters (ϵ, δ) .

B. IMPLEMENTATION

As in [20] and [80], demodulation is implemented via a neural network probabilistic model $p(y|x, \phi)$ consisting of a fully connected network with real inputs $x[i]$ of dimension 2 as per (26), followed by three hidden layers with 10, 30, and 30 neurons having ReLU activations in each layer. The last layer implements a softmax classification for the $|\mathcal{Y}|$ possible constellation points.

We adopt the standard NC score (15), where the trained model $\phi_{\mathcal{D}}$ for frequentist learning is obtained via $I = 120$ GD update steps for the minimization of the cross-entropy training loss with learning rate $\eta = 0.2$; while for Bayesian learning we implement a gradient-based MC method, namely Langevin MC, with burn-in period of $R_{\min} = 100$, ensemble size $R = 20$, learning rate $\eta = 0.2$, and temperature parameter $T = 20$. We assume standard Gaussian distribution for the prior distribution [31]. Details on Langevin MC can be found in Appendix IX.

We compare the naïve set predictor (9), also studied in [20] and [80], which provides no formal coverage guarantees, with the CP set prediction methods reviewed in Sec. IV. VB-CP uses equal set sizes for the training and validation sets. We target the miscoverage level as $\alpha = 0.1$.

C. RESULTS

We consider the Amplitude-Phase-Shift-Keying (APSK) modulation with $|\mathcal{Y}| = 8$. The SNR level is set to $\text{SNR} = 5$ dB. The amplitude and phase imbalance parameters are independent and distributed as $\epsilon \sim \text{Beta}(\epsilon/0.15|5, 2)$ and $\delta \sim \text{Beta}(\delta/15^\circ|5, 2)$, respectively [80].

Fig. 7 shows the

$$\text{empirical coverage} = \frac{1}{N^{\text{te}}} \sum_{j=1}^{N^{\text{te}}} \mathbb{1}(y^{\text{te}}[j] \in \Gamma(x^{\text{te}}[j]|\mathcal{D})), \quad (29)$$

and Fig. 8 shows the

$$\text{empirical inefficiency} = \frac{1}{N^{\text{te}}} \sum_{j=1}^{N^{\text{te}}} |\Gamma(x^{\text{te}}[j]|\mathcal{D})|, \quad (30)$$

both evaluated on a test set $\mathcal{D}^{\text{te}} = \{(x^{\text{te}}[j], y^{\text{te}}[j])\}_{j=1}^{N^{\text{te}}}$ with $N^{\text{te}} = 100$ data points, as a function of the size of the




FIGURE 7. Coverage for naïve set predictor (9), VB-CP (16), CV-CP, and K-CV-CP (18) with $K = 4$, for symbol demodulation problem (Section VI). For every set predictors, the NC scores are evaluated either using frequentist learning (dashed lines) or Bayesian learning (solid lines). The coverage level is set to $1 - \alpha = 0.9$, and each numerical evaluation is averaged over 50 independent trials (new channel state c) with $N^{\text{te}} = 100$ test points.




FIGURE 8. Average set prediction size (inefficiency) for the same setting of Fig. 7.

available data set \mathcal{D} . We average the results for 50 independent trials, each corresponding to independent draws of the variables $\{\mathcal{D}, \mathcal{D}^{\text{te}}\}$ from the ground truth distribution. This way, the metrics (29)-(30) provide an estimate of the coverage (7) and of the inefficiency (8), respectively [53].

From Fig. 7, we first observe that the naïve set predictor, with both frequentist and Bayesian learning, does not meet the desired coverage level in the regime of a small number N of available samples. In contrast, confirming the theoretical calibration guarantees presented in Sec. IV, all CP methods provide coverage guarantees, achieving coverage rates above $1 - \alpha$. Furthermore, as seen in Fig. 8, coverage guarantees

are achieved by suitably increasing the size of prediction sets, which is reflected by the larger inefficiency. The size of the prediction sets, and hence the inefficiency, decreases as the data set size, N , increases. In this regard, due to their more efficient use of the available data, CV-CP and K -CV-CP predictors have a lower inefficiency as compared to VB predictors, with CV-CP offering the best performance. Finally, Bayesian NC scores are generally seen to yield set predictors with lower inefficiency, confirming the merits of Bayesian learning in terms of calibration.

VII. MODULATION CLASSIFICATION

In this section, we propose and evaluate the application of offline CP to the problem of modulation classification [45], [46].

A. PROBLEM FORMULATION

Due to the scarcity of frequency bands, electromagnetic spectrum sharing among licensed and unlicensed users is of special interest to improve the efficiency of spectrum utilization. In sensing-based spectrum sharing, a transmitter scans the prospective frequency bands to identify, for each band, if the spectrum is occupied, and, if so, if the signal is from a licensed user or not. A key enabler for this operation is the ability to classify the modulation of the received signal [84]. The modulation classification task is made challenging by the dimensionality of the baseband input signal and by the distortions caused by the propagation channel. Data-driven solutions [85] have shown to be effective for this problem in terms of accuracy, while the focus here is on calibration performance.

Accordingly, we aim at designing *set modulation classifiers* that output a subset of the set of all possible modulation schemes with the property that the true modulation scheme is contained in the subset with a desired probability level $1 - \alpha$. Set predictions provide actionable information in several applications of modulation classifiers. For instance, in cognitive radio systems, there may be protected modulation schemes adopted by a primary user [86], [87]. If the predicted set produced by a secondary user includes such schemes, the secondary user should, e.g., refrain from transmitting or use lower transmission powers, guaranteeing that the primary user is not affected with probability at least $1 - \alpha$. As another example, in military or emergency response applications, modulation classification is used a preliminary step to synchronize a receiver [88], [89]. In this case, having a list of possible modulation schemes enables the receiver to attempt decoding in parallel on all possible identified schemes, with the guarantee that the correct modulation scheme is evaluated with probability at least $1 - \alpha$.

We adopt the data set provided by [46], which has approximately 2.5×10^6 baseband signals, each produced using one out of 24 possible digital and analog modulations across different SNR values and channel models. The data set contains data for single-carrier modulations such as OOK, ASK, PSK, APSK, QAM, AM, FM, GMSK, and OQPSK,

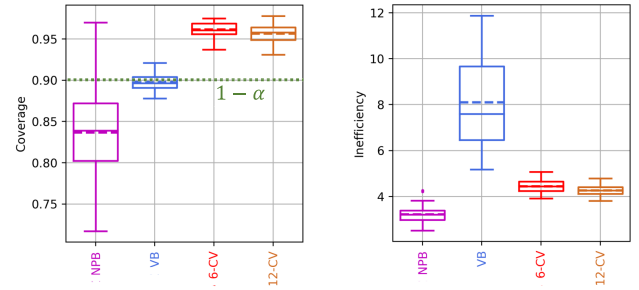


FIGURE 9. Coverage and inefficiency for NPB (9), VB-CP (16), and K -CV-CP (18) with $K = 6$ and $K = 12$, for the modulation classification problem (implementation details in Section VII-B). The boxes represent the 25% (lower edge), 50% (solid line within the box), and 75% (upper edge) percentiles of the empirical performance metrics evaluated over 32 different experiments, with average value shown by the dashed line.

including several modulation orders for each scheme. For each modulation, several short-time signal observations of 1024 I/Q samples are synthesized using random realizations of Rayleigh fading, carrier frequency mismatch, sampling rate mismatch, and shaping roll-off values. Out of the whole data set, in this paper we focus on the high SNR regime (≥ 6 dB). Accordingly, the data set \mathcal{D} consists of approximately 1.28×10^6 pairs (x, y) , where x is the channel output signal of 2048 interleaved I/Q samples and y is the index of one of the $|\mathcal{Y}| = 24$ possible modulations. The SNR value itself is not available to the classifier.

B. IMPLEMENTATION

We use a neural network architecture similar to the one used in [46], which has 7 one-dimensional convolutional layers with kernel size 3 and 64 channels for all layers, except for the first layer with has 2 channels. The convolution layers are followed by 3 fully-connected linear layers. A scaled exponential linear unit (SELU) is used for all inner layers, and a softmax is used at the last, fully connected, layer. We assume availability of $N = 4800$ pairs (x, y) for the data set \mathcal{D} , while gauging the empirical inefficiency and coverage level with $N^{\text{te}} = 1000$ held-out pairs. A total number of $I = 4000$ GD steps with fixed learning rate of 0.02 are carried out, and the target miscoverage rate is set to $\alpha = 0.1$. VB partitions its available data into equal sets for training and validation.

C. RESULTS

In this problem, due to computational cost, we exclude CV-CP and we focus on K -CV-CP with a moderate number of folds, namely $K = 6$ and $K = 12$. In Fig. 9, box plots show the quartiles of the empirical coverage (29) and of the empirical inefficiency (30) from 32 independent runs, with different realizations of data set and test examples. The lower edge of the box represents the 0.25-quantile; the solid line within the box the median; the dashed line within the box the average; and the upper edge of the box the 0.75-quantile.


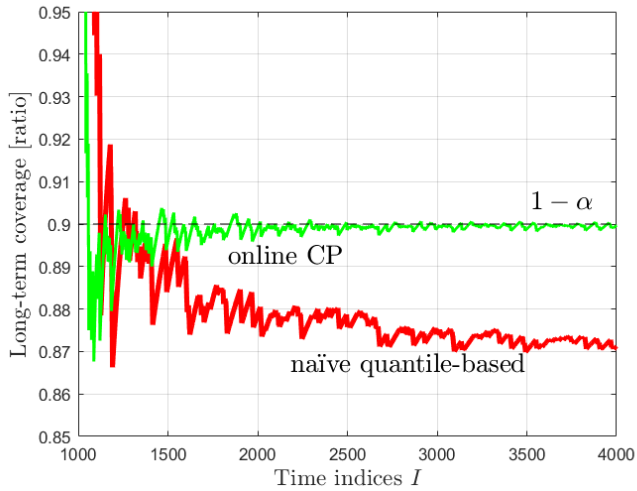


FIGURE 10. CP for time-series `Outdoor` of [69]: (top) Time-average coverage (32) of naive set prediction and online CP; (bottom) Time-averaged inefficiency (33) of naive set prediction and online CP.

As can be seen in the figure, the naive set predictor is invalid (see average shown as dashed line), and it exhibits a wide spread of the coverage rates across the trials. On the other hand, all CP set predictors are valid, meeting the predetermined coverage level $1 - \alpha = 0.9$, and have less spread-out coverage rates.

As also noted in the previous section, VB-CP suffers from larger predicted set size as compared to K -CV-CP, due to poor sample efficiency. A small number of folds, as low as $K = 6$, is sufficient for K -CV-CP to outperform VB-CP. This improvement in efficiency comes at the computational cost of training six models, as compared to the single model trained by VB-CP.

VIII. ONLINE CHANNEL PREDICTION

In this section, we investigate the use of online CP, as described in Sec. V, for the problem of channel prediction. We specifically focus on the prediction of the received signal

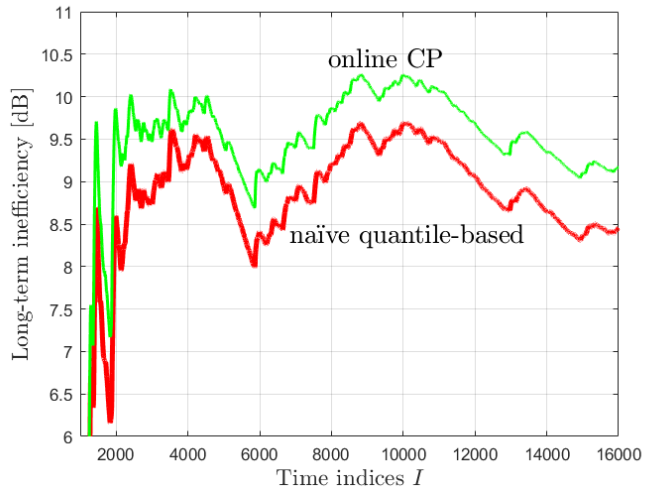
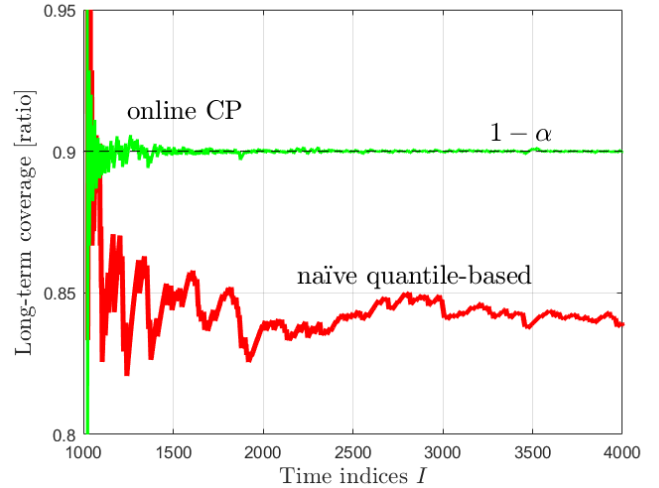


FIGURE 11. CP for time-series `NLOS_Head_Indoor_1kHz` in [68]: (top) Time-average coverage (32) of naive set prediction and online CP; (bottom) Time-average inefficiency (33) of naive set prediction and online CP.

strength (RSS), which is a key primitive at the physical layer, supporting important functionalities such as resource allocation [90], [91].

A. PROBLEM FORMULATION

Consider a receiver that has access to a sequence of RSS samples from a given device. We aim at designing a predictor that, given a sequence of past samples from the RSS sequence, produces an *interval* of values for the next RSS sample. To meet calibration requirements, the interval must contain the correct future RSS value with the desired rate level $1 - \alpha$. Unlike the previous applications, here the rate of coverage is evaluated based on the time average

$$\frac{1}{t} \sum_{i=1}^t \mathbb{1}\left(y[i] \in \Gamma_i(x[i] | \{z[j]\}_{j=1}^{i-1})\right). \quad (31)$$

This is computed as the fraction previous time instants $i \in \{1, \dots, t\}$ at which the set predictor Γ_i includes the true RSS value $y[i]$.

Interval predictions for communication channels can support resource allocation. For instance, a scheduler may select users whose lowest possible RSS in the predicted set is sufficiently large to enable communication at some desired level [92]. A scheduler may also use larger values in the predicted set to facilitate exploration via optimism [35], [93]. Finally, set prediction can enable the detection of outage events, or anomalies, by checking if the lowest allowed RSS is not in line with the past observations for a given link.

We consider two data sets of RSS sequences. The first data set records RSS samples $y[i]$ in logarithmic scale for an IEEE 802.15.4 radio over time index i [69]. We further use the available side information on the time-variant channel ID, which determines the carrier frequency used at time i out of the 16 possible bands, as the input $x[i]$. At time i , we observe a sequence of RSS samples $z[1], \dots, z[i-1]$ with $z[i] = (x[i], y[i])$, and the goal is to predict the next RSS sample $y[i]$ via the online set predictor Γ_i^{RCI} (24).

The second data set [68] reports samples $y[i]$, measured in dBm, on a 5.8 GHz device-to-device link without additional input. Hence, in this case, we predict the next RSS sample $y[i]$ using the previous RSS samples $y[1], \dots, y[i-1]$. Note that the prior works [68], [69] adopted standard probabilistic predictors, while here we focus on set predictors that produce a prediction interval $\Gamma_i^{\text{RCI}}(x[i]|\{z[j]\}_{j=1}^{i-1})$.

B. IMPLEMENTATION

We build the CP set predictor by leveraging the probabilistic neural network used in [61] as the model class for the quantile predictors in (13)-(14). Each quantile predictor consists of a multi-layer neural network that pre-processes the most recent K pairs $\{z[i-K], \dots, z[i-1]\}$; of a stacked long short-term memory (LSTM) [94] with two layers; and of a post-processing neural network, which maps the last LSTM hidden vector into a scalar that estimates the quantile used in (14). For details of the implementation, we refer to Appendix 13.

C. RESULTS

Fig. 10 and Fig. 11 report the

$$\text{time-average coverage} = \frac{1}{T} \sum_{i=1}^T \mathbb{1}(y[i] \in \Gamma_i(x[i]|\{z[j]\}_{j=1}^{i-1})) \quad (32)$$

and the

$$\text{time-average inefficiency} = \frac{1}{T} \sum_{i=1}^T \left| \Gamma_i(x[i]|\{z[j]\}_{j=1}^{i-1}) \right| \quad (33)$$

for online CP (24), compared to a baseline of the naïve quantile-based predictor (14), as a function of the time window size I for data sets [68] and [68], respectively. We have

discarded 1000 samples for a warm-up period for both metrics (32) and (33).

In both cases, the naïve predictor is seen to fail to satisfy the coverage condition (21) for both data sets, while online CP converges to the target level $1 - \alpha = 0.9$. This result is obtained by online CP with a modest increase of around 8% for both data sets in terms of inefficiency.

IX. CONCLUSION

AI in communication engineering should not only target accuracy, but also calibration, ensuring a reliable and safe adoption of machine learning within the overall telecommunication ecosystem. In this paper, we have proposed the adoption of a general framework, known as conformal prediction (CP), to transform any existing AI model into a well-calibrated model via post-hoc calibration for communication engineering. Depending on the situation of interest, post-hoc calibration leverages either an held-out (cross) validation set or previous samples. Unlike calibration approaches that do not formally guarantee reliability, such as Bayesian learning or temperature scaling, CP provides formal guarantees of calibration, defined either in terms of ensemble averages or long-term time averages. Calibration is retained irrespective of the accuracy of the trained models, with more accurate models producing smaller set predictions.

To validate the reliability of CP-based set predictors, we have provided extensive comparisons with conventional methods based on Bayesian or frequentist learning. Focusing on demodulation, modulation classification, and channel prediction, we have demonstrated that AI models calibrated by CP provide formal guarantees of reliability, which are practically essential to ensure calibration in the regime of limited data availability.

Future work may consider applications of CP to other use cases in wireless communication systems, as well as extension of involving training-based calibration [55], [58] and/or meta-learning [95].

APPENDIX A FREQUENTIST AND BAYESIAN LEARNING

Given access to training data set $\mathcal{D} = \{z[i] = (x[i], y[i])\}_{i=1}^N$ with N examples, frequentist learning finds a model parameter vector $\phi_{\mathcal{D}}^*$ by tackling the following empirical risk minimization (ERM) problem

$$\begin{aligned} \min_{\phi} \left\{ L_{\mathcal{D}}(\phi) &= -\frac{1}{N} \sum_{(x,y) \in \mathcal{D}} \log p(y|x, \phi) \right. \\ &= \mathbb{E}_{p_{\mathcal{D}}(x,y)} \left[-\log p(y|x, \phi) \right] \left. \right\}, \quad (34) \end{aligned}$$

with empirical distribution $p_{\mathcal{D}}(x, y)$ defined by the data set \mathcal{D} .

Bayesian learning addresses epistemic uncertainty by treating the model parameter vector as a random vector ϕ with prior distribution $\phi \sim p(\phi)$. Ideally, Bayesian learning updates the prior $p(\phi)$ to produce the posterior distribution

$p(\phi|\mathcal{D})$ as

$$p(\phi|\mathcal{D}) \propto p(\phi) \prod_{i=1}^N p(y[i]|x[i], \phi) \quad (35)$$

and obtains the ensemble predictor for the test point (x, y) by averaging over multiple models, i.e.,

$$p(y|x, \mathcal{D}) = \mathbb{E}_{p(\phi|\mathcal{D})}[p(y|x, \phi)]. \quad (36)$$

In practice, as the true posterior distribution is generally intractable due to the normalizing factor in (35), approximate Bayesian approaches are considered via VI or MC techniques (see, e.g., [23]).

In the experiments, we adopted Langevin MC to approximate the Bayesian posterior [23], [31]. Langevin MC adds Gaussian noise to each standard GD update for frequentist learning (see, e.g., [23, Sec. 4.10]). The noise has power $2\eta/T$, where η is the GD learning rate and $T > 0$ is a temperature parameter. Langevin MC produces R model parameters $\{\phi[r]\}_{r=1}^R$ across R consecutive iterations. We specifically retain only the last R samples, discarding an initial burn-in period of R_{\min} iterations. The temperature parameter T is typically chosen to be larger than 1 [96], [97]. With the R samples, the expectation term in (36) is approximated as the empirical average $\frac{1}{R} \sum_{r=1}^R p(y|x, \phi[r])$.

We observe that Langevin MC is a probabilistic training algorithm, and that it satisfies the permutation-invariance property in terms of the distribution of the random output models discussed in Sec. IV-C.

APPENDIX B ALGORITHMIC DETAILS FOR ROLLING CONFORMAL INFERENCE

The RCI algorithm is reproduced from [61] in Algorithm 1.

APPENDIX C IMPLEMENTATION OF ONLINE CHANNEL PREDICTION

The architecture of the set predictor is inspired by [61], and made out of three artificial neural networks. The first, $f_{\text{pre}}(\cdot)$, is a multi-layer perceptron (MLP) network with hidden layers of 16, 32 neurons each, parametrized by vector $\phi_{\text{pre}}[i]$. It is meant to apply a pre-process over the most recent observed $K = 20$ pairs $\{z[i-K], \dots, z[i-1]\}$ to be transformed element-wise into a length- K vector $w[i] = [w_1[i], \dots, w_K[i]]^\top$, in which the k -th element ($k = 1, \dots, K$) is $w_k[i] = f_{\text{pre}}(z[i-K+k-1]|\phi_{\text{pre}}[i])$. Effectively, this will serve as a temporal sliding K -length window, with a time-evolving pre-processing function. The second neural network, $f_{\text{LSTM}}(\cdot)$ has two layers with model parameter vectors $\phi_{\text{LSTM}}^1[i]$ (first layer) and $\phi_{\text{LSTM}}^2[i]$ (second layer), which retains a memory via the hidden state vectors h and c , initialized at every time index i as $c_0^1[i] = c_0^2[i] = h_0^1[i] = h_0^2[i] = 0$. By accessing the previous K pairs via the vector $w[i]$, this recurrent neural network extracts temporal patterns by sequentially transferring information via LSTM cells (with shared parameter vectors) in the image of hidden and cell state vectors $c_k[i], h_k[i]$ via the LSTM

Algorithm 1 Rolling Conformal Inference (for Regression) [61]

Inputs : $\alpha =$ long-term target miscoverage level
 $\theta[1] =$ initial calibration parameter
 $\phi^{\text{lo}}[1], \phi^{\text{hi}}[1] =$ initial models

Parameters : $I =$ number of online iterations
 $\gamma =$ learning rate for calibration parameter
 $\eta =$ learning rate for model updates

Output : $\{\Gamma_i^{\text{RCI}}(x[i]|\{z[j]\}_{j=1}^{i-1})\}_{i=1}^I =$ predicted sets for $\{x[i]\}_{i=1}^I$

1 **for** $i = 1, \dots, I$ time instants **do**
2 Retrieve a new data sample $(x[i], y[i])$
3 \triangleleft Set prediction of new input \triangleright
4 Calculate set $\Gamma_i^{\text{RCI}}(x[i]|\{z[j]\}_{j=1}^{i-1})$ using (24)
5 $\left[\hat{y}(x[i]|\phi^{\text{lo}}[i]) - \varphi(\theta[i]), \hat{y}(x[i]|\phi^{\text{hi}}[i]) + \varphi(\theta[i]) \right]$
6 \triangleleft Check if prediction is unsuccessful \triangleright
7 $\text{err}[i] \leftarrow \mathbb{1}(y[i] \notin \Gamma_i^{\text{RCI}}(x[i]|\{z[j]\}_{j=1}^{i-1}))$
8 \triangleleft Update calibration parameter \triangleright
9 $\theta[i+1] \leftarrow \theta[i] + \gamma(\text{err}[i] - \alpha)$
10 \triangleleft Update models using new sample \triangleright
11 $\phi^{\text{lo}}[i+1] \leftarrow \phi^{\text{lo}}[i] - \eta \nabla_{\phi} \ell_{\alpha/2}(y[i], \hat{y}(x[i]|\phi^{\text{lo}}[i]))$
12 $\phi^{\text{hi}}[i+1] \leftarrow \phi^{\text{hi}}[i] - \eta \nabla_{\phi} \ell_{1-\alpha/2}(y[i], \hat{y}(x[i]|\phi^{\text{hi}}[i]))$

13 **return** predicted sets $\{\Gamma_i^{\text{RCI}}(x[i]|\{z[j]\}_{j=1}^{i-1})\}_{i=1}^I$

cells. These vectors flow along the LSTM by concatenating $k = 1, \dots, K$ cells, and forming vectors of length 32 each

$$[c_k^1[i], h_k^1[i]] = f_{\text{LSTM}}(w_k[i], c_{k-1}^1[i], h_{k-1}^1[i]|\phi_{\text{LSTM}}^1[i]) \quad (37)$$

$$[c_k^2[i], h_k^2[i]] = f_{\text{LSTM}}(h_k^1[i], c_{k-1}^2[i], h_{k-1}^2[i]|\phi_{\text{LSTM}}^2[i]). \quad (38)$$

The third and last network is a post-processing MLP $f_{\text{post}}(\cdot)$ with one hidden layer of 32 neurons, and with parameter vector $\phi_{\text{post}}[i]$, which maps the last LSTM hidden 64-length vector $h_K[i] = [h_K^1[i], h_K^2[i]]^\top$ into a scalar $f_{\text{post}}(x[i], h_K[i]|\phi_{\text{post}}[i]) \in \mathbb{R}$ that estimates the quantile for the output $y[i]$. Accordingly, the time evolving model parameter is the tuple

$$\phi[i] = (\phi_{\text{pre}}[i], \phi_{\text{LSTM}}^1[i], \phi_{\text{LSTM}}^2[i], \phi_{\text{post}}[i]). \quad (39)$$

This model is instantiated twice for the regression problem: one for the $\alpha/2$ lower quantile and the other for $1-\alpha/2$ upper quantile. For every time instant i , after the new output $y[i]$ is observed, continual learning of the models is taken place by training the models with corresponding pinball losses (11) using the new pair $(x[i], y[i])$, while initializing the models as the previous models at time instant $i-1$.

The miscoverage rate was set to $\alpha = 0.1$, the learning rate to $\eta = 0.01$, and we chose $\gamma = 0.03$ for the calibration parameter θ in (23).

ACKNOWLEDGMENT

The authors acknowledge the use King's Computational Research, Engineering and Technology Environment (CREATE) Environment (CREATE). Retrieved October 2, 2023, from <https://doi.org/10.18742/rnvf-m076>.

An earlier version of this paper was presented in part at the 2023 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2023) [DOI: 10.1109/ICASSP49357.2023.10096780].

REFERENCES

- [1] B. Wang and K. J. R. Liu, "Advances in cognitive radio networks: A survey," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 1, pp. 5–23, Feb. 2011.
- [2] O. Simeone, I. Stanojev, S. Savazzi, Y. Bar-Ness, U. Spagnolini, and R. Pickholtz, "Spectrum leasing to cooperating secondary ad hoc networks," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 1, pp. 203–213, Jan. 2008.
- [3] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1321–1330.
- [4] A. Masegosa, "Learning under model misspecification: Applications to variational and ensemble methods," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 33, 2020, pp. 5479–5491.
- [5] W. R. Morningstar, A. Alemi, and J. V. Dillon, "PACm-Bayes: Narrowing the empirical risk gap in the misspecified Bayesian regime," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2022, pp. 8270–8298.
- [6] M. Zecchin, S. Park, O. Simeone, M. Kountouris, and D. Gesbert, "Robust PAC^m: Training ensemble models under misspecification and outliers," *IEEE Trans. Neural Netw. Learn. Syst.*, early access, Jul. 24, 2023, doi: 10.1109/TNNLS.2023.3295168.
- [7] P. Cannon, D. Ward, and S. M. Schmon, "Investigating the impact of model misspecification in neural simulation-based inference," 2022, *arXiv:2209.01845*.
- [8] D. T. Frazier, C. P. Robert, and J. Rousseau, "Model misspecification in approximate Bayesian computation: Consequences and diagnostics," *J. Roy. Stat. Soc. B, Stat. Methodol.*, vol. 82, no. 2, pp. 421–444, Apr. 2020.
- [9] J. Ridgway, "Probably approximate Bayesian computation: Nonasymptotic convergence of ABC under misspecification," 2017, *arXiv:1707.05987*.
- [10] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*, 2nd ed. New York, NY, USA: Springer, Jan. 2022.
- [11] M. Fontana, G. Zeni, and S. Vantini, "Conformal prediction: A unified review of theory and new challenges," *Bernoulli*, vol. 29, no. 1, pp. 1–23, Feb. 2023, doi: 10.3150/21-BEJ1447.
- [12] A. N. Angelopoulos and S. Bates, "Conformal prediction: A gentle introduction," *Found. Trends Mach. Learn.*, vol. 16, no. 4, pp. 494–591, 2023.
- [13] J. Vaicenavicius, D. Widmann, C. Andersson, F. Lindsten, J. Roll, and T. Schön, "Evaluating model calibration in classification," in *Proc. 22nd Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2019, pp. 3459–3467.
- [14] I. Gibbs and E. Candes, "Adaptive conformal inference under distribution shift," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 1660–1672.
- [15] O. Simeone, "A very brief introduction to machine learning with applications to communication systems," *IEEE Trans. Cogn. Commun. Netw.*, vol. 4, no. 4, pp. 648–664, Dec. 2018.
- [16] L. Dai, R. Jiao, F. Adachi, H. V. Poor, and L. Hanzo, "Deep learning for wireless communications: An emerging interdisciplinary paradigm," *IEEE Wireless Commun.*, vol. 27, no. 4, pp. 133–139, Aug. 2020.
- [17] Y. Eldar, A. Goldsmith, D. Gündüz, and H. Poor, *Machine Learning and Wireless Communications*. Cambridge, U.K.: Cambridge Univ. Press, 2022. [Online]. Available: <https://books.google.co.uk/books?id=JN59EAAAQBAJ>
- [18] T. Erpek, T. J. O'Shea, Y. E. Sagduyu, Y. Shi, and T. C. Clancy, "Deep learning for wireless communications," in *Development and Analysis of Deep Learning Architectures*. Cham, Switzerland: Springer, 2020, pp. 223–266, doi: 10.1007/978-3-030-31764-5_9.
- [19] Z. Sun, J. Wu, X. Li, W. Yang, and J.-H. Xue, "Amortized Bayesian prototype meta-learning: A new probabilistic meta-learning approach to few-shot image classification," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2021, pp. 1414–1422.
- [20] K. M. Cohen, S. Park, O. Simeone, and S. S. Shitz, "Bayesian active meta-learning for reliable and efficient AI-based demodulation," *IEEE Trans. Signal Process.*, vol. 70, pp. 5366–5380, 2022.
- [21] M. Zecchin, S. Park, O. Simeone, M. Kountouris, and D. Gesbert, "Robust Bayesian learning for reliable wireless AI: Framework and applications," *IEEE Trans. Cogn. Commun. Netw.*, vol. 9, no. 4, pp. 897–912, Aug. 2023.
- [22] E. Angelino, M. J. Johnson, and R. P. Adams, "Patterns of scalable Bayesian inference," *Found. Trends Mach. Learn.*, vol. 9, nos. 2–3, pp. 119–247, 2016.
- [23] O. Simeone, *Machine Learning for Engineers*. Cambridge, U.K.: Cambridge Univ. Press, 2022.
- [24] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural network," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1613–1622.
- [25] S. Ravi and A. Beaton, "Amortized Bayesian meta-learning," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Vancouver, BC, Canada, 2018, pp. 1–14.
- [26] A. Graves, "Practical variational inference for neural networks," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 24, Granada, Spain, 2011, pp. 1–9.
- [27] M. Dusenberry et al., "Efficient and scalable Bayesian neural nets with rank-1 factors," in *Proc. Int. Conf. Mach. Learn.*, Baltimore, MD, USA, 2020, pp. 2782–2792.
- [28] E. Daxberger, A. Kristiadi, A. Immer, R. Eschenhagen, M. Bauer, and P. Hennig, "Laplace redux—effortless Bayesian deep learning," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 34, 2021, pp. 20089–20103.
- [29] S. Farquhar, L. Smith, and Y. Gal, "Liberty or depth: Deep Bayesian neural nets do not need complex weight posterior approximations," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 33, 2020, pp. 4346–4357.
- [30] R. M. Neal, "MCMC using Hamiltonian dynamics," in *Handbook of Markov Chain Monte Carlo*. London, U.K.: Chapman & Hall, 2011, pp. 139–188.
- [31] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient Langevin dynamics," in *Proc. 28th Int. Conf. Mach. Learn. (ICML)*, Bellevue, WA, USA: Citeseer, 2011, pp. 681–688.
- [32] R. Zhang, C. Li, J. Zhang, C. Chen, and A. G. Wilson, "Cyclical stochastic gradient MCMC for Bayesian deep learning," in *Proc. 8th Int. Conf. Learn. Represent.*, Addis Ababa, Ethiopia, 2020, pp. 1–27.
- [33] O. Narmanlioglu, E. Zeydan, M. Kandemir, and T. Kranda, "Prediction of active UE number with Bayesian neural networks for self-organizing LTE networks," in *Proc. 8th Int. Conf. Netw. Future (NOF)*, Nov. 2017, pp. 73–78.
- [34] N. K. Jha and V. K. N. Lau, "Transformer-based online Bayesian neural networks for grant-free uplink access in CRAN with streaming variational inference," *IEEE Internet Things J.*, vol. 9, no. 9, pp. 7051–7064, May 2022.
- [35] L. Maggi, A. Valcarce, and J. Hoydis, "Bayesian optimization for radio resource management: Open loop power control," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 7, pp. 1858–1871, Jul. 2021, doi: 10.1109/JSAC.2021.3078490.
- [36] Z. Wu and H. Li, "Stochastic gradient Langevin dynamics for massive MIMO detection," *IEEE Commun. Lett.*, vol. 26, no. 5, pp. 1062–1065, May 2022.
- [37] N. Zilberstein, C. Dick, R. Doost-Mohammady, A. Sabharwal, and S. Segarra, "Annealed Langevin dynamics for massive MIMO detection," *IEEE Trans. Wireless Commun.*, vol. 22, no. 6, pp. 3762–3776, Jun. 2023.
- [38] Z. Tao and S. Wang, "Improved downlink rates for FDD massive MIMO systems through Bayesian neural networks-based channel prediction," *IEEE Trans. Wireless Commun.*, vol. 21, no. 3, pp. 2122–2134, Mar. 2022.
- [39] N. K. Jha and V. K. N. Lau, "Online downlink multi-user channel estimation for mmWave systems using Bayesian neural network," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 8, pp. 2374–2387, Aug. 2021.
- [40] R. Prasad, C. R. Murthy, and B. D. Rao, "Joint channel estimation and data detection in MIMO-OFDM systems: A sparse Bayesian learning approach," *IEEE Trans. Signal Process.*, vol. 63, no. 20, pp. 5369–5382, Oct. 2015.
- [41] X. Lv, Y. Li, Y. Wu, X. Wang, and H. Liang, "Joint channel estimation and impulsive noise mitigation method for OFDM systems using sparse Bayesian learning," *IEEE Access*, vol. 7, pp. 74500–74510, 2019.
- [42] J. Xu, Y. Shen, E. Chen, and V. Chen, "Bayesian neural networks for identification and classification of radio frequency transmitters using power amplifiers' nonlinearity signatures," *IEEE Open J. Circuits Syst.*, vol. 2, pp. 457–471, 2021.

- [43] X. Zhang, Y.-C. Liang, and J. Fang, "Bayesian learning based multiuser detection for M2M communications with time-varying user activities," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Paris, France, May 2017, pp. 1–6.
- [44] X. Zhang, Y.-C. Liang, and J. Fang, "Novel Bayesian inference algorithms for multiuser detection in M2M communications," *IEEE Trans. Veh. Technol.*, vol. 66, no. 9, pp. 7833–7848, Sep. 2017.
- [45] T. J. O'Shea, J. Corgan, and T. C. Clancy, "Convolutional radio modulation recognition networks," in *Engineering Applications of Neural Networks*. Cham, Switzerland: Springer, 2016, pp. 213–226.
- [46] T. J. O'Shea, T. Roy, and T. C. Clancy, "Over-the-air deep learning based radio signal classification," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 1, pp. 168–179, Feb. 2018.
- [47] X. Song et al., "A novel convolutional neural network based indoor localization framework with WiFi fingerprinting," *IEEE Access*, vol. 7, pp. 110698–110709, 2019.
- [48] N. Dvorecki, O. Bar-Shalom, L. Banin, and Y. Amizur, "A machine learning approach for Wi-Fi RTT ranging," in *Proc. Int. Tech. Meeting Inst. Navigat.*, Feb. 2019, pp. 435–444.
- [49] J. C. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," in *Advances in Large Margin Classifiers*. Cambridge, MA, USA: MIT Press, 1999, pp. 61–74.
- [50] B. Zadrozny and C. Elkan, "Transforming classifier scores into accurate multiclass probability estimates," in *Proc. 8th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2002, pp. 694–699.
- [51] A. Kumar, P. S. Liang, and T. Ma, "Verified uncertainty calibration," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1–21.
- [52] X. Ma and M. B. Blaschko, "Meta-cal: Well-controlled post-hoc calibration by ranking," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 7235–7245.
- [53] R. F. Barber, E. J. Candès, A. Ramdas, and R. J. Tibshirani, "Predictive inference with the jackknife+," *Ann. Statist.*, vol. 49, no. 1, pp. 486–507, Feb. 2021.
- [54] D. Stutz, K. D. Dvijotham, A. T. Cemgil, and A. Doucet, "Learning optimal conformal classifiers," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–27.
- [55] B.-S. Einbinder, Y. Romano, M. Sesia, and Y. Zhou, "Training uncertainty-aware classifiers with conformalized deep learning," in *Proc. Adv. Neural Inf. Process. Syst.*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022, pp. 1–16. [Online]. Available: https://openreview.net/forum?id=NaZwgxp-mT_
- [56] R. J. Tibshirani, R. F. Barber, E. Candès, and A. Ramdas, "Conformal prediction under covariate shift," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1–11.
- [57] Y. Yang and A. K. Kuchibhotla, "Finite-sample efficient conformal prediction," 2021, *arXiv:2104.13871*.
- [58] A. Kumar, S. Sarawagi, and U. Jain, "Trainable calibration measures for neural networks from kernel mean embeddings," in *Proc. 35th Int. Conf. Mach. Learn.*, vol. 80, Jul. 2018, pp. 2805–2814.
- [59] M. J. Holland, "Making learning more transparent using conformalized performance prediction," 2020, *arXiv:2007.04486*.
- [60] A. Perez-Lebel, M. L. Morvan, and G. Varoquaux, "Beyond calibration: Estimating the grouping loss of modern neural networks," in *Proc. 11th Int. Conf. Learn. Represent. (ICLR)*, Kigali, Rwanda, May 2023, pp. 1–47. [Online]. Available: <https://hal.science/hal-03829870>
- [61] S. Feldman, L. Ringel, S. Bates, and Y. Romano, "Achieving risk control in online learning settings," 2022, *arXiv:2205.09095*.
- [62] V. Vovk, G. Shafer, and I. Nourtdinov, "Self-calibrating probability forecasting," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 16, 2003, pp. 1–8.
- [63] V. Vovk and I. Petej, "Venn–Abers predictors," in *Proc. 13th Conf. Uncertainty Artif. Intell.*, 2014, pp. 829–838.
- [64] A. Lambrou, I. Nourtdinov, and H. Papadopoulos, "Inductive venn prediction," *Ann. Math. Artif. Intell.*, vol. 74, nos. 1–2, pp. 181–201, Jun. 2015.
- [65] H. Lofstrom, T. Lofstrom, U. Johansson, and C. Sonstrod, "Calibrated explanations: With uncertainty information and counterfactuals," 2023, *arXiv:2305.02305*.
- [66] I. C. Covert, S. Lundberg, and S.-I. Lee, "Explaining by removing: A unified framework for model explanation," *J. Mach. Learn. Res.*, vol. 22, no. 1, pp. 9477–9566, 2021.
- [67] K. M. Cohen, S. Park, O. Simeone, and S. S. Shitz, "Calibrating AI models for few-shot demodulation VIA conformal prediction," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Jun. 2023, pp. 1–5.
- [68] N. Simmons, S. B. F. Gomes, M. D. Yacoub, O. Simeone, S. L. Cotton, and D. E. Simmons, "AI-based channel prediction in D2D links: An empirical validation," *IEEE Access*, vol. 10, pp. 65459–65472, 2022.
- [69] A. Zanella and A. Bardella, "RSS-based ranging by multichannel RSS averaging," *IEEE Wireless Commun. Lett.*, vol. 3, no. 1, pp. 10–13, Feb. 2014.
- [70] B. de Finetti, *Theory of Probability*, vol. 1. New York, NY, USA: Wiley, 1974.
- [71] S. Park, O. Bastani, J. Weimer, and I. Lee, "Calibrated prediction with covariate shift via unsupervised domain adaptation," in *Proc. Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2020, pp. 3219–3229.
- [72] R. Koenker and G. Bassett Jr., "Regression quantiles," *Econometrica, J. Econ. Soc.*, vol. 46, no. 1, pp. 33–50, 1978.
- [73] I. Steinwart and A. Christmann, "Estimating conditional quantiles with the help of the pinball loss," *Bernoulli*, vol. 17, no. 1, pp. 211–225, Feb. 2011.
- [74] J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman, "Distribution-free predictive inference for regression," *J. Amer. Stat. Assoc.*, vol. 113, no. 523, pp. 1094–1111, Jul. 2018.
- [75] R. F. Barber, E. J. Candès, A. Ramdas, and R. J. Tibshirani, "Conformal prediction beyond exchangeability," *Ann. Statist.*, vol. 51, no. 2, pp. 816–845, Apr. 2023.
- [76] Y. Romano, M. Sesia, and E. Candès, "Classification with valid and adaptive coverage," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 3581–3591.
- [77] Z. Wang, R. Gao, M. Yin, M. Zhou, and D. Blei, "Probabilistic conformal prediction using conditional random samples," in *Proc. 26th Int. Conf. Artif. Intell. Statist.*, in Proceedings of Machine Learning Research, vol. 206, F. Ruiz, J. Dy, and J.-W. van de Meent, Eds. PMLR, Apr. 2023, pp. 8814–8836. [Online]. Available: <https://proceedings.mlr.press/v206/wang23n.html>
- [78] M. Zaffran, O. Féron, Y. Goude, J. Josse, and A. Dieuleveut, "Adaptive conformal predictions for time series," in *Proc. Int. Conf. Mach. Learn.*, 2022, pp. 25834–25866.
- [79] Z. Lin, S. Trivedi, and J. Sun, "Conformal prediction with temporal quantile adjustments," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 31017–31030.
- [80] S. Park, H. Jang, O. Simeone, and J. Kang, "Learning to demodulate from few pilots via offline and online meta-learning," *IEEE Trans. Signal Process.*, vol. 69, pp. 226–239, 2021.
- [81] N. Seshadri and C.-E.-W. Sundberg, "List Viterbi decoding algorithms with applications," *IEEE Trans. Commun.*, vol. 42, no. 2, pp. 313–323, Feb. 1994.
- [82] I. Tal and A. Vardy, "List decoding of polar codes," *IEEE Trans. Inf. Theory*, vol. 61, no. 5, pp. 2213–2226, May 2015.
- [83] D. Tandur and M. Moonen, "Joint adaptive compensation of transmitter and receiver IQ imbalance under carrier frequency offset in OFDM-based systems," *IEEE Trans. Signal Process.*, vol. 55, no. 11, pp. 5246–5252, Nov. 2007.
- [84] Z. Zhu and A. K. Nandi, *Automatic Modulation Classification: Principles, Algorithms and Applications*. Hoboken, NJ, USA: Wiley, 2015.
- [85] R. Zhou, F. Liu, and C. W. Gravelle, "Deep learning for modulation recognition: A survey with a demonstration," *IEEE Access*, vol. 8, pp. 67366–67376, 2020.
- [86] *Reconfigurable Radio Systems (RRS): Feasibility Study on Radio Frequency (RF) Performance for Cognitive Radio Systems Operating in UHF TV Band White Spaces*, Standard TR 103 067, Version 1.1.1, ETSI, May 2013.
- [87] P. S. Aparna and M. Jayasheela, "Cyclostationary feature detection in cognitive radio using different modulation schemes," *Int. J. Comput. Appl.*, vol. 47, no. 21, pp. 12–16, Jun. 2012.
- [88] J. Nieminen, R. Jantti, and L. Qian, "Time synchronization of cognitive radio networks," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Nov. 2009, pp. 1–6.
- [89] Z. Kollar and P. Horvath, "Physical layer considerations for cognitive radio: Synchronization point of view," in *Proc. IEEE 73rd Veh. Technol. Conf. (VTC Spring)*, May 2011, pp. 1–5.
- [90] I. C. Wong and B. L. Evans, "Optimal resource allocation in the OFDMA downlink with imperfect channel knowledge," *IEEE Trans. Commun.*, vol. 57, no. 1, pp. 232–241, Jan. 2009.
- [91] I. Nikoloska and O. Simeone, "Modular meta-learning for power control via random edge graph neural networks," *IEEE Trans. Wireless Commun.*, vol. 22, no. 1, pp. 457–470, Jan. 2023.

[92] H. Shokri-Ghadikolaei, C. Fischione, G. Fodor, P. Popovski, and M. Zorzi, "Millimeter wave cellular networks: A MAC layer perspective," *IEEE Trans. Commun.*, vol. 63, no. 10, pp. 3437–3458, Oct. 2015.

[93] Y. Zhang, O. Simeone, S. T. Jose, L. Maggi, and A. Valcarce, "Bayesian and multi-armed contextual meta-optimization for efficient wireless radio resource management," *IEEE Trans. Cogn. Commun. Netw.*, early access, Jun. 19, 2023, doi: 10.1109/TCCN.2023.3287240.

[94] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[95] S. Park, K. M. Cohen, and O. Simeone, "Few-shot calibration of set predictors via meta-learned cross-validation-based conformal prediction," in *Proc. 6th Workshop Meta-Learn. Conf. Neural Inf. Process. Syst.*, 2022, pp. 1–22. [Online]. Available: <https://openreview.net/forum?id=S0ItikPStJy>

[96] F. Wenzel et al., "How good is the Bayes posterior in deep neural networks really?" in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 10248–10259.

[97] N. Ye, Z. Zhu, and R. Mantiuk, "Langevin dynamics with continuous tempering for training deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 618–626.



KFIR M. COHEN (Student Member, IEEE) received the B.Sc. (summa cum laude) and M.Sc. degrees from the Faculty of Electrical Engineering, Technion—Israel Institute for Technology, Haifa, Israel, in 2006 and 2013, respectively. He is currently pursuing the Ph.D. degree with the King's Communications, Learning and Information Processing Laboratory, Department of Engineering, King's College London (KCL). Before joining KCL, his last position was as a Communication Signal Processing Engineer. He served for 15 years in different research and development roles. His research interests include Bayesian and reliable machine learning, conformal prediction, signal processing, and their applications to communications systems.



SANGWOO PARK (Member, IEEE) received the B.S. degree in physics and the M.S.E. and Ph.D. degrees in electrical engineering from the Korea Advanced Institute of Science and Technology, Daejeon, South Korea, in 2014, 2016, and 2020, respectively. From April 2019 to December 2019, he was a Visiting Ph.D. Student with the King's Communications, Learning and Information Processing Laboratory (KCLIP), King's College London (KCL), U.K., where he is currently a

Research Associate. His research interests include practical, reliable AI, and its application to engineering systems, including wireless communications and quantum information processing.



OSVALDO SIMEONE (Fellow, IEEE) received the M.Sc. (Hons.) and Ph.D. degrees in information engineering from Politecnico di Milano, Milan, Italy, in 2001 and 2005, respectively. From 2006 to 2017, he was a Faculty Member of the Department of Electrical and Computer Engineering (ECE), New Jersey Institute of Technology (NJIT), where he was affiliated with the Center for Wireless Information Processing (CWIP). He is currently a Professor of information

engineering with the Centre for Telecommunications Research, Department of Engineering, King's College London, where he directs the King's Communications, Learning and Information Processing Laboratory. His research has been supported by the U.S. National Science Foundation, the European Commission, the European Research Council, the Vienna Science and Technology Fund, and the European Space Agency, as well as by a number of industrial collaborations, including with Intel Labs and InterDigital. He is the author of the textbook *Machine Learning for Engineers* published by Cambridge University Press, four monographs, two edited books, and more than 200 research journals and magazine papers. His research interests include information theory, machine learning, wireless communications, neuromorphic computing, and quantum machine learning. He is a fellow of the IET and EPSRC. He was a co-recipient of the 2022 IEEE Communications Society Outstanding Paper Award, the 2021 IEEE Vehicular Technology Society Jack Neubauer Memorial Award, the 2019 IEEE Communication Society Best Tutorial Paper Award, the 2018 IEEE Signal Processing Best Paper Award, the 2017 JCN Best Paper Award, the 2015 IEEE Communication Society Best Tutorial Paper Award, and the Best Paper Awards of IEEE SPAWC 2007 and IEEE WRECOM 2007. He was awarded an Open Fellowship by the EPSRC in 2022 and a Consolidator Grant by the European Research Council (ERC) in 2016. He was the Chair of the Signal Processing for Communications and Networking Technical Committee of the IEEE Signal Processing Society in 2022 and the U.K. and Ireland Chapter of the IEEE Information Theory Society from 2017 to 2022. He was a Distinguished Lecturer of the IEEE Communications Society from 2021 to 2022 and the IEEE Information Theory Society from 2017 to 2018.



SHLOMO SHAMAI (SHITZ) (Life Fellow, IEEE) is currently with the Department of Electrical Engineering, Technion—Israel Institute of Technology, where he is currently a Technion Distinguished Professor, and holds the William Fondiller Chair of Telecommunications. He is also a URSI Fellow, a member of the Israeli Academy of Sciences and Humanities, and a Foreign Member of the U.S. National Academy of Engineering. He was a recipient of the 2011 Claude E. Shannon

Award, the 2014 Rothschild Prize in Mathematics/Computer Sciences and Engineering, the 2017 IEEE Richard W. Hamming Medal, and numerous technical and paper awards and recognitions of the IEEE (Donald G. Fink Prize Paper Award), Information Theory, Communications, and Signal Processing Societies as well as EURASIP. He was also a co-recipient of the 2018 Third Bell Labs Prize for Shaping the Future of Information and Communications Technology. He was listed as a Highly Cited Researcher (Computer Science) from 2013 to 2018. He has served as an Associate Editor for the Shannon Theory of the IEEE Transactions on Information Theory. He has also served twice on the Board of Governors of the Information Theory Society and an Executive Editorial Board for the IEEE Transactions on Information Theory, the IEEE Information Theory Society Nominations and Appointments Committee, and the IEEE Information Theory Society, Shannon Award Committee.