

Bayesian Active Meta-Learning for Reliable and Efficient AI-Based Demodulation

Kfir M. Cohen , *Student Member, IEEE*, Sangwoo Park , *Member, IEEE*, Osvaldo Simeone , *Fellow, IEEE*, and Shlomo Shamai (Schitz) , *Life Fellow, IEEE*

Abstract—Two of the main principles underlying the life cycle of an artificial intelligence (AI) module in communication networks are *adaptation* and *monitoring*. Adaptation refers to the need to adjust the operation of an AI module depending on the current conditions; while monitoring requires measures of the reliability of an AI module’s decisions. Classical frequentist learning methods for the design of AI modules fall short on both counts of adaptation and monitoring, catering to one-off training and providing overconfident decisions. This paper proposes a solution to address both challenges by integrating meta-learning with Bayesian learning. As a specific use case, the problems of demodulation and equalization over a fading channel based on the availability of few pilots are studied. Meta-learning processes pilot information from multiple frames in order to extract useful shared properties of effective demodulators across frames. The resulting trained demodulators are demonstrated, via experiments, to offer better calibrated soft decisions, at the computational cost of running an ensemble of networks at run time. The capacity to quantify uncertainty in the model parameter space is further leveraged by extending Bayesian meta-learning to an active setting. In it, the designer can select in a sequential fashion channel conditions under which to generate data for meta-learning from a channel simulator. Bayesian active meta-learning is seen in experiments to significantly reduce the number of frames required to obtain efficient adaptation procedure for new frames.

Index Terms—Bayesian meta-learning, uncertainty quantification, Bayesian active meta-learning, demodulation.

I. INTRODUCTION

A. Motivation

Artificial intelligence (AI) is seen as a key enabler for next-generation wireless systems [2]. Emerging solutions, such as

Manuscript received 9 March 2022; revised 22 July 2022; accepted 23 October 2022. Date of publication 11 November 2022; date of current version 21 November 2022. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Yuxin Chen. The work of Kfir M. Cohen, Sangwoo Park, and Osvaldo Simeone was supported by the European Research Council (ERC) through European Union’s Horizon 2020 Research and Innovation Programme under Grant 725731. The work of Osvaldo Simeone was also supported by the Open Fellowship of the EPSRC. The work of Shlomo Shamai was supported by the European Union’s Horizon 2020 Research and Innovation Programme under Grant 694630. An earlier version of this work was presented in WSA 2021 - 25th International ITG Workshop on Smart Antennas. (Corresponding author: Sangwoo Park.)

Kfir M. Cohen, Sangwoo Park, and Osvaldo Simeone are with the King’s Communication, Learning, & Information Processing (KCLIP) lab, Department of Engineering, King’s College London, WC2R 2LS London, U.K. (e-mail: kfir.cohen@kcl.ac.uk; sangwoo.park@kcl.ac.uk; osvaldo.simeone@kcl.ac.uk).

Shlomo Shamai is with the Viterbi Faculty of Electrical and Computing Engineering, Technion–Israel Institute of Technology, Haifa 3200003, Israel (e-mail: sshlomo@ee.technion.ac.il).

Digital Object Identifier 10.1109/TSP.2022.3220035

Open-Radio Access Network (O-RAN), incorporate AI modules as native components of a modular architecture that can be fine-tuned to meet the requirements of specific deployments [3]. Two of the main principles underlying the life cycle of an AI module in communication networks are *adaptation* and *monitoring* [4]. Adaptation refers to the need to adjust the operation of an AI module depending on the current conditions, particularly for real-time applications at the frame level. At run time, an AI model should ideally enable monitoring of the quality of its outputs by providing measures of the reliability of its decisions. The availability of such reliability measures is instrumental in supporting several important functionalities, from the combination of multiple models to decisions about retraining [5].

Classical frequentist learning methods for the design of AI modules fall short on both counts of adaptation and monitoring (see, e.g., [6], [7]). First, conventional frequentist learning is well known to provide inaccurate measures of reliability, typically producing overconfident decisions [7]. Second, the standard learning approach prescribes the one-off optimization of an AI model, hence failing to capture the need for adaptation. This paper investigates the integration of meta-learning and Bayesian learning as a means to address both challenges. As we detail in the next section, Bayesian learning can provide well-calibrated, and hence reliable, measures of uncertainty of a model’s decision; while meta-learning can reduce the amount of data required for adaptation to a new task, thus improving efficiency. As a specific use case, we focus on the problems of demodulation and equalization over a fading channel based on the availability of few pilots (see Fig. 1). The goal is to develop AI solutions that are capable of adapting the demodulator/equalizer to changing conditions based on few training symbols, while also being able to quantify the uncertainty of the AI model’s output.

B. Background

As illustrated in Fig. 2, frequentist learning assigns a single value to each model parameter as a result of training. This neglects (epistemic) uncertainty that exists at the level of model parameters due to the limited availability of data. In contrast, *Bayesian learning* can express uncertainty about the true value of the model parameter vector by optimizing over a distribution, rather than over a single point value [8]. By averaging predictions over the distribution of the model parameters, Bayesian learning is known to be capable of providing decisions that are well calibrated [9], [10], [11]. Calibration refers to the capacity of

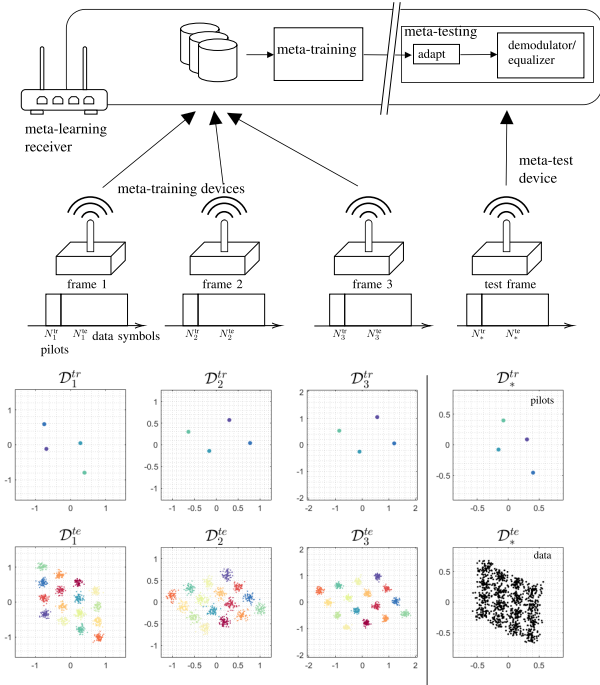


Fig. 1. Illustration of the meta-learning problem studied in this work for the example of 16-ary quadrature amplitude modulation (16-QAM). A receiver has available data corresponding to frames previously received from multiple devices, each possibly experiencing different channel conditions. Given meta-training data sets $\{\mathcal{D}_\tau\}_{\tau=1}^t$ of pilots from previous frames, partitioned into training data and test data, the demodulator optimizes a hyperparameter vector ξ . For a newly received frame, the receiver uses the few pilots therein to adapt the demodulator/equalizer parameter vector ϕ_* . In the Bayesian meta-learning framework, instead of a single parameter vector ϕ_* , the receiver optimizes over an ensemble of parameter vectors through the hyperparameter vector ξ of a posterior distribution $p(\phi_*|\mathcal{D}_*^{\text{tr}}, \xi)$.

a model to produce confidence levels that reproduce well the actual accuracy of the decisions.

Meta-learning, also known as *learning to learn*, optimizes training strategies that can fine-tune a model based on few samples for a new task by transferring knowledge across different learning tasks [12], [13], [14], [15], [16], [17], [18]. Meta-learning is a natural tool to produce AI solutions that are optimized for adaptation. Prior work on meta-learning for communication systems, including [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], is limited to standard frequentist learning. Therefore, existing art is unable to produce models that provide well-calibrated estimations of reliability. Most related to our work is [19], which proposes to leverage pilot information from previous frames in order to optimize training procedures to be applied to the pilots of new frames (see Fig. 1).

Bayesian meta-learning aims at optimizing the procedure that produces the posterior distribution for new learning tasks. Accordingly, the goal of Bayesian meta-learning is to enhance the efficiency of Bayesian learning by reducing the number of training points needed to obtain accurate and well calibrated Bayesian models. The optimization of the Bayesian learning process is carried out by transferring knowledge from previously encountered tasks for which data are assumed to be available [30], [31], [32]. To the best of our knowledge, with the

exception of the conference version of this paper [1], this is the first work to consider the application of Bayesian meta-learning to communication systems.

Beside meta-learning, another approach to reduce the number of required training data points is *active learning* [33], [34], [35], [36], [37]. Active learning amounts to the process of choosing which samples should be annotated next and incrementally added to the training set [38]. Through this process, active learning can select relevant samples at which the model is currently most uncertain in order to speed up the training process.

A much less studied area is *active meta-learning*, which aims at reducing the number of tasks a meta-learner must collect data from, before it can adapt efficiently to new tasks [37], [39]. Reference [37] proposes a method based on Bayesian meta-learning via empirical Bayes; while the paper [39] takes a hierarchical Bayesian approach, generalizing the Bayesian active learning by disagreements (BALD) criterion introduced in [33] to meta-learning. While [37] assumes labeled training sets, reference [39] considers unlabeled data during active meta-learning. As such, the setting in it is not applicable to the problem under study here in which data consists of supervised pairs of pilots and received signals (see Fig. 1). A summary of the relevant approaches built upon in this work is given in Table I.

C. Contributions

This paper introduces the use of Bayesian meta-learning to enable both adaptation and monitoring for the tasks of demodulation and equalization. Unlike prior works that considered either frequentist meta-learning [6], [19], [20], [21], [22], [23], [24], [25], [26] or Bayesian learning [40], [41], [42], [43], the proposed Bayesian meta-learning methodology enables both resource-efficient adaptation and a reliable quantification of uncertainty. To further improve the efficiency of Bayesian meta-learning we propose the use of active meta-learning, which reduces the number of required meta-training data from previously received frames. Specific contributions are as follows.

- We introduce Bayesian meta-learning for the problems of demodulation and equalization from few pilots. The proposed implementation is derived based on parametric VI.
- We introduce Bayesian active meta-learning as a solution to reduce the number of frames required by meta-learning. Active meta-learning selects in a sequential fashion channel conditions under which to generate data for meta-learning from a channel simulator.
- Extensive experimental results demonstrate that Bayesian meta-learning produces demodulators and equalizers that offer better calibrated soft decisions. Furthermore, they show that for a target meta-testing loss, active meta-learning can reduce the number of simulated meta-training frames required.

Part of this paper was presented in [1], which presented the idea of Bayesian meta-learning with some preliminary experiments. This journal version presents full technical details, new results and introduces for the first time Bayesian active meta-learning for communication systems.

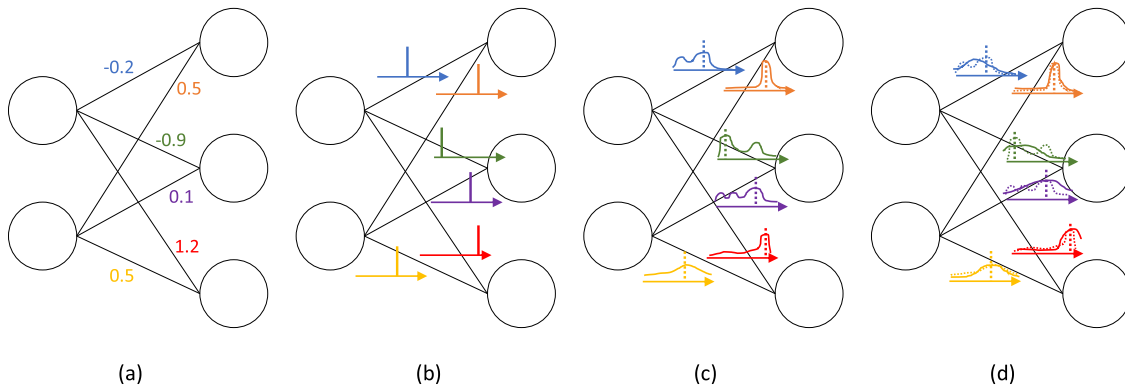


Fig. 2. Network weights in frequentist and Bayesian learning: (a) in frequentist learning, each weight is described by a scalar value; (b) the scalar value can be viewed as random variable having a degenerated probabilistic distribution concentrated at a simple prior; (c) in Bayesian learning, the weights are assigned a probability distribution, which, unlike the frequentist point estimate (dashed vertical line), provides information about the uncertainty on the weight; (d) in variational inference (VI), the posterior is approximated with a parameter distribution.

TABLE I
A SUMMARY OF THE RELEVANT TECHNIQUES CONSIDERED IN THIS WORK

Approach	Goal	Methodology
frequentist learning	accurate data-driven predictions	minimize the training loss over the model parameter vector ϕ
Bayesian learning via variational inference	reliable and accurate data-driven predictions	minimize the free energy over the parameters φ of the variational distribution $q(\phi \varphi)$
frequentist meta-learning	sample efficiency	minimize the meta-training loss over hyperparameters ξ to be used for frequentist learning
Bayesian meta-learning	sample efficiency and reliability	minimize the meta-training loss over hyperparameters ξ to be used for Bayesian learning
Bayesian active meta-learning	task efficiency, sample efficiency, and reliability	minimize the meta-training loss over the hyperparameters ξ and over the sequential selection of meta-learning tasks

The rest of the paper is organized as follows. Section II introduces the channel model, along with background material on standard frequentist learning and frequentist meta-learning. Section III expands on Bayesian meta-learning. Then, we present Bayesian active meta-learning in Section IV. Numerical results are presented in Section V, and Section VI concludes the paper.

D. Related Work

For scalability, Bayesian learning can be implemented via approximate methods based on variational inference (VI) or Monte Carlo (MC). VI methods approximate the exact Bayesian posterior distribution with a tractable variational density [9], [44], [45], [46], [47], while Monte Carlo techniques obtain approximate samples from the Bayesian posterior distribution [48], [49], [50]. Each class of methods comes with its own set of technical challenges and engineering choices. For instance, VI requires the selection of a variational distribution family, such as mean-field Gaussian models, and the specification of a stochastic optimization algorithm. There are also non-parametric VI methods such as Stein variational gradient descent (SVGD) [51], which optimize over deterministic and interacting particles. For MC techniques, solutions range from first-order Langevin dynamics techniques [49] to more complex methods such as Hamiltonian Monte Carlo (HMC) [48]. Implementing any of these schemes for a specific engineering application is a non-trivial task.

Bayesian learning has been applied in reference [52] to the problem of predicting the number of active users in LTE system; papers [53], [54] applied MC-based Bayesian learning for MIMO detection; the works [55], [56], [57] addressed channel prediction/estimation for massive MIMO systems; reference [58] studied the identification of IoT transmitters; and the authors of [59] proposed the use of robust Bayesian learning for modulation classification, localization, and channel modeling.

As for active learning, applications to communication systems include paper [60], which proposed a sample-efficient retransmission protocol; reference [61], which tackled initial beam alignment for massive MIMO system; work [62], which aimed at mitigating the problem of scarce training data in wireless cyber-security attack; and reference [63], which addressed resource allocation problems in vehicular communication systems.

Like Bayesian learning, meta-learning also provides a general design principle, which can be implemented by following different approaches. Optimization-based methods design the hyperparameters used by training algorithms; model-based techniques optimize an additional neural network model to guide adaptation of the main AI model; and metric-based schemes identify metric spaces for non-parametric inference (see, e.g., [64] and references therein).

The integration of meta-learning and Bayesian learning is highly non-trivial, and is an active topic of research in the

machine learning literature. References [65], [66], [67], [68], [69] addressed Bayesian meta-learning via empirical Bayes using parametric VI [65], [66], particle-based VI [67], deep-kernels [68], and expectation-maximization [69]; while the papers [70], [71], [72], [73] studied full Bayesian meta-learning that treats also the hyperparameters as random variables. Lastly, the work [74] proposed the use of quantum machine learning models as parameterized variational distributions.

II. CHANNEL MODEL AND BACKGROUND

A. Channel Model and Soft Demodulation or Equalization

In this paper, we consider frame-based transmission over a memoryless block fading channel model with constellation \mathcal{X} and channel output's alphabet \mathcal{Y} . The channel is characterized by a conditional distribution $p(y|x, c)$ of received symbol $y \in \mathcal{Y}$ given transmitted symbol $x \in \mathcal{X}$ and channel state c . In the case of demodulation, we treat the set \mathcal{X} as discrete; while for equalization we view it as the space of vectors of a certain size. In both cases, we will refer to channel input x as symbol. The channel state c is constant within each frame, and it is independently and identically distributed (i.i.d.) across frames according to an unknown distribution $p(c)$. At frame τ , the transmitter sends a packet consisting of N_τ symbols $x_\tau = \{x_\tau[i]\}_{i=1}^{N_\tau}$. Given the channel state c_τ and the transmitted symbols, collected in a vector x_τ , the received samples $y_\tau = \{y_\tau[i]\}_{i=1}^{N_\tau}$ are conditionally independent and each received i -th sample is distributed as $y_\tau[i] \sim p(y_\tau[i]|x_\tau[i], c_\tau)$.

A *soft demodulator/equalizer* is a conditional distribution $p(x|y, \phi)$ that maps channel outputs $y \in \mathcal{Y}$ to estimated probabilities for channel input symbol $x \in \mathcal{X}$. The demodulator/equalizer depends on a vector of parameters ϕ , and is applied separately to each received sample $y[i]$ in a memoryless fashion as $p(x|y_\tau[i], \phi)$. The ideal frame-specific parameter vector ϕ_τ for the frame τ is the one that best approximates the channel conditional distribution $p(x_\tau|y_\tau, c_\tau)$, within its model class, obtained from the Bayes rule as

$$p(x_\tau|y_\tau, \phi_\tau) \approx p(x_\tau|y_\tau, c_\tau) = \frac{p(y_\tau|x_\tau, c_\tau)p(x_\tau)}{\sum_{x'_\tau \in \mathcal{X}} p(y_\tau|x'_\tau, c_\tau)p(x'_\tau)}, \quad (1)$$

where $p(x_\tau)$ is the distribution of the input symbol vector x_τ . In practice, as we detail below, the demodulator/equalizer is optimized based on pilot symbols. To simplify the terminology, we will also refer to demodulation/equalization as *prediction* henceforth.

B. Conventional Data-Driven Demodulators/Equalizers

Pilot-aided schemes utilize available pilot symbols to adapt the predictor $p(x|y, \phi)$ to the unknown channel state c in each frame τ . A typical choice for a predictor is a multi-layer neural network [75]. With L layers, given received sample y , this class of models produces a vector

$$a(y|\phi) = W_L \cdot f_{W_{L-1}, b_{L-1}} \circ \dots \circ f_{W_1, b_1}(y) + b_L, \quad (2)$$

where \circ is the composition operator; the weights $\{W_l\}_{l=1}^L$ and biases $\{b_l\}_{l=1}^L$ define the model parameter vector $\phi :=$

$\{W_l, b_l\}_{l=1}^L$ for a total of D parameters; and the function for the l -th layer f_{W_l, b_l} is a linear mapping followed by an entry-wise activation function $h(\cdot)$, i.e., $y_l = f_{W_l, b_l}(y_{l-1}) = h(W_l \cdot y_{l-1} + b_l)$ with $y_0 = y$. In the last, L -th layer, a *soft demodulator* applies the softmax function to vector $a(y|\phi)$, producing the probability distribution

$$p(x|y, \phi) = [\text{softmax}(a(y|\phi))]_x = \frac{\exp([a(y|\phi)]_x)}{\sum_{x' \in \mathcal{X}} \exp([a(y|\phi)]_{x'})}, \quad (3)$$

using $[\cdot]_x$ as the x -th element of the vector. In contrast, a *soft equalizer* typically defines the conditional distribution

$$p(x|y, \phi) = \mathcal{N}(x|a(y|\phi), \beta^{-1}), \quad (4)$$

where the precision β is fixed. Throughout this paper, we use $\mathcal{N}(x|\mu, \Sigma)$ to indicate the probability density function of a Gaussian vector with mean μ and covariance matrix Σ .

In each frame τ , *conventional learning* optimizes the model parameters ϕ_τ using N_τ^{tr} i.i.d. pilots $\mathcal{D}_\tau^{\text{tr}} = \{(y_\tau^{\text{tr}}[i], x_\tau^{\text{tr}}[i])\}_{i=1}^{N_\tau^{\text{tr}}}$ as training data. Optimization of the prediction aims at minimizing the *training log-loss*

$$\mathcal{L}_{\mathcal{D}_\tau^{\text{tr}}}(\phi_\tau) := -\frac{1}{N_\tau^{\text{tr}}} \sum_{i=1}^{N_\tau^{\text{tr}}} \log p(x_\tau^{\text{tr}}[i]|y_\tau^{\text{tr}}[i], \phi_\tau), \quad (5)$$

which amounts to the cross entropy for demodulation (3) and the quadratic prediction loss for equalization (4). Minimization of (5) can be done via *gradient descent* (GD), or stochastic GD (SGD), a variant thereof [76].

GD updates model parameter vector ϕ_τ for I iterations with learning rate $\eta > 0$ starting from an initialization vector ξ . Accordingly, the updated parameters $\phi_\tau := \phi^{\text{GD}}(\mathcal{D}_\tau^{\text{tr}}|\xi)$ are obtained via the iterations

$$\begin{aligned} \phi_\tau^{(0)} &= \xi, \\ \forall i = 1, \dots, I: \quad \phi_\tau^{(i)} &\leftarrow \phi_\tau^{(i-1)} - \eta \nabla_{\phi_\tau^{(i-1)}} \mathcal{L}_{\mathcal{D}_\tau^{\text{tr}}}(\phi_\tau^{(i-1)}), \\ \phi^{\text{GD}}(\mathcal{D}_\tau^{\text{tr}}|\xi) &= \phi_\tau^{(I)}. \end{aligned} \quad (6)$$

The resulting prediction for a test input-output pair $(y_\tau^{\text{te}}[i], x_\tau^{\text{te}}[i])$ is given as $p(x_\tau^{\text{te}}[i]|y_\tau^{\text{te}}[i], \phi^{\text{GD}}(\mathcal{D}_\tau^{\text{tr}}|\xi))$.

C. Frequentist Meta-Learning

The most prominent shortcoming of conventional learning is its potentially high sample complexity, which translates into the need for a large number of pilots, N_τ^{tr} , per frame. Meta-learning addresses this issue by transferring knowledge acquired over previous frames. Specifically, frequentist meta-learning, as proposed in [19], treats the initialization vector ξ in (6) as a hyperparameter vector to be optimized based on the availability of pilots from t previous transmission frames.

As a preliminary step, we decompose the available pilots from each frame τ into a disjoint training set $\mathcal{D}_\tau^{\text{tr}}$ and test set $\mathcal{D}_\tau^{\text{te}}$ as $\mathcal{D}_\tau = \{\mathcal{D}_\tau^{\text{tr}}, \mathcal{D}_\tau^{\text{te}}\}$. Furthermore, the data sets for all previous t frames are stacked as $\mathcal{D}_{1:t} = \{\mathcal{D}_\tau\}_{\tau=1}^t$, and similarly for $\mathcal{D}_{1:t}^{\text{te}} = \{\mathcal{D}_\tau^{\text{te}}\}_{\tau=1}^t$, having a total of $N_{1:t}^{\text{te}} = \sum_{\tau=1}^t N_\tau^{\text{te}}$

samples. Meta-learning has two phases: meta-training and meta-testing. These are defined next by following the frequentist meta-learning strategy of [19].

Meta-training tackles the bi-level optimization problem

$$\begin{aligned} \min_{\xi} \quad & \frac{1}{N_{1:t}^{\text{te}}} \sum_{\tau=1}^t N_{\tau}^{\text{te}} \mathcal{L}_{\mathcal{D}_{\tau}^{\text{te}}}(\phi_{\tau}(\mathcal{D}_{\tau}^{\text{tr}}|\xi)) \quad (7a) \\ \text{s.t.} \quad & \phi_{\tau}(\mathcal{D}_{\tau}^{\text{tr}}|\xi) = \underset{\phi}{\text{argmin}} \mathcal{L}_{\mathcal{D}_{\tau}^{\text{te}}}(\phi), \quad \tau = 1, \dots, t. \quad (7b) \end{aligned}$$

The notation $\phi(\xi)$ in (7b) indicates the dependence of the optimizer on the initialization vector ξ . By (7), the goal of frequentist meta-training is to find a hyperparameter vector ξ such that for any frame τ , the optimized model parameter vector $\phi_{\tau}(\mathcal{D}_{\tau}^{\text{tr}}|\xi)$ fits well the test data set $\mathcal{D}_{\tau}^{\text{te}}$.

Problem (7) is addressed via a nested loop optimization involving SGD-based inner updates and SGD-based outer updates, which are also referred as meta-iterations. The inner loop tackles the inner optimization (7b) in a per-frame manner via (6) for a randomly selected subset $\mathcal{T} \subset \{1, \dots, t\}$ of frames, which are redrawn independently at each meta-iteration. The outer loop addresses the outer optimization (7a) via an SGD step of the meta-loss with learning-rate $\kappa > 0$, i.e.,

$$\xi \leftarrow \xi - \kappa \frac{1}{N_{\mathcal{T}}^{\text{te}}} \sum_{\tau \in \mathcal{T}} N_{\tau}^{\text{te}} \nabla_{\xi} \mathcal{L}_{\mathcal{D}_{\tau}^{\text{te}}}(\phi^{\text{GD}}(\mathcal{D}_{\tau}^{\text{tr}}|\xi)), \quad (8)$$

based on data from the batch \mathcal{T} of selected frames, and using the notation $N_{\mathcal{T}}^{\text{te}} = \sum_{\tau \in \mathcal{T}} N_{\tau}^{\text{te}}$ for the total samples within the batch of selected frames. Meta-training updates the initialization vector ξ across multiple meta-iterations. When meeting some stopping criterion, here determined by a predefined number of meta-iterations I_{meta} , meta-training stops, and the hyperparameter vector ξ is stored to be used for future learning tasks.

Upon deployment, i.e., during *meta-testing*, the meta-test frames also include pilots and data as the meta-training frames. Accordingly, each meta-test device loads the hyperparameter vector ξ for initialization, and produces the adapted model parameter vector $\phi_{*} = \phi^{\text{GD}}(\mathcal{D}_{*}^{\text{tr}}|\xi)$ as in (6) using N_{*}^{tr} pilots symbols $\mathcal{D}_{*}^{\text{tr}} = \{(y_{*}^{\text{tr}}[i], x_{*}^{\text{tr}}[i])\}_{i=1}^{N_{*}^{\text{tr}}}$. Then, it applies the learned model to the payload data symbols $\{y_{*}^{\text{te}}[i]\}_{i=1}^{N_{*}^{\text{te}}}$ to carry out demodulation or equalization

$$p(x_{*}^{\text{te}}[i]|y_{*}^{\text{te}}[i], \phi_{*}). \quad (9)$$

III. THE BAYESIAN FRAMEWORK

A. Bayesian Learning

Bayesian learning treats the model parameter vector ϕ_{τ} for some frame τ as a random vector, rather than as a deterministic optimization variable as in frequentist learning framework. As illustrated in Fig. 2, instead of producing a single demodulator parameters $\phi_{\tau} = \phi^{\text{GD}}(\mathcal{D}_{\tau}^{\text{tr}}|\xi)$ as in (6), Bayesian learning produces a distribution $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)$ over the space of the demodulator parameters ϕ_{τ} . This distribution is computed based on training data $\mathcal{D}_{\tau}^{\text{tr}}$ and on *predetermined* prior distribution $p(\phi_{\tau}|\xi)$, which depends in turn on the hyperparameter vector ξ , also fixed *a priori*.

Having obtained the distribution $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)$, the *ensemble prediction* of a test point $(y_{\tau}^{\text{te}}[i], x_{\tau}^{\text{te}}[i])$ is given by the ensemble average of the predictions $p(x_{\tau}^{\text{te}}[i]|y_{\tau}^{\text{te}}[i], \phi_{\tau})$ with random vector ϕ_{τ} having distribution $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)$, i.e.,

$$p(x_{\tau}^{\text{te}}[i]|y_{\tau}^{\text{te}}[i], \mathcal{D}_{\tau}^{\text{tr}}, \xi) = \mathbb{E}_{p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)} [p(x_{\tau}^{\text{te}}[i]|y_{\tau}^{\text{te}}[i], \phi_{\tau})]. \quad (10)$$

The frequentist prediction (9), reviewed in the previous section, can be viewed as a special case in which one is limited to the choice $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi) = \delta(\phi_{\tau} - \phi^{\text{GD}}(\mathcal{D}_{\tau}^{\text{tr}}|\xi))$, with $\delta(\cdot)$ indicating the Dirac Delta. With this choice, the distribution $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)$ is concentrated at one point, namely the GD solution (6). The frequentist approach is therefore inherently limited in its capacity to express uncertainty on the model parameters due to limited data.

Ideally, the distribution $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)$ should be obtained as the posterior distribution

$$p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi) \propto p(\phi_{\tau}|\xi) p(\mathcal{D}_{\tau}^{\text{tr}}|\phi_{\tau}), \quad (11)$$

where $p(\mathcal{D}_{\tau}^{\text{tr}}|\phi_{\tau}) = \prod_{i=1}^{N_{\tau}^{\text{tr}}} p(x_{\tau}^{\text{tr}}[i]|y_{\tau}^{\text{tr}}[i], \phi_{\tau})$ is the likelihood function for the training data. However, computing the posterior $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)$ in (11) is generally intractable for high dimensional vector ϕ_{τ} .

To address this challenge, we follow VI and introduce a *variational distribution approximation*

$$q(\phi_{\tau}|\varphi_{\tau}) \approx p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi), \quad (12)$$

which depends on a variational parameter vector φ_{τ} . A typical choice is given by the Gaussian mean-field approximation [77] which can be expressed as

$$q(\phi_{\tau}|\varphi_{\tau}) = \mathcal{N}(\phi_{\tau}|\nu_{\tau}, \text{Diag}(\exp(2\rho_{\tau}))), \quad (13)$$

with variational parameter vector $\varphi_{\tau} = [\nu_{\tau}^{\top}, \rho_{\tau}^{\top}]^{\top}$, and the exponent function is applied element-wise. The variational parameter vector includes the mean vector $\nu_{\tau} \in \mathbb{R}^D$ and the vector of the logarithm of the standard deviations $\rho_{\tau} \in \mathbb{R}^D$ for the Gaussian random vector ϕ_{τ} . Note that vector ρ_{τ} models uncertainty in the model parameter space.

To describe VI, we will use the Kullback-Liebler (KL) divergence $\text{KL}(p(z)||q(z))$ [78], which is a measure of the distance between two distributions $p(z)$ and $q(z)$. It is defined as the average of the log-likelihood ratio $\log(p(z)/q(z))$ as

$$\text{KL}(p(z)||q(z)) = \mathbb{E}_{p(z)} \left[\log \left(\frac{p(z)}{q(z)} \right) \right]. \quad (14)$$

VI-based Bayesian learning prescribes that the variational parameter vectors ϕ_{τ} be obtained via the minimization of the KL divergence $\text{KL}(q(\phi_{\tau}|\varphi_{\tau})||p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi))$ between the variational distribution $q(\phi_{\tau}|\varphi_{\tau})$ and the posterior distribution $p(\phi_{\tau}|\mathcal{D}_{\tau}^{\text{tr}}, \xi)$. This problem can be equivalently formulated as the minimization [76], [77]

$$\varphi_{\tau} = \underset{\varphi}{\text{argmin}} F_{\mathcal{D}_{\tau}^{\text{tr}}}(\varphi|\xi), \quad (15)$$

where the *variational free energy* [79] is defined as

$$\begin{aligned} F_{\mathcal{D}_{\tau}^{\text{tr}}}(\varphi_{\tau}|\xi) &= N_{\tau}^{\text{tr}} \mathbb{E}_{q(\phi_{\tau}|\varphi_{\tau})} [\mathcal{L}_{\mathcal{D}_{\tau}^{\text{tr}}}(\phi_{\tau})] + \text{KL}(q(\phi_{\tau}|\varphi_{\tau})||p(\phi_{\tau}|\xi)) \\ &= N_{\tau}^{\text{tr}} L_{\mathcal{D}_{\tau}^{\text{tr}}}(\varphi_{\tau}) + \text{KL}(q(\phi_{\tau}|\varphi_{\tau})||p(\phi_{\tau}|\xi)). \quad (16) \end{aligned}$$

Algorithm 1: Reparametrization Trick [80].

Inputs : $\mathcal{G}(\cdot)$ = a function over vector ϕ_τ
 φ_τ = variational parameter
Parameters : R = ensemble size
Output : $\hat{G}(\varphi_\tau)$ = approximation of $\mathbb{E}_{q(\phi_\tau|\varphi_\tau)}[\mathcal{G}(\phi_\tau)]$

1 **for** $r = 1, \dots, R$ **do**
2 Draw $e_{\tau,r} \sim \mathcal{N}(0, I_D)$
3 $\phi_{\tau,r}(\varphi_\tau, e_{\tau,r}) \leftarrow \nu_\tau + \exp(\varrho_\tau) \odot e_{\tau,r}$
4 **return** $\hat{G}(\varphi_\tau) \leftarrow \frac{1}{R} \sum_{r=1}^R \mathcal{G}(\phi_{\tau,r}(\varphi_\tau, e_{\tau,r}))$

In (16), we have defined as $L_{\mathcal{D}_\tau^{\text{tr}}}(\varphi_\tau)$ the expectation of loss function $\mathcal{L}_{\mathcal{D}_\tau^{\text{tr}}}(\phi_\tau)$ (5) over variational distribution $q(\phi_\tau|\varphi_\tau)$, i.e.,

$$L_{\mathcal{D}_\tau^{\text{tr}}}(\varphi_\tau) = \mathbb{E}_{q(\phi_\tau|\varphi_\tau)}[\mathcal{L}_{\mathcal{D}_\tau^{\text{tr}}}(\phi_\tau)]. \quad (17)$$

In (16), the second summand is a regularizer that restricts the variational distribution to be close to the prior distribution. Note that, if the variational distribution has ability to express the posterior distribution in (11), the minimizer of the problem (15) becomes the Bayesian posterior $p(\phi_\tau|\mathcal{D}_\tau^{\text{tr}}, \xi)$, since the KL divergence $\text{KL}(q(\phi_\tau|\varphi)||p(\phi_\tau|\mathcal{D}_\tau^{\text{tr}}, \xi))$ is minimized (and it equals zero) when the two distributions are the same.

A typical choice for the prior distribution $p(\phi_\tau|\xi)$ is the Gaussian distribution. In this case, we have

$$p(\phi_\tau|\xi) = \mathcal{N}(\phi_\tau|\nu, \text{Diag}(\exp(2\varrho))), \quad (18)$$

which is defined by the hyperparameter vector $\xi = [\nu^\top, \varrho^\top]^\top$, where $\nu \in \mathbb{R}^D$ and $\varrho \in \mathbb{R}^D$ stand for the mean and logarithm of the standard deviation vector of the Gaussian random vector ϕ_τ .

Assuming the Gaussian variational distribution in (13) and the Gaussian prior (18), the regularizer term in (16) can be computed in closed-form as

$$\text{KL}(q(\phi_\tau|\varphi_\tau)||p(\phi_\tau|\xi)) = \frac{1}{2} \sum_{d=1}^D \left(2(\varrho[d] - \varrho_\tau[d]) + \frac{\exp(2\varrho_\tau[d]) + (\nu_\tau[d] - \nu[d])^2}{\exp(2\varrho[d])} - 1 \right),$$

which is a differentiable function for φ_τ .

With these choices of variational posterior and prior, problem (15) can be addressed via gradient-descent methods by using the reparametrization trick [80]. This is done by writing the random model parameter vector $\phi_\tau \sim q(\phi_\tau|\varphi_\tau)$ as $\phi_\tau = \nu_\tau + \exp(\varrho_\tau) \odot e$, with random vector $e \sim \mathcal{N}(0, I_D)$ and \odot being the element-wise multiplication. An estimate of the gradient of the objective (17) using the reparametrization trick is done with the aid of R drawn independently samples of the standard normal Gaussian random vector e , and differentiating the resulting empirical estimate of (17).

Specifically, we estimate the free energy in (16) by replacing the training loss $L_{\mathcal{D}_\tau^{\text{tr}}}(\varphi_\tau)$ with the empirical estimate

$$\hat{L}_{\mathcal{D}_\tau^{\text{tr}}}(\varphi_\tau) = \frac{1}{R} \sum_{r=1}^R \mathcal{L}_{\mathcal{D}_\tau^{\text{tr}}}(\nu_\tau + \exp(\varrho_\tau) \odot e_{\tau,r}), \quad (19)$$

obtained by drawing samples $e_{\tau,r} \sim \mathcal{N}(0, I_D)$ for $r = 1, 2, \dots, R$. This yields the estimated free energy

$$\hat{F}_{\mathcal{D}_\tau^{\text{tr}}}(\varphi_\tau|\xi) = N_\tau^{\text{tr}} \hat{L}_{\mathcal{D}_\tau^{\text{tr}}}(\varphi_\tau) + \text{KL}(q(\phi_\tau|\varphi)||p(\phi_\tau|\xi)). \quad (20)$$

This is a special case of Algorithm 1 with input $\mathcal{G}(\phi_\tau) = \mathcal{L}_{\mathcal{D}_\tau^{\text{tr}}}(\phi_\tau)$. The function (20) can be directly differentiated and used in SGD updates.

Once the variational parameter φ_τ is inferred using Bayesian training, ensemble prediction for a payload data symbol $(y_\tau^{\text{te}}[i], x_\tau^{\text{te}}[i])$ can be obtained via (10) by replacing $p(\phi_\tau|\mathcal{D}_\tau^{\text{tr}}, \xi)$ with $q(\phi_\tau|\varphi_\tau)$ to yield the ensemble predictor

$$p(x_\tau^{\text{te}}[i]|y_\tau^{\text{te}}[i], \varphi_\tau) = \mathbb{E}_{q(\phi_\tau|\varphi_\tau)} [p(x_\tau^{\text{te}}[i]|y_\tau^{\text{te}}[i], \phi_\tau)]. \quad (21)$$

Practically, it uses Monte Carlo sampling with R model vectors, producing the approximated soft predictor $\hat{p}(x_\tau^{\text{te}}[i]|y_\tau^{\text{te}}[i], \varphi_\tau)$ via Algorithm 1 with $\mathcal{G}(\phi_\tau) = p(x_\tau^{\text{te}}[i]|y_\tau^{\text{te}}[i], \phi_\tau)$.

B. Bayesian Meta-Learning

While conventional Bayesian learning assumes that the random model parameter vector ϕ_τ has a fixed prior distribution $p(\phi_\tau|\xi)$ parametrized by a predefined hyperparameter vector ξ , Bayesian meta-learning leverages the stronger assumption that there is a shared prior distribution $p(\phi_\tau|\xi)$ across all frames that can be optimized through a hyperparameter vector ξ .

In this section, we formulate Bayesian meta-learning by following *empirical Bayes* [81], with the aim of selecting a distribution $p(\phi_\tau|\xi)$ that provides a useful prior for the design of the predictor on new frames. Mathematically, Bayesian meta-training optimizes over the hyperparameter vector ξ by addressing the bi-level problem

$$\min_{\xi} \frac{1}{N_{1:t}^{\text{te}}} \sum_{\tau=1}^t N_\tau^{\text{te}} \mathbb{E}_{q(\phi_\tau|\varphi_\tau(\mathcal{D}_\tau^{\text{tr}}|\xi))} [\mathcal{L}_{\mathcal{D}_\tau^{\text{te}}}(\phi_\tau)] \quad (22a)$$

$$\text{s.t. } \varphi_\tau(\mathcal{D}_\tau^{\text{tr}}|\xi) = \underset{\varphi}{\text{argmin}} F_{\mathcal{D}_\tau^{\text{tr}}}(\varphi|\xi), \tau = 1, \dots, t. \quad (22b)$$

Problem (22) chooses the hyperparameter vector ξ that minimizes the average test loss on the meta-training frames $\tau \in \{1, \dots, t\}$ that is obtained with the variational posterior via (15). The subproblems in (22b) correspond to Bayesian learning applied separately to each frame as explained in Section III-A. An illustration of all the quantities involved in problem (22) can be found in Fig. 3 by using the formalism of Bayesian networks [82].

To address problem (22) in a tractable manner, we apply the reparametrization trick for both outer (22a) and inner optimization (22b) by following the same steps described in Section III-A. Details on the optimization can be found in Algorithm 2. In short, the inner loop updates the frame-specific variational parameters φ_τ by minimizing the approximated free energy (20) separately for each frame τ within a mini-batch \mathcal{T} via GD (dashed blue line in Fig. 3(b)). Following [30], [66], the prior's parameter vector ξ plays two roles in the inner loop, namely (i) as the initialization for the inner GD update in Algorithm 2 line 9; and (ii) as the regularizer for the same update via the prior $p(\phi_\tau|\xi)$. The outer optimization (22a) is addressed via SGD

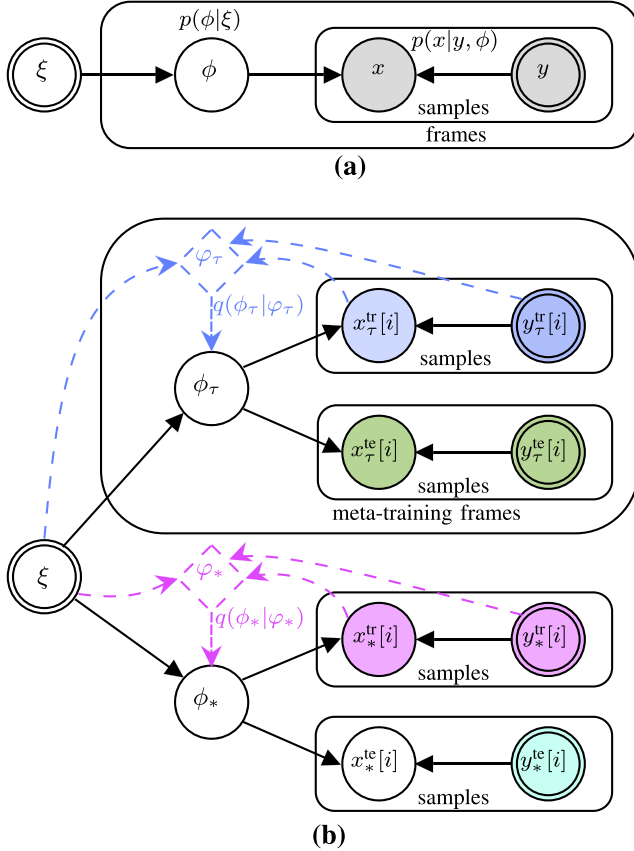


Fig. 3. Probabilistic graphical model (Bayesian network) [82] for Bayesian meta-learning. Circles represent random variables; double-lined circles represent deterministic variables or (hyper)parameters; shaded circles represent observations; dashed diamonds represent variational parameter vectors; and plaques indicate multiple instances (the outer plaques represent frames, whereas the inner represent multiple sample, e.g., symbols across time): (a) High level representation, assuming a prior $p(\phi|\xi)$ and predictor $p(x|y, \phi)$; (b) Model using the train/test splits, with variational inference $q(\phi_\tau|\varphi_\tau) \approx p(\phi_\tau|\mathcal{D}_\tau^{\text{tr}}, \xi)$ indicated as dashed arrows.

to minimize the average log-likelihood for test set using Algorithm 1 with $\mathcal{G}(\phi_\tau) = \mathcal{L}_{\mathcal{D}_\tau^{\text{te}}}(\phi_\tau)$, shown as dashed green line in Fig. 3(b).

After obtaining meta-trained hyperparameter ξ , meta-testing takes place, starting with the adaptation of the variational parameter $\varphi_*(\mathcal{D}_*^{\text{tr}}|\xi)$ via (22b) using the available pilot data $\mathcal{D}_*^{\text{tr}}$ at the current frame, to obtain ensemble prediction

$$p(x_*^{\text{te}}[i]|y_*^{\text{te}}[i], \varphi_*) = \mathbb{E}_{q(\phi_*|\varphi_*)} [p(x_*^{\text{te}}[i]|y_*^{\text{te}}[i], \phi_*)], \quad (23)$$

as done in (21). Bayesian meta-learning is illustrated comparatively to meta-learning in Fig. 4.

C. Computational Complexity

We now briefly elaborate on the complexity of meta-learning by analyzing the complexity of meta-training and of meta-testing. To this end, let us define as C the complexity of obtaining the probability $p(x|y, \phi)$ for a data sample (y, x) . This baseline complexity depends on the model dimensionality, and it accounts for the amount of time needed to carry out the forward pass on the neural network implementing the model $p(x|y, \phi)$.

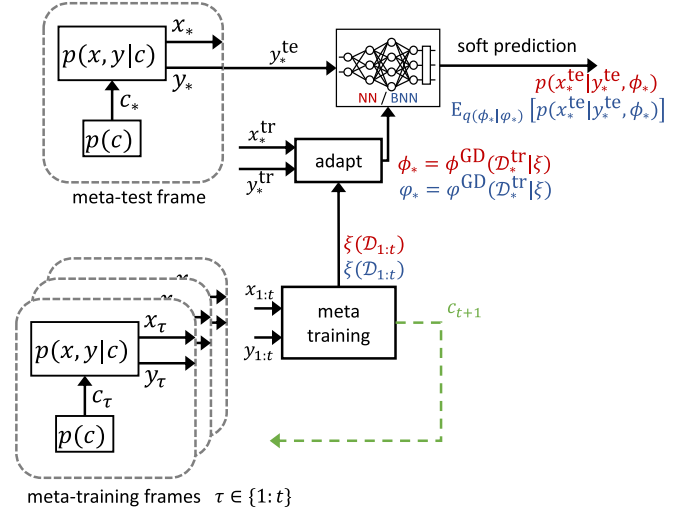


Fig. 4. Bayesian meta-learning (blue) as compared to frequentist meta-learning (red). The frequentist predictor uses a single predictor, depicted as a neural network (NN), whereas Bayesian meta-learning uses an ensemble of predictors, e.g., a Bayesian NN (BNN). The dashed line represents the operation of the active meta-learning introduced in Section IV. The data for each frame is generated by following the distribution $p(x, y|c) = p(x)p(y|x, c)$, with input distribution $p(x)$ and conditional distribution $p(y|x, c)$ for channel state c .

Algorithm 2: Bayesian Meta-Training.

Inputs	: $\mathcal{D}_{1:t}$ = labeled data sets of t meta-training frames
Parameters	: B = number of frames per meta-update batch I = number of inner update steps η, κ = inner/outer updates learning rates
Output	: ξ = learned hyperparameter vector

```

1 initialize  $\xi$ 
2 while meta-learning not done do
3    $\mathcal{T} \leftarrow$  random batch of  $B$  frames
4   for  $\tau \in \mathcal{T}$  do
5     randomly divide  $\mathcal{D}_\tau = \{\mathcal{D}_\tau^{\text{tr}}, \mathcal{D}_\tau^{\text{te}}\}$ 
6      $\triangleleft$  frame-specific update  $\triangleright$ 
7      $\varphi_\tau^{(0)} \leftarrow \xi$ 
8     for  $i = 1, 2, \dots, I$  inner update steps do
9        $\varphi_\tau^{(i)} \leftarrow \varphi_\tau^{(i-1)} - \frac{\eta}{N_\tau^{\text{tr}}} \nabla_{\varphi_\tau} \hat{F}_{\mathcal{D}_\tau^{\text{tr}}}(\varphi_\tau^{(i-1)}|\xi)$  using (20)
10       $\varphi^{\text{GD}}(\mathcal{D}_\tau^{\text{tr}}|\xi) \leftarrow \varphi_\tau^{(I)}$ 
11     $\triangleleft$  meta-update  $\triangleright$ 
12     $\xi \leftarrow \xi - \kappa \frac{1}{N_\tau^{\text{te}}} \sum_{t \in \mathcal{T}} N_\tau^{\text{te}} \nabla_\xi \hat{L}_{\mathcal{D}_\tau^{\text{te}}}(\varphi^{\text{GD}}(\mathcal{D}_\tau^{\text{tr}}|\xi))$ 
13 return  $\xi$ 

```

Accordingly, as seen in Table II, the per-data point complexity of meta-testing equals C for frequentist learning, and CR^{te} for Bayesian learning, where R^{te} is the size of the ensemble used for inference.

The complexity of computing the first-order gradient via backpropagation per-sample is given by $G_1 C$, with G_1 being a constant in the range between 2 and 5 [83], [84]. Furthermore, computing the Hessian-vector product (HVP) has a complexity of the order $G_2 G_1 C$, where constant G_2 is also between 2 to 5 [19, Appendix A], [85, Appendix C]. Assume that all tasks have data sets of equal size, i.e., $N_\tau^{\text{tr}} = N^{\text{tr}}$ and $N_\tau^{\text{te}} = N^{\text{te}}$ for any task τ . Therefore, for each meta-training iteration, for a batch of

TABLE II
COMPUTATIONAL COMPLEXITY OF FREQUENTIST AND BAYESIAN
META-LEARNING. (SEE TEXT IN SECTION III-C FOR DETAILS)

	inference [per-test sample]	meta-training [per-meta-iteration]
frequentist	C	$BG_1C(IN^{\text{tr}}(G_2 + 1) + N^{\text{te}})$
Bayesian	CR^{te}	$BG_1C(IN^{\text{tr}}(G_2 + 1)R^{\text{tr}} + N^{\text{te}}R^{\text{te}})$

B tasks with I local updates, the complexity of the frequentist meta-update (8) is of the order

$$B \left(\underbrace{IN^{\text{tr}}G_1C}_{\text{frame-specific update}} + \underbrace{IN^{\text{tr}}G_2G_1C}_{\text{HVPs in meta-update}} + \underbrace{N^{\text{te}}G_1C}_{\text{gradient in meta-update}} \right). \quad (24)$$

For Bayesian meta-learning, the complexity increases linearly with the training ensemble size that is used for estimating the loss functions in (22a) and (22b). Note that the impact of the size R^{tr} of the training ensemble used for meta-training is different from the size R^{te} used for inference, as the first determines the variance of the stochastic loss functions, while the latter determines the quality of Bayesian prediction (see, e.g., [59] and references therein). Ignoring the constant cost of differentiating the KL term in the free energy and for sampling from the Gaussian distribution, the complexity analysis is summarized in Table II.

IV. BAYESIAN ACTIVE META-LEARNING

In the previous sections, we have considered a passive meta-learning setting in which the meta-learner is given a number of meta-training data sets, each corresponding to a different channel state c . In this section, we study the situation in which the meta-learner has access to a simulator that can be used to generate random data sets for any channel state c via the channel $p(y|x, c)$. The problem of interest is to minimize the use of the simulator by actively selecting the channels $\{c_\tau\}$ for which meta-training data is generated. To this end, we devise a sequential approach, whereby the meta-learner optimizes the next channel state c_{t+1} , given all t meta-training data sets of frames $\tau = 1, \dots, t$.

At the core of the proposed active meta-learning strategy, are mechanisms used by the meta-learner to discover model parameter vectors ϕ that have been underexplored so far, and to relate model parameter vector ϕ to a channel state.

A. Active Selection of Channel States

After having collected t meta-training data sets $\mathcal{D}_{1:t} = \{\mathcal{D}_\tau\}_{\tau=1}^t$, the proposed active meta-learning scheme selects the next channel state, c_{t+1} , to use for the generation of the $(t+1)$ -th meta-training data set \mathcal{D}_{t+1} . We adopt the general principle of maximizing the amount of “knowledge” that can be extracted from the data set associated with selected channel c_{t+1} , when added to the t available data sets $\mathcal{D}_{1:t}$. This is done via the following three steps: (i) searching in the space of model parameter vectors for a vector ϕ_{t+1} that is most “surprising” given the available meta-training data $\mathcal{D}_{1:t}$; (ii) translating the selected model parameter vector ϕ_{t+1} into a channel c_{t+1} ; and

(iii) generating data set \mathcal{D}_{t+1} by using the simulator with input c_{t+1} .

As illustrated in Fig. 5, in step (i), we adopt the scoring function introduced in [37], i.e.,

$$s_t(\phi|\varphi_{1:t}) := -\log \left(\frac{1}{t} \sum_{\tau=1}^t q(\phi|\varphi_\tau) \right) \quad (25)$$

in order to select the next model parameter vector as

$$\phi_{t+1} = \underset{\phi}{\operatorname{argmax}} s_t(\phi|\varphi_{1:t}). \quad (26)$$

The criterion (25) measures how incompatible model parameter vector ϕ is with the available data $\mathcal{D}_{1:t}$. In fact, by the derivations in the previous section: the mixture of variational distributions $\frac{1}{t} \sum_{\tau=1}^t q(\phi|\varphi_\tau)$ quantifies how likely a vector ϕ is on the basis of the data $\mathcal{D}_{1:t}$ (Fig. 5(b)); and the negative logarithm in (25) evaluates the information-theoretic “surprise” associated with that mixture. Problem (26) can be addressed either by grid search for low-dimensional model parameter space, or by using gradient ascent due to the differentiability nature of the scoring function (25), as illustrated in Fig. 5(c).

In step (ii), we need to convert the selected model parameter vector ϕ_{t+1} , i.e., the outcome of (26), into channel state c_{t+1} . We choose the channel state c_{t+1} that minimizes the cross entropy loss when evaluated at ϕ_{t+1} , i.e.,

$$c_{t+1} \in \underset{c}{\operatorname{argmin}} \{ \mathcal{L}_p(\phi_{t+1}|c) = \mathbb{E}_{p(x,y|c)} [-\log p(x|y, \phi_{t+1})] \}, \quad (27)$$

where we set $p(x, y|c) = p(x)p(y|x, c)$, with $p(x)$ being some fixed distribution and $p(y|x, c)$ being the distribution of the output of the simulator. In (27), we have emphasized that there may be more than one solution to the problem. The rationale behind problem (27) is that data generated from the distribution $p(x, y|c_{t+1})$ can be interpreted as being the most compatible with the demodulator $p(x|y, \phi_{t+1})$, where compatibility is measured by the average of the cross entropy $\mathbb{E}_{p(y|c_{t+1})} [\mathbb{H}(p(x|y, c_{t+1}), p(x|y, \phi_{t+1}))]$.

We emphasize that the proposed approach is different from the methodology introduced by [37], which uses another variational distribution in problem (22). In our experiments, we found the method in [37] to be ineffective and complex for the problem under study here. The main issue appears to be overfitting for the additional variational distribution, which is overcome by leveraging the availability of the channel simulator implementing the model $p(y|x, c)$.

In some models, problem (27) can be solved analytically. For more complex models, SGD-based approaches can be used, either by differentiating an estimate of the loss in a manner similar to the discussion in Section III (i.e., Algorithm 1 with $\mathcal{G}(\phi_{t+1}) = \mathcal{L}_p(\phi_{t+1}|c)$), or by directly estimating its gradient [86].

Finally, in step (iii), meta-training data set $\mathcal{D}_{t+1} = \{(y_{t+1}[i], x_{t+1}[i])\}_{i=1}^{N_{t+1}}$ is generated using the simulator in an i.i.d. fashion following the distribution

$$\prod_{i=1}^{N_{t+1}} p(x_{t+1}[i]) p(y_{t+1}[i]|x_{t+1}[i], c_{t+1}). \quad (28)$$

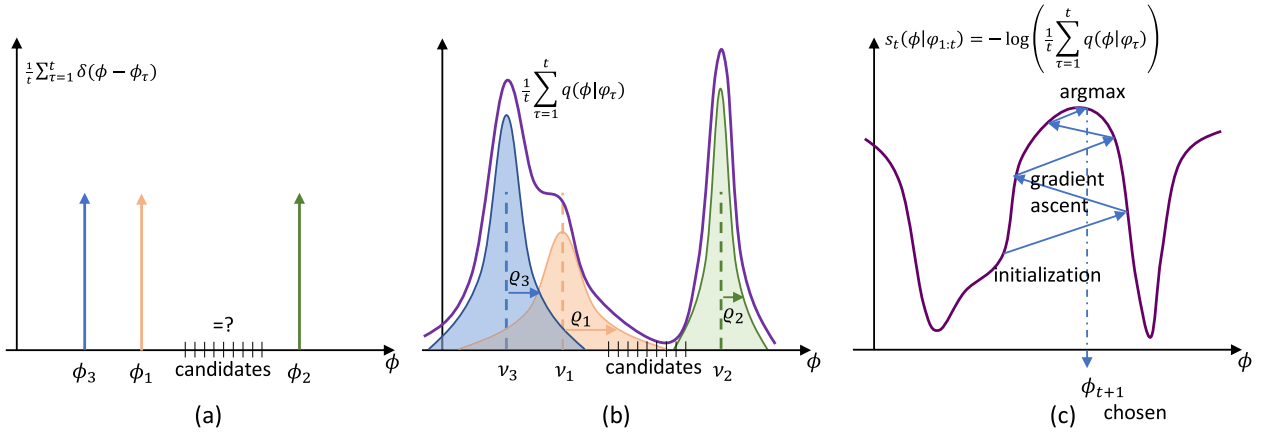


Fig. 5. Illustration of how model parameter vectors are scored to enable active meta-learning provided $t = 3$ meta-training sets. (a) Frequentist meta-learning relies on point estimates, and is hence unable to score as-of-yet unexplored model parameters; (b) Bayesian meta-learning can associate a score to each model parameter vector ϕ based on the variational distributions $\{q(\phi|\varphi_\tau)\}$ evaluated in the previously observed frames $\tau = 1, \dots, t$; (c) The scoring function can be maximized to obtain the next model parameter vector ϕ_{t+1} as the most “surprising” one.

Algorithm 3: Bayesian Active Meta-Training.

Inputs : $p(y|x, c) =$ channel model
 $p(x) =$ generative symbols distribution
Parameters : $t_{\text{init}} =$ number of prior-based first frames
Output : $\xi =$ shared hyperparameter vector

```

1  $\triangleleft$  Generate initial experience  $\triangleright$ 
2 for  $t = 1, 2, \dots, t_{\text{init}}$  do
3   Draw using the prior  $c_t \sim p(c_t)$ 
4   Acquire data  $\mathcal{D}_t \sim \prod_{i=1}^{N_t} p(x_t[i])p(y_t[i]|x_t[i], c_t)$ 
5 while data acquisition not done do
6    $\xi \leftarrow$  BayesianMetaTraining( $\mathcal{D}_{1:t}$ ) using Algorithm 2
7    $\triangleleft$  frame-specific update with updated  $\xi$   $\triangleright$ 
8   for  $\tau = 1, 2, \dots, t$  do
9      $\varphi_\tau \leftarrow \varphi^{\text{GD}}(\mathcal{D}_\tau^x|\xi)$  using (20)
10   $\triangleleft$  step (i), choose surprising model parameter  $\triangleright$ 
11   $\phi_{t+1} = \text{argmax}_\phi s_t(\phi|\varphi_{1:t})$  using (25)
12   $\triangleleft$  step (ii), choose next channel  $\triangleright$ 
13   $c_{t+1} \in \text{argmin}_c \mathcal{L}_p(\phi_{t+1}|c)$  as in (27)
14   $\triangleleft$  step (iii), generate data set  $\triangleright$ 
15   $\mathcal{D}_{t+1} \sim \prod_{i=1}^{N_{t+1}} p(x_{t+1}[i])p(y_{t+1}[i]|x_{t+1}[i], c_{t+1})$ 
16   $t \leftarrow t + 1$ 
17 return  $\xi$ 

```

As a final note, we adopt the proposal in [37] of implementing active selection only after $t_{\text{init}} > 1$ channel states that are generated at random, as a means to avoid being overconfident at early stages. The overall proposed Bayesian active meta-learning scheme is summarized in Algorithm 3.

V. EXPERIMENTS

In this section, we present experimental results to evaluate the performance of Bayesian meta-learning for demodulation/equalization.

A. Performance Metrics

Apart from the standard measures of symbol error rate (SER) and mean squared error (MSE), we will also evaluate metrics quantifying the performance in terms of the reliability of the

confidence measures provided by the predictor. While such measures can be defined for both classification and regression problems, we will focus here on uncertainty quantification for demodulation via calibration metrics (see [87] for discussion on regression).

As discussed in the previous sections, for a new frame, we need to make a prediction for the payload symbols $\{y_*^{\text{te}}[i]\}_{i=1}^{N_*^{\text{te}}}$ via the demodulator $p(x_*^{\text{te}}[i]|y_*^{\text{te}}[i], \phi_*)$ for frequentist meta-learning (9), or $p(x_*^{\text{te}}[i]|y_*^{\text{te}}[i], \varphi_*)$ for Bayesian meta-learning (21). The *confidence level* assigned by the model to the hard predicted symbol

$$\hat{x}_*^{\text{te}}[i] = \underset{x \in \mathcal{X}}{\text{argmax}} p(x|y_*^{\text{te}}[i], \theta) \quad (29)$$

given the received symbol $y_*^{\text{te}}[i]$, can be defined as the corresponding probability [7]

$$\hat{p}[i] = \max_{x \in \mathcal{X}} p(x|y_*^{\text{te}}[i], \theta) = p(\hat{x}_*^{\text{te}}[i]|y_*^{\text{te}}[i], \theta), \quad (30)$$

where we have $\theta = \phi_*$ for frequentist meta-learning and $\theta = \varphi_*$ for Bayesian meta-learning. *Perfect calibration* [7] can be defined as the condition where symbols that are assigned a confidence level $\hat{p}[i]$ are also characterized by a probability of correct detection equal to p .

Two standard means of quantifying the extent to which the perfect calibration is satisfied are reliability diagrams [88] and expected calibration error (ECE) [7]. To introduce them, the probability interval $[0, 1]$ is first divided into M equal length intervals, with the m -th interval $(\frac{m-1}{M}, \frac{m}{M}]$ referred to as the m -th bin henceforth. Let us denote as \mathcal{B}_m the subset of the payload data symbol indices whose associated confidence level $\hat{p}[i]$ lie within the m -th bin, i.e.,

$$\mathcal{B}_m = \{i | \hat{p}[i] \in (\frac{m-1}{M}, \frac{m}{M}]\}, \text{ with } i = 1, 2, \dots, N_*^{\text{te}}. \quad (31)$$

Note this is a partition of the data set $\mathcal{D}_*^{\text{te}}$ since we have $\bigcup_{m=1}^M \mathcal{B}_m = \{i = 1, 2, \dots, N_*^{\text{te}}\}$ and $\mathcal{B}_m \cap \mathcal{B}_{m'} = \emptyset$ for any $m' \neq m$.

The within-bin empirical average accuracy of the predictor for the m -th bin is defined as

$$\text{acc}(\mathcal{B}_m) = \frac{1}{|\mathcal{B}_m|} \sum_{i \in \mathcal{B}_m} \mathbf{1}(\hat{x}_*^{\text{te}}[i] = x_*^{\text{te}}[i]), \quad (32)$$

with $\mathbf{1}(\cdot)$ being indicator function and $|\mathcal{B}_m|$ denoting the number of total samples in \mathcal{B}_m . The within-bin empirical average confidence of the predictor for the m -th bin is

$$\text{conf}(\mathcal{B}_m) = \frac{1}{|\mathcal{B}_m|} \sum_{i \in \mathcal{B}_m} \hat{p}[i]. \quad (33)$$

A perfectly calibrated demodulator $p(x|y, \theta)$ would have $\text{acc}(\mathcal{B}_m) = \text{conf}(\mathcal{B}_m)$ for all $m \in \{1, \dots, M\}$ in the limit of a sufficiently large payload data set, i.e., $N_*^{\text{te}} \rightarrow \infty$.

Reliability diagrams plot the accuracy $\text{acc}(\mathcal{B}_m)$ and the confidence $\text{conf}(\mathcal{B}_m)$ over the binned probability interval $[0, 1]$. Ideal calibration would yield $\text{acc}(\mathcal{B}_m) = \text{conf}(\mathcal{B}_m)$ in a reliability plot. If in the m -th bin, the empirical accuracy and empirical confidence are different, the predictor is considered to be over-confident when $\text{conf}(\mathcal{B}_m) > \text{acc}(\mathcal{B}_m)$, and under-confident when $\text{conf}(\mathcal{B}_m) < \text{acc}(\mathcal{B}_m)$.

The ECE quantifies the overall amount of miscalibration by computing the weighted average of the differences between within-bin accuracy and within-bin confidence levels across all M bins, i.e.,

$$\text{ECE} = \frac{1}{N_*^{\text{te}}} \sum_{m=1}^M |\mathcal{B}_m| |\text{acc}(\mathcal{B}_m) - \text{conf}(\mathcal{B}_m)|. \quad (34)$$

B. Frequentist and Bayesian Meta-Learning for Demodulation

For the first set of experiments, we focus on a demodulation problem at the symbol level in the presence of transmitter I/Q imbalance [89], [90], as considered also in [19]. The main reason for this choice is that channel decoding typically requires a hard decision on the transmitted codeword, whose accuracy can be validated via a cyclic redundancy check. In contrast, demodulation is usually a preliminary step at the receiver side, and downstream blocks, such as channel decoding, expect soft inputs that are well calibrated. For each frame τ , the transmitted symbols $x_\tau[i]$ are drawn uniformly at random from the 16-QAM constellation $\mathcal{X} = 1/\sqrt{10}(\{\pm 1, \pm 3\} + j\{\pm 1, \pm 3\})$. The received symbol $y_\tau[i] \in \mathcal{Y} = \mathbb{C}$ is given as

$$y_\tau[i] = h_\tau f_{\text{IQ}, \tau}(x_\tau[i]) + z_\tau[i], \quad (35)$$

for a unit energy fading channel coefficient h_τ , where the additive noise is $z_\tau[i] \sim \mathcal{CN}(0, \text{SNR}^{-1})$ for some signal-to-noise ratio (SNR) level SNR, and the I/Q imbalance function [91] $f_{\text{IQ}, \tau} : \mathcal{X} \rightarrow \tilde{\mathcal{X}}_\tau$ is

$$f_{\text{IQ}, \tau}(x_\tau[i]) = \bar{x}_{\text{I}, \tau}[i] + j\bar{x}_{\text{Q}, \tau}[i]$$

$$\begin{bmatrix} \bar{x}_{\text{I}, \tau}[i] \\ \bar{x}_{\text{Q}, \tau}[i] \end{bmatrix} = \begin{bmatrix} 1 + \epsilon_\tau & 0 \\ 0 & 1 - \epsilon_\tau \end{bmatrix} \begin{bmatrix} \cos \delta_\tau & -\sin \delta_\tau \\ -\sin \delta_\tau & \cos \delta_\tau \end{bmatrix} \begin{bmatrix} x_{\text{I}, \tau}[i] \\ x_{\text{Q}, \tau}[i] \end{bmatrix}, \quad (36)$$

which depends on the imbalance parameters ϵ_τ and δ_τ . In (35), $x_{\text{I}, \tau}[i]$ and $x_{\text{Q}, \tau}[i]$ refer to the real and imaginary parts of the modulated symbol $x_\tau[i]$; and $\bar{x}_{\text{I}, \tau}[i]$ and $\bar{x}_{\text{Q}, \tau}[i]$ stand for the

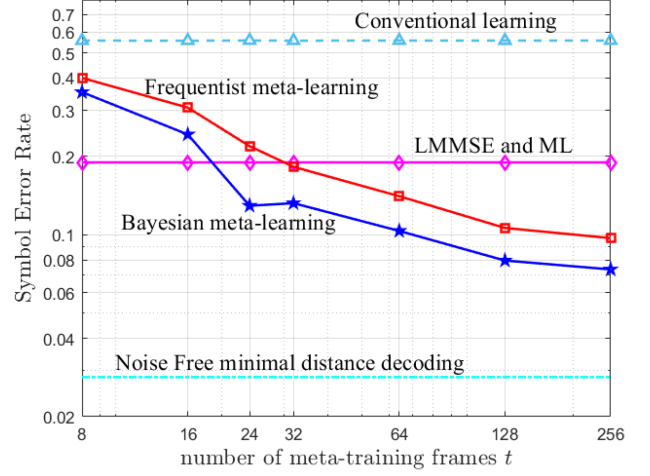


Fig. 6. Symbol error rate (SER) as a function of the number t of meta-training frames with 16-QAM, Rayleigh fading, and I/Q imbalance for $N_\tau^{\text{tr}} = 4$, $N_\tau^{\text{tr}} = 8$. The symbol error rate is averaged over by $N_*^{\text{te}} = 4000$ data symbols and 50 meta-test frames with ensemble of size 100.

real and imaginary parts of the transmitted symbol $f_{\text{IQ}, \tau}(x_\tau[i])$. Note that the constellation $\tilde{\mathcal{X}}_\tau$ of the transmitted symbols $\bar{x}_\tau[i]$ is also composed of 16 points via (36).

By (35) and (36), the channel state c_τ consists of the tuple: (a) amplitude imbalance factor $\epsilon_\tau \in [0, 0.15]$; (b) phase imbalance factor $\delta_\tau \in [0, 15^\circ]$; and (c) channel realization $h_\tau \in \mathbb{C}$. All of the variables are drawn i.i.d. across different frames and are fixed during each frame. We consider the channel state distribution for frame τ as

$$p(c_\tau) = \text{Beta}\left(\frac{\epsilon_\tau}{0.15} \middle| 5, 2\right) \text{Beta}\left(\frac{\delta_\tau}{0.15^\circ} \middle| 5, 2\right) \mathcal{CN}(h_\tau | 0, 1). \quad (37)$$

We set our base learner to be a multi-layer fully-connected neural network (3) with $L = 5$ layers. The real and imaginary parts of the input $y[i] \in \mathbb{C}$ are treated as a vector in \mathbb{R}^2 , which is fed to layers with 10, 30, and 30 neurons, all with ReLU activations, while the last linear layer implements a softmax function that produces probabilities for the 16QAM constellation points.

To address the ability of meta-learning to adapt the demodulator using only few pilots, we set the number of pilots as $N_\tau^{\text{tr}} = 4$ during meta-training and $N_\tau^{\text{tr}} = 8$ for meta-testing [19]. Fig. 6 shows the SER as a function of the number of total meta-training frames t . Since only half of the constellation points are available as pilots during meta-test ($N_\tau^{\text{tr}} = 8$ different symbols out of 16), conventional learning cannot obtain a SER lower than of 0.5. In fact, conventional learning performs worse than a standard model-based receiver applying linear minimal mean square error (LMMSE), followed by maximum likelihood (ML) demodulation, while disregarding the presence of I/Q imbalance function f_{IQ} . Both meta-learning schemes are clearly superior to conventional learning and to the mentioned model-based solution, showing that useful knowledge has been transferred from previous frames to a new frame. Furthermore, Bayesian meta-learning obtains a slightly lower SER as compared to frequentist meta-learning. This advantage stems from the capacity of ensemble predictors to implement more complex decision boundaries [59].




Fig. 7. Expected calibration error (ECE) over meta-test data $\mathcal{D}_*^{\text{te}}$ as a function of the number t of meta-training frames, for the same setting as in Fig. 6.

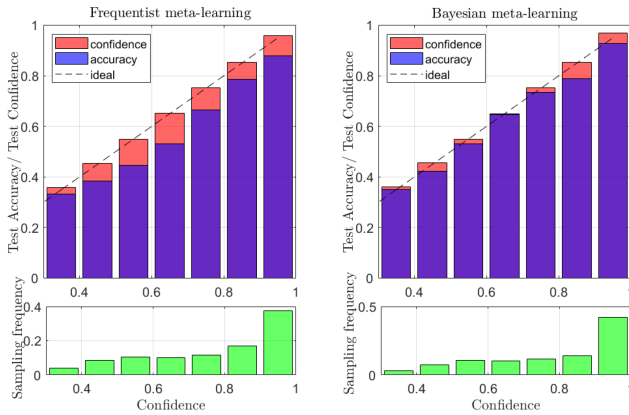


Fig. 8. Reliability diagrams (top) for frequentist meta-learning (left) and Bayesian meta-learning (right) with SNR = 18 dB, using $t = 16$ meta-training frames and predictions averaged over 50 meta-test frames. Frequentist meta-learning tends to be over-confident, whereas the Bayesian soft predictions are better matched to the true accuracy. The bottom figure shows the histogram of $|\mathcal{B}_m|/N$ of prediction over $M = 10$ bins. Full details in Appendix A.

To gain insights into the reliability of the uncertainty quantification provided by the demodulator, we use the metrics defined in Section V-A by setting the total number of bins to $M = 10$. We plot the ECE as a function of the number of total meta-training frames t in Fig. 7. Bayesian meta-learning is seen to achieve a lower ECE than frequentist meta-learning, indicating that Bayesian meta-learning provides more reliable estimates of uncertainty. Furthermore, the increase in ECE as the number t of available meta-training frames increases may be interpreted as a consequence of meta-overfitting [92]. This suggests that meta-learning may be considered as complete after a number of frames that depends on the complexity of the propagation environment. In practice, this can be assessed by evaluating the performance of the demodulator on pilots (see the online strategy in [19] for further discussion on this point).

To further elaborate on the quality of uncertainty quantification, Fig. 8 depicts reliability diagrams for frequentist and Bayesian meta-learning. The within-bin accuracy levels $\text{acc}(\mathcal{B}_m)$ in (32) and the within-bin empirical confidence

$\text{conf}(\mathcal{B}_m)$ in (33) are depicted as dark (blue) and light (red) bars, respectively. Frequentist meta-learning is observed to produce generally over-confident predictions, while Bayesian meta-learning provides better calibrated predictions with well-matching confidence and accuracy levels.

C. Bayesian Active Meta-Learning for Equalization

In this sub section, we illustrate the operation of active meta-learning by investigating a single-input multiple-output (SIMO) Rayleigh block fading real channel model. At frame τ , the modulator uses a 4-PAM to produce symbols $x_\tau[i], i = 1, 2, \dots, N_\tau$, taken uniformly from the set $\mathcal{X} \in 1/\sqrt{5}\{-3, -1, +1, +3\}$. Given channel state $c_\tau = [c_\tau^0, c_\tau^1]^\top \in \mathbb{R}^2$, the i -th channel output symbol $y_\tau[i] \in \mathbb{R}^2$ for $i = 1, 2, \dots, N_\tau$ is defined as the two-dimensional real vector

$$y_\tau[i] = c_\tau x_\tau[i] + z_\tau[i], \quad (38)$$

where both the additive noise $z_\tau[i] \sim \mathcal{N}(0, \frac{1}{2\text{SNR}} I_2)$ and the normalized real block fading coefficients $c_\tau \sim p(c) = \mathcal{N}(c|0, I_2)$ are i.i.d. We adopt the linear equalizer

$$\hat{x}_\tau[i] = \phi_\tau^\top \cdot y_\tau[i] \quad (39)$$

with linear equalizer weight vector $\phi_\tau = [\phi_\tau^0, \phi_\tau^1]^\top \in \mathbb{R}^2$. To obtain a soft equalization, we account for a precision level β via the conditional distribution

$$p(x_\tau[i]|y_\tau[i], \phi_\tau) = \mathcal{N}(\phi_\tau^\top \cdot y_\tau[i], \beta^{-1}). \quad (40)$$

The next model parameter ϕ_{t+1} is chosen to maximize the scoring function as in (26) by restricting the optimization to the domain $\|\phi\| \leq 1$. This restricted optimization domain is selected in order to match the circular symmetry of the problem. Furthermore, the corresponding next channel state c_{t+1} is selected by tackling problem (27), which amounts to the minimization

$$c_{t+1}(\phi) \in \underset{c}{\text{argmin}} E_{p(x)p(y|x,c)} [-\log p(x|y, \phi)] \quad (41a)$$

$$\begin{aligned} &= \underset{c}{\text{argmin}} E_{p(x)p(z)} \left[\frac{\beta}{2} (x - \phi^\top \cdot (cx + z))^2 \right] \\ &= \underset{c}{\text{argmin}} E_{p(x)p(z)} \left[\frac{\beta}{2} ((1 - \phi^\top \cdot c)x - \phi^\top z)^2 \right] \\ &= \{c | \phi^\top \cdot c = 1\}. \end{aligned} \quad (41b)$$

In the set of solutions of problem (41b), we select the minimum-norm solution $c_{t+1} = \phi_{t+1} / \|\phi_{t+1}\|^2$. This way, the selected channel focuses on the more challenging low-SNR regime. Details of this experiment are provided in Appendix A.

Fig. 9 illustrates the scoring function (25) used to select the next model parameter ϕ_{t+1} as a heat map in the space of model parameter ϕ . Specifically, the figure shows the scoring functions after observing $t = 4$ and $t = 5$ meta-training frames. The optimized next model parameter vector ϕ_{t+1} (26) is shown as a star, while the previously selected model parameter vectors $\phi_{1:t}$ are shown as squares. Fig. 9 illustrates how active meta-learning efficiently explores the model parameter space. It does so by avoiding the inclusion of channel states that are similar to those already considered (i.e., the squares in the figure). This

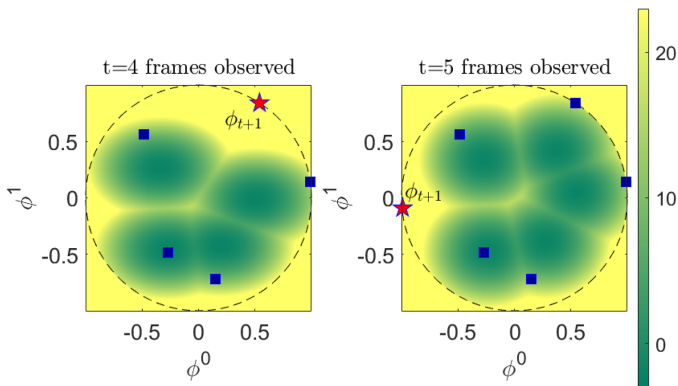


Fig. 9. Scoring function (25) used by Bayesian active meta-learning to select the next model parameter vector ϕ_{t+1} at the fourth and fifth iterations. The scoring function is shown as a heat map over the two dimensional space of the model parameter vector ϕ for the example detailed in Section V-C.




Fig. 10. Meta-test mean squared error (MSE) loss as function of the number of frames t . Bayesian active meta-training is able to achieve lower meta-test loss levels by using fewer meta-training tasks t . Solid lines are the mean test loss over 100 channel states. The confidence levels account for one standard deviation.

way, the model parameter space can be covered with fewer meta-training frames t , leading to a larger frame efficiency of active meta-learning.

Finally, to numerically validate the advantage of active meta-learning, we plot the meta-test MSE loss in Fig. 10 for both passive and active Bayesian meta-learning versus the number of frames t . For passive meta-learning, we have generated random channel realizations by drawing from the distribution $p(c) = \mathcal{N}(c|0, I_2)$. We have repeated the experiment 100 times, and show the confidence interval of one standard deviation for the meta-test loss. The results in the figure confirm that active meta-learning requires far fewer meta-training frames. Furthermore, the increased randomness of passive meta-learning is due to the random selection of channel states at each iteration.

VI. CONCLUSION

In this paper, we have introduced tools for reliable and efficient AI in communication systems via Bayesian meta-learning. Bayesian learning has the advantage of producing well-calibrated decisions whose confidence levels are a close

match for the corresponding test accuracy. This property facilitates monitoring of the quality of the outputs of an AI module. Meta-learning optimizes models that can quickly adapt based on few pilots, producing sample-efficient AI solutions. This paper has focused on the application of Bayesian meta-learning to the basic problems of demodulation/equalization from few pilots. We have demonstrated via experiments that the demodulator/equalizer obtained via Bayesian meta-learning not only achieves a higher accuracy, but it also enjoys better calibration performance than its standard frequentist counterpart. Furthermore, thanks to meta-learning, such performance levels can be obtained based on a limited number of pilots per frame.

To reduce the number of past frames required by meta-learning, we have also introduced Bayesian active meta-learning, which leverages the uncertainty estimates produced by Bayesian learning to actively explore the space of channel conditions. We have shown via numerical results that active meta-learning can indeed significantly speed up meta-training in terms of number of frames.

Future work may consider a fully Bayesian meta-learning implementation that also accounts for uncertainty at the level of hyperparameters (see, e.g., [73] and references therein). This may be particularly useful in the regime of low number of frames. Another direction for research would be to investigate different scoring functions for active meta-learning (see, e.g., [39]). A study on the impact of well-calibrated decisions obtained via Bayesian learning on downstream blocks at the receiver, such as channel decoding, is also of interest. Finally, the proposed tools may find applications to other problems in communications, such as power control [25] and channel coding [28], [93].

APPENDIX A

EXPERIMENTS DETAILS

Table III summarizes the parameters used for the numerical experiments in Section V for demodulation and equalization. Throughout the simulations, we used PyTorch [94] adopting autograd's option `create_graph = True` to allow the computational graph to calculate second-order derivatives.

For the demodulation problem in Section V-B (Figs. 6–8), the complex input space $\mathcal{Y} = \mathbb{C}$ is treated as a two-dimensional real vector space \mathbb{R}^2 when is fed into the neural network demodulator. The KL term in (20) is suppressed by a multiplicative coefficient of 0.1, as a means to emphasize the average log-likelihood term should have over the prior. This is an approach known as generalized Bayesian inference [79], [95]. To handle the discrepancy in the number of pilots for adaptation during meta-training and meta-testing, i.e., $N_*^{\text{tr}} > N_*^{\text{tr}}$, we consider the following strategy akin to burn-in phase [49] during meta-testing as done in [19]: (i) start with I updates using learning rate η utilizing N_*^{tr} pilots among the available N_*^{tr} pilots; (ii) then, additional $I_* - I$ updates are performed with reduced learning rate (5% of the original learning rate) with all available N_*^{tr} pilots. This strategy becomes particularly useful in practical scalable systems in which the number of pilots may change depending on the deployment environments.

As for the equalization setting in Section V-C (Figs. 9–10), we observe that reinitializing the hyperparameter ξ to a random

TABLE III
PARAMETERS FOR THE DEMODULATION AND EQUALIZATION META-LEARNING

Description	Symbol	Demodulation (Sec. V-B)	Equalization (Sec. V-C)
Signal to noise ratio [dB]	SNR	18	6
Modulation	-	16-QAM	4-PAM
Neural network input dimension	integer	$\dim(\mathbb{C}) = 2$	$\dim(\mathbb{R}^2) = 2$
Neural network layers size [{hidden},output]	neurons per layer	[10, 30, 30, 16]	[, 1]
Activation of hidden layers	-	ReLU	NA
Activation of last layer	-	softmax	no activation
Meta-training frames mini-batch size	B	16	full batch (t)
Frame-specific learning rate	η	10^{-1}	$2 \cdot 10^{-3}$
Meta-learning rate	κ	10^{-3}	$5 \cdot 10^{-2}$
Number of pilots for frame-specific updates while meta-training	N_T^{tr}	4	4
Number of pilots for meta-updates while meta-training	N_T^{te}	3000	4
Number of pilots for frame-specific updates while meta-testing	N_T^{tr}	8	4
Number of symbols with each channel states while meta-testing	N_*^{te}	4000	1000
Number of inner SGD updates while meta-training	I	2	2
Number of inner SGD updates while meta-testing	I_*	200	2
Number of meta-updates while meta-training	I_{meta}	200	no. of tasks t
Ensemble size (for Bayesian framework only)	R	100	100
Assumed precision of equalizer	β	-	150
Number of frames forming the initial experience	t_{init}	-	3
Number of meta-training iterations	-	200	100
Number of meta-testing frames averaged over	-	50	100

value at each data acquisition iteration benefits meta-training in practice. While using the previous iteration's optimized hyperparameter vector ξ as the starting point for the current iteration is useful in reducing the computational complexity [19], [96], we found it beneficial not to do so in our equalization problem to avoid meta-overfitting especially in the few-frames (e.g., 10 frames) regime of interest.

ACKNOWLEDGEMENT

The authors would like to thank the use of the Research Computing Facility at King's College London (2022), King's Computational Research, Engineering and Technology Environment (CREATE). Retrieved October 25, 2022, from <https://doi.org/10.18742/rnvf-m076>, and Rosalind (<https://rosalind.kcl.ac.uk>)

REFERENCES

- [1] K. M. Cohen, S. Park, O. Simeone, and S. Shamai, "Learning to learn to demodulate with uncertainty quantification via Bayesian meta-learning," in *Proc. WSA 25th Int. ITG Workshop Smart Antennas EURECOM*, 2021, pp. 202–207.
- [2] K. B. Letaief, W. Chen, Y. Shi, J. Zhang, and Y.-J. A. Zhang, "The roadmap to 6G: AI empowered wireless networks," *IEEE Commun. Mag.*, vol. 57, no. 8, pp. 84–90, Aug. 2019.
- [3] L. Bonati, S. D'Oro, M. Polese, S. Basagni, and T. Melodia, "Intelligence and learning in O-RAN for data-driven NextG cellular networks," *IEEE Commun. Mag.*, vol. 59, no. 10, pp. 21–27, Oct. 2021.
- [4] P. H. Masur, J. H. Reed, and N. Tripathi, "Artificial intelligence in open-radio access network," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 37, no. 9, pp. 6–15, Sep. 2022.
- [5] O-RAN Working Group 2, "O-RAN AI/ML workflow description and requirements-v1. 03," *O-RAN. WG2. AIML-v01. 03 Tech. Specification*, 2021.
- [6] O. Simeone, S. Park, and J. Kang, "From learning to meta-learning: Reduced training overhead and complexity for communication systems," in *Proc. 2nd 6G Wireless Summit*, 2020, pp. 1–5.
- [7] 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.
- [8] D. Barber, *Bayesian Reasoning and Machine Learning*. Cambridge, MA, USA: Cambridge Univ. Press, 2012.
- [9] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural network," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1613–1622.
- [10] H. Wang and D.-Y. Yeung, "A survey on Bayesian deep learning," *ACM Comput. Surv.*, vol. 53, no. 5, pp. 1–37, 2020.
- [11] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- [12] S. Thrun, "Lifelong Learning algorithms," in *Learning to Learn*. Boston, MA, USA: Springer, 1998, pp. 181–209.
- [13] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. 34th Int. Conf. Mach. Learn.*, vol. 70, 2017, pp. 1126–1135.
- [14] L. Zintgraf, K. Shiari, V. Kurin, K. Hofmann, and S. Whiteson, "Fast context adaptation via meta-learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 7693–7702.
- [15] D. Maclaurin, D. Duvenaud, and R. Adams, "Gradient-based hyperparameter optimization through reversible learning," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2113–2122.
- [16] Z. Li, F. Zhou, F. Chen, and H. Li, "Meta-SGD: Learning to learn quickly for few-shot learning," 2017, *arXiv:1707.09835*.
- [17] H. S. Behl, A. G. Baydin, and P. H. Torr, "Alpha MAML: Adaptive model-agnostic meta-learning," in *Proc. 6th ICML Workshop Automated Mach. Learn.*, Long Beach, California, USA, 2019, pp. 1–11.
- [18] J. Baxter, "Theoretical Models of Learning to Learn," in *Learning to Learn*. Boston, MA, USA: Springer, 1998, pp. 71–94.
- [19] 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.
- [20] M. Goutay, F. Ait Aoudia, and J. Hoyer, "Deep HyperNetwork-based MIMO detection," in *Proc. IEEE 21st Int. Workshop Signal Process. Adv. Wireless Commun.* Atlanta, USA, 2020, pp. 1–5.
- [21] Y. Yuan, G. Zheng, K.-K. Wong, B. Ottersten, and Z.-Q. Luo, "Transfer learning and meta learning-based fast downlink beamforming adaptation," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1742–1755, Mar. 2021.
- [22] Y. Hu, M. Chen, W. Saad, H. V. Poor, and S. Cui, "Meta-reinforcement learning for trajectory design in wireless UAV networks," in *Proc. IEEE Glob. Commun. Conf.*, Taipei, Taiwan, 2020, pp. 1–6.
- [23] T. Raviv, S. Park, N. Shlezinger, O. Simeone, Y. C. Eldar, and J. Kang, "Meta-ViterbiNet: Online meta-learned viterbi equalization for non-stationary channels," in *Proc. IEEE Int. Conf. Commun. Workshops*, Montreal, Canada, 2021, pp. 1–6.
- [24] A. E. Kalor, O. Simeone, and P. Popovski, "Prediction of mmWave/THz link blockages through meta-learning and recurrent neural networks," in *IEEE Wireless Commun. Lett.*, vol. 10, no. 12, 2021, pp. 2815–2819, doi: [10.1109/LWC.2021.3118269](https://doi.org/10.1109/LWC.2021.3118269).
- [25] I. Nikoloska and O. Simeone, "Fast power control adaptation via meta-learning for random edge graph neural networks," in *Proc. 2021 IEEE 22nd Int. Workshop Signal Process. Adv. Wireless Commun.*, Lucca, Italy, 2021, pp. 146–150.

- [26] J. Zhang, Y. Yuan, G. Zheng, I. Krikidis, and K.-K. Wong, "Embedding model-based fast meta learning for downlink beamforming adaptation," *IEEE Trans. Wireless Commun.*, vol. 21, no. 1, pp. 149–162, Jan. 2022.
- [27] A. E. Kalør, O. Simeone, and P. Popovski, "Prediction of mmWave/THz link blockages through meta-learning and recurrent neural networks," *IEEE Wireless Commun. Lett.*, vol. 10, no. 12, pp. 2815–2819, Dec. 2021.
- [28] Y. Jiang, H. Kim, H. Asnani, and S. Kannan, "Mind: Model independent neural decoder," in *Proc. IEEE 20th Int. Workshop Signal Process. Adv. Wireless Commun.*, Cannes, France, 2019, pp. 1–5.
- [29] J. Zhang, Y. He, Y.-W. Li, C.-K. Wen, and S. Jin, "Meta learning-based MIMO detectors: Design, simulation, and experimental test," *IEEE Trans. Wireless Commun.*, vol. 20, no. 2, pp. 1122–1137, Feb. 2021.
- [30] C. Nguyen, T.-T. Do, and G. Carneiro, "Uncertainty in model-agnostic meta-learning using variational inference," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2020, pp. 3090–3100.
- [31] K. Posch, J. Steinbrener, and J. Pilz, "Variational inference to measure model uncertainty in deep neural networks," 2019, *arXiv:1902.10189*.
- [32] 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.
- [33] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, "Bayesian active learning for classification and preference learning," 2011, *arXiv:1112.5745*.
- [34] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian active learning with image data," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1183–1192.
- [35] F. Sohrabi, T. Jiang, W. Cui, and W. Yu, "Active sensing for communications by learning," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 6, pp. 1780–1794, Jun. 2022.
- [36] H. Sahbi, S. Deschamps, and A. Stoian, "Active learning for interactive satellite image change detection," 2021, *arXiv:2110.04250*.
- [37] J. Kaddour et al., "Probabilistic active meta-learning," in *Proc. Adv. Neural Inf. Process. Syst. Virtual-Only Conf.*, vol. 33, 2020, pp. 20813–20822.
- [38] A. Kirsch, J. Van Amersfoort, and Y. Gal, "BatchBALD: Efficient and diverse batch acquisition for deep Bayesian active learning," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 7026–7037.
- [39] I. Nikoloska and O. Simeone, "Bayesian active meta-learning for black-box optimization," in *Proc. IEEE 23rd Int. Workshop Signal Process. Adv. Wireless Commun.*, Oulu, Finland, 2022, pp. 1–5.
- [40] 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.*, Paris, France, 2017, pp. 1–6.
- [41] 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.
- [42] 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.
- [43] 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.
- [44] A. Graves, "Practical variational inference for neural networks," in *Proc. 24th Int. Conf. Neural Inf. Process. Syst.*, 2011, pp. 2348–2356.
- [45] M. Dusenberry et al., "Efficient and scalable Bayesian neural nets with rank-1 factors," in *Proc. Int. Conf. Mach. Learn.*, Baltimore, USA, 2020, pp. 2782–2792.
- [46] 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. Virtual-Only Conf.*, vol. 34, 2021, pp. 20089–20103.
- [47] 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. Virtual-Only Conf.*, vol. 33, 2020, pp. 4346–4357.
- [48] R. M. Neal et al., "MCMC using Hamiltonian dynamics," in *Handbook of Markov Chain Monte Carlo*. Boca Raton, FL, USA: CRC Press, 2011, pp. 113–160.
- [49] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient Langevin dynamics," in *Proc. 28th Int. Conf. Mach. Learn.*, 2011, pp. 681–688.
- [50] 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. Representations*, 2020, pp. 1–27.
- [51] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose Bayesian inference algorithm," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 2378–2386.
- [52] 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*, 2017, pp. 73–78.
- [53] 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.
- [54] N. Zilberstein, C. Dick, R. Doost-Mohammady, A. Sabharwal, and S. Segarra, "Annealed Langevin dynamics for massive MIMO detection," 2022, *arXiv:2205.05776*.
- [55] 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.
- [56] N. K. Jha and V. K. 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.
- [57] N. K. Jha and V. K. 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.
- [58] 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.
- [59] M. Zecchin, S. Park, O. Simeone, M. Kountouris, and D. Gesbert, "Robust Bayesian learning for reliable wireless AI: Framework and applications," 2022, *arXiv:2207.00300*.
- [60] D. Liu, G. Zhu, Q. Zeng, J. Zhang, and K. Huang, "Wireless data acquisition for edge learning: Data-importance aware retransmission," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 406–420, Jan. 2021.
- [61] S.-E. Chiu, N. Ronquillo, and T. Javidi, "Active learning and CSI acquisition for mmWave initial alignment," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 11, pp. 2474–2489, Nov. 2019.
- [62] K. Yang, J. Ren, Y. Zhu, and W. Zhang, "Active learning for wireless IoT intrusion detection," *IEEE Wireless Commun.*, vol. 25, no. 6, pp. 19–25, Dec. 2018.
- [63] M. K. Abdel-Aziz, S. Samarakoon, M. Bennis, and W. Saad, "Ultra-reliable and low-latency vehicular communication: An active learning approach," *IEEE Commun. Lett.*, vol. 24, no. 2, pp. 367–370, Feb. 2020.
- [64] T. M. Hospedales, A. Antoniou, P. Micaelli, and A. J. Storkey, "Meta-learning in neural networks: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 9, pp. 5149–5169, Sep. 2021.
- [65] C. Finn, K. Xu, and S. Levine, "Probabilistic model-agnostic meta-learning," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 9537–9548.
- [66] S. Ravi and A. Beaton, "Amortized Bayesian meta-learning," in *Proc. Int. Conf. Learn. Representations*, Vancouver, Canada, 2018, pp. 1–14. [Online]. Available: <https://openreview.net/forum?id=rkgpy3C5tX>
- [67] J. Yoon, T. Kim, O. Dia, S. Kim, Y. Bengio, and S. Ahn, "Bayesian model-agnostic meta-learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 7343–7353.
- [68] M. Patacchiola, J. Turner, E. J. Crowley, M. O'Boyle, and A. J. Storkey, "Bayesian meta-learning for the few-shot setting via deep kernels," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 16108–16118.
- [69] Y. Zou and X. Lu, "Gradient-EM Bayesian meta-learning," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 20865–20875.
- [70] R. Amit and R. Meir, "Meta-learning by adjusting priors based on extended PAC-Bayes theory," in *Proc. Int. Conf. Mach. Learn.*, Stockholm, Sweden, PMLR, 2018, pp. 205–214.
- [71] J. Rothfuss, V. Fortuin, M. Josifoski, and A. Krause, "PACOH: Bayes-optimal meta-learning with PAC-Guarantees," in *Proc. Int. Conf. Mach. Learn. Virtual-Only Conf.*, 2021, pp. 9116–9126.
- [72] J. Rothfuss et al., "Meta-learning reliable priors in the function space," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021, vol. 34, pp. 280–293.
- [73] S. T. Jose, S. Park, and O. Simeone, "Information-theoretic analysis of epistemic uncertainty in Bayesian meta-learning," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2022, pp. 9758–9775.
- [74] I. Nikoloska and O. Simeone, "Quantum-aided meta-learning for Bayesian binary neural networks via born machines," 2022, *arXiv:2203.17089*.
- [75] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [76] O. Simeone, "A brief introduction to machine learning for engineers," *Found. Trends Signal Process.*, vol. 13, no. 12, pp. 200–431, Aug. 2018.

- [77] E. Angelino et al., "Patterns of scalable Bayesian inference," *Found. Trends Mach. Learn.*, vol. 9, no. 2-3, pp. 119–247, 2016.
- [78] T. M. Cover and J. A. Thomas, "Information theory and statistics," *Elements Inf. Theory*, vol. 1, no. 1, pp. 279–335, 1991.
- [79] S. T. Jose and O. Simeone, "Free energy minimization: A unified framework for modeling, inference, learning, and optimization [lecture notes]," *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 120–125, Mar. 2021.
- [80] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013, *arXiv:1312.6114*.
- [81] E. Grant, C. Finn, S. Levine, T. Darrell, and T. Griffiths, "Recasting gradient-based meta-learning as hierarchical bayes," in *Proc. 6th Int. Conf. Learn. Representations*, Vancouver, Canada, 2018, pp. 1–13.
- [82] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA, USA: MIT Press, 2009.
- [83] O. Simeone, *Machine Learning for Engineers*. Cambridge, U.K.: Cambridge University Press, 2022.
- [84] A. Griewank, "Some bounds on the complexity of gradients, Jacobians, and Hessians," in *Complexity in Numerical Optimization*. Singapore: World Scientific, 1993, pp. 128–162.
- [85] A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine, "Meta-learning with implicit gradients," in *Proc. Neural Inf. Process. Syst.*, Vancouver, Canada, 2019, pp. 113–124.
- [86] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih, "Monte Carlo gradient estimation in machine learning," *J. Mach. Learn. Res.*, vol. 21, no. 132, pp. 1–62, 2020.
- [87] T. Melliush, C. Saunders, I. Nourtdinov, and V. Vovk, "Comparing the Bayes and typicalness frameworks," in *Proc. Eur. Conf. Mach. Learn.*, Freiburg, Germany, Springer, 2001, pp. 360–371.
- [88] M. H. DeGroot and S. E. Fienberg, "The comparison and evaluation of forecasters," *J. Roy. Stat. Soc.: Ser. D*, vol. 32, no. 1-2, pp. 12–22, 1983.
- [89] Y. Zhang et al., "DeepWiPHY: Deep learning-based receiver design and dataset for IEEE 802.11 ax systems," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1596–1611, Mar. 2021.
- [90] A. G. Helmy, M. Di Renzo, and N. Al-Dhahir, "On the robustness of spatial modulation to I/Q imbalance," *IEEE Commun. Lett.*, vol. 21, no. 7, pp. 1485–1488, Jul. 2017.
- [91] 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.
- [92] M. Yin, G. Tucker, M. Zhou, S. Levine, and C. Finn, "Meta-learning without memorization," in *Proc. Int. Conf. Learn. Representations*, 2019, pp. 1–21.
- [93] R. Li et al., "A channel coding benchmark for meta-learning," in *Proc. Adv. Neural Inf. Process. Syst. Virtual-Only Conf., Track Datasets Benchmarks*, 2021, pp. 1–12. [Online]. Available: <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/e96ed478dab8595a7dbda4c3bbee168f-Abstract-round2.html>
- [94] A. Paszke et al., "Automatic differentiation in pytorch," in *Proc. Conf. Neural Inf. Process. Syst. Workshop Autodiff*, 2017.
- [95] J. Knoblauch, J. Jewson, and T. Damoulas, "Generalized variational inference: Three arguments for deriving new posteriors," 2019, *arXiv:1904.02063*.
- [96] C. Finn, A. Rajeswaran, S. Kakade, and S. Levine, "Online meta-learning," in *Proc. Int. Conf. Mach. Learn.*, Long Beach, CA, USA, 2019, pp. 1920–1930.



interests include Bayesian and reliable machine learning, signal processing, and their applications to communications systems.

Kfir M. Cohen (Student Member, IEEE) received the B.Sc. (*summa cum laude*) and M.Sc. degrees from the Electrical Engineering Faculty, Technion–Israel Institute for Technology, Haifa, Israel, in 2006 and 2013, respectively. He is currently working toward the Ph.D. degree with King’s Communications, Learning and Information Processing Lab, Department of Engineering of King’s College London (KCL), London, U.K. He worked for 15 years in different R&D roles. Before joining KCL, his last position was as a Communication Signal Processing Engineer. His research



Sangwoo Park (Member, IEEE) received the B.Sc. degree in physics, and the M.S.E and Ph.D. degrees in electrical engineering from Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea, in 2014, 2016, and 2020, respectively. He is currently a Research Associate with the Department of Engineering, King’s Communications, Learning and Information Processing Lab, King’s College London, London, U.K. His research interests include practical, reliable AI and its application for wireless communication systems and quantum information processing.



Osvaldo Simeone (Fellow, IEEE) received the M.Sc. degree (with Hons.) and the Ph.D. degree in information engineering from Politecnico di Milano, Milan, Italy, in 2001 and 2005, respectively. He is currently a Professor of information engineering with the Centre for Telecommunications Research, Department of Engineering, King’s College London, London, U.K., where he directs the King’s Communications, Learning and Information Processing Lab. From 2006 to 2017, he was a Faculty Member with the Electrical and Computer Engineering Department, New Jersey

Institute of Technology, Newark, NJ, USA, where he was affiliated with the Center for Wireless Information Processing. His research interests include information theory, machine learning, wireless communications, neuromorphic computing, and quantum machine learning. Dr Simeone was a co-recipient of the 2022 IEEE Communications Society Outstanding Paper Award, 2021 IEEE Vehicular Technology Society Jack Neubauer Memorial Award, 2019 IEEE Communication Society Best Tutorial Paper Award, 2018 IEEE Signal Processing Best Paper Award, 2017 JCN Best Paper Award, 2015 IEEE Communication Society Best Tutorial Paper Award, Best Paper awards of IEEE SPAWC 2007 and IEEE WRECOM 2007, and was the recipient of an Open Fellowship by the EPSRC in 2022 and a Consolidator Grant by the European Research Council (ERC) in 2016. His research has been also supported by the U.S. National Science Foundation, European Commission, European Research Council, Vienna Science and Technology Fund, European Space Agency, and by the number of industrial collaborations including with Intel Labs and InterDigital. He is the Chair of the Signal Processing for Communications and Networking Technical Committee of the IEEE Signal Processing Society and of the U.K. & Ireland Chapter of the IEEE Information Theory Society. He is currently a Distinguished Lecturer of the IEEE Communications Society, and he was a Distinguished Lecturer of the IEEE Information Theory Society in 2017 and 2018. He is also the author of the textbook *Machine Learning for Engineers* published by Cambridge University Press, four monographs, two edited books, and more than 180 research journal and magazine papers. He is a Fellow of the IET and EPSRC.



Shlomo Shamai (Shitz) (Life Fellow, IEEE) is currently with the Department of Electrical Engineering, Technion–Israel Institute of Technology, where he is also a Technion Distinguished Professor, and holds the William Fondiller Chair of telecommunications. He was the recipient of the 2011 Claude E. Shannon Award, 2014 Rothschild Prize in Mathematics/Computer Sciences and Engineering, and the 2017 IEEE Richard W. Hamming Medal and numerous Technical and Paper awards and recognitions of the IEEE (Donald G. Fink Prize Paper Award), Informa-

tion Theory, Communications and Signal Processing Societies and EURASIP. He was also the co-recipient of the 2018 Third Bell Labs Prize for Shaping the Future of Information and Communications Technology. He is listed as a Highly Cited Researcher (computer science) for the years 2013/4/5/6/7/8. He was an Associate Editor for the Shannon Theory of the IEEE TRANSACTIONS ON INFORMATION THEORY and has also been twice on the Board of Governors of the Information Theory Society. He was on the Executive Editorial Board of the IEEE TRANSACTIONS ON INFORMATION THEORY, IEEE Information Theory Society Nominations and Appointments Committee, and IEEE Information Theory Society (Shannon Award Committee). Dr. Shamai is an URSI Fellow, the Member of the Israeli Academy of Sciences and Humanities, and Foreign Member of the U.S. National Academy of Engineering.