

# ViterbiNet: A Deep Learning Based Viterbi Algorithm for Symbol Detection

Nir Shlezinger<sup>1</sup>, Member, IEEE, Nariman Farsad<sup>2</sup>, Member, IEEE, Yonina C. Eldar, Fellow, IEEE, and Andrea J. Goldsmith, Fellow, IEEE

**Abstract**—Symbol detection plays an important role in the implementation of digital receivers. In this work, we propose ViterbiNet, which is a data-driven symbol detector that does not require channel state information (CSI). ViterbiNet is obtained by integrating deep neural networks (DNNs) into the Viterbi algorithm. We identify the specific parts of the Viterbi algorithm that depend on the channel model, and design a DNN to implement only those computations, leaving the rest of the algorithm structure intact. We then propose a meta-learning based approach to train ViterbiNet online based on recent decisions, allowing the receiver to track dynamic channel conditions without requiring new training samples for every coherence block. Our numerical evaluations demonstrate that the performance of ViterbiNet, which is ignorant of the CSI, approaches that of the CSI-based Viterbi algorithm, and is capable of tracking time-varying channels without needing instantaneous CSI or additional training data. Moreover, unlike conventional Viterbi detection, ViterbiNet is robust to CSI uncertainty, and it can be reliably implemented in complex channel models with constrained computational burden. More broadly, our results demonstrate the conceptual benefit of designing communication systems that integrate DNNs into established algorithms.

**Index Terms**—Symbol detection, machine learning (ML).

## I. INTRODUCTION

A FUNDAMENTAL task of a digital receiver is to reliably recover the transmitted symbols from the observed channel output. This task is commonly referred to as *symbol detection*. Conventional symbol detection algorithms, such as those based on the maximum a-posteriori probability (MAP) rule, require complete knowledge of the underlying channel model and its parameters. Consequently, in some cases, conventional methods cannot be applied, particularly when the

channel model is highly complex, poorly understood, or does not capture well the underlying physics of the system. Furthermore, when the channel models are known, many detection algorithms rely on the instantaneous channel state information (CSI), i.e., the instantaneous parameters of the channel model, for detection. Therefore, conventional techniques require the instantaneous CSI to be estimated. However, this process entails overhead, which decreases the data transmission rate. Moreover, inaccurate CSI estimation typically degrades detection performance.

One of the most common CSI-based symbol detection methods is the iterative scheme proposed by Viterbi in [1], known as the Viterbi algorithm. The Viterbi algorithm is an efficient symbol detector that is capable of achieving the minimal probability of error in recovering the transmitted symbols for channels obeying a Markovian input-output stochastic relationship, which is encountered in many practical channels [2]. Since the Viterbi algorithm requires the receiver to know the exact statistical relationship relating the channel input and output, the receiver must have full instantaneous CSI.

An alternative data-driven approach is based on machine learning (ML). ML methods, and in particular, deep neural networks (DNNs), have been the focus of extensive research in recent years due to their empirical success in various applications, including image processing and speech processing [3]. There are several benefits in using ML schemes over traditional model-based approaches: First, ML methods are independent of the underlying stochastic model, and thus can operate efficiently in scenarios where this model is unknown or its parameters cannot be accurately estimated. Second, when the underlying model is extremely complex, ML algorithms have demonstrated the ability to extract and disentangle meaningful semantic information from the observed data [4], a task which is very difficult to carry out using traditional model-based approaches. Finally, ML techniques often lead to faster convergence compared to iterative model-based approaches, even when the model is known [5], [6].

Recent years have witnessed growing interest in the application of DNNs for receiver design. Due to the large amount of recent works on this topic, we only discuss a few representative papers; detailed surveys can be found in [7]–[10]: The work [11] used deep learning for decoding linear codes. Decoding of structured codes using DNNs was considered in [12], while symbol recovery in multiple-input multiple-output (MIMO) systems was treated in [13]. The work [14] used variational autoencoders for equalizing linear

Manuscript received May 26, 2019; revised November 6, 2019; accepted January 30, 2020. Date of publication February 14, 2020; date of current version May 8, 2020. This work was supported in part by the U.S.—Israel Binational Science Foundation under Grant 2026094, in part by the Israel Science Foundation under Grant 0100101, and in part by the Office of the Naval Research under Grant 18-1-2191. This article was presented at the 2019 IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), Cannes, France. The associate editor coordinating the review of this article and approving it for publication was X. Chen. (Corresponding author: Nir Shlezinger.)

Nir Shlezinger and Yonina C. Eldar are with the Faculty of Math and CS, Weizmann Institute of Science, Rehovot 7610001, Israel (e-mail: nirshlezinger1@gmail.com; yonina@weizmann.ac.il).

Nariman Farsad and Andrea J. Goldsmith are with the Department of EE, Stanford University, Stanford, CA 94305 USA (e-mail: nfarsad@stanford.edu; andrea@wsl.stanford.edu).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TWC.2020.2972352

multipath channels. Sequence detection using bi-directional recurrent neural networks (RNNs) was proposed in [15], while [16] considered a new ML-based channel decoder by combining convolutional neural networks with belief propagation. RNN structures for decoding sequential codes were studied in [17]. The work [18] proposed a neural network architecture that enables parallel implementation of the Viterbi algorithm with binary signals in hardware. DNN-based MIMO receivers for mitigating the effect of low-resolution quantization were studied in [19]. These approaches were shown to yield good performance when sufficient training is available. However, previous applications of DNNs typically treat the network as a black box, hoping to achieve the desired performance by using sufficient training and relying on methods that were developed to treat other tasks such as image processing. This gives rise to the question of whether additional gains, either in performance, complexity, or training size, can be achieved by combining model-based methods, such as the Viterbi algorithm, with ML-based techniques.

In this work, we design and study ML-based symbol detection for finite-memory causal channels based on the Viterbi algorithm. Our design is inspired by deep unfolding, which is a common method for obtaining ML architectures from model-based iterative algorithms [5], [20], [21]. Unfolding was shown to yield efficient and reliable DNNs for applications such as sparse recovery [5], recovery from one-bit measurements [6], matrix factorization [22], image deblurring [23], and robust principal component analysis [24]. However, there is a fundamental difference between our approach and conventional unfolding: The main rationale of unfolding is to convert each iteration of the algorithm into a layer, namely, to design a DNN in light of a model-based algorithm, or alternatively, to *integrate the algorithm into the DNN*. Our approach to symbol detection implements the Viterbi channel-model-based algorithm, while only removing its channel dependence by replacing the CSI-based computations with dedicated DNNs, i.e., we *integrate ML into the Viterbi algorithm*.

In particular, we propose ViterbiNet, which is an ML-based symbol detector integrating deep learning into the Viterbi algorithm. The resulting system approximates the mapping carried out by the channel-model-based Viterbi method in a data-driven fashion without requiring CSI, namely, without knowing the exact channel input-output statistical relationship. ViterbiNet combines established properties of DNN-based sequence classifiers to implement the specific computations in the Viterbi algorithm that are CSI-dependent. Unlike direct application of ML algorithms, the resulting detector is capable of exploiting the underlying Markovian structure of finite-memory causal channels in the same manner as the Viterbi algorithm. ViterbiNet consists of a simple network architecture which can be trained with a relatively small number of training samples.

Next, we propose a method for adapting ViterbiNet to dynamic channel conditions online without requiring new training data every time the statistical model of the channel changes. Our approach is based on meta-learning [25], namely, automatic learning and adaptation of DNNs, typically by acquiring training from related tasks. The concept of

meta-learning was recently used in the context of ML for communications in the work [26], which used pilots from neighboring devices as meta-training in an Internet of Things (IOT) network. Our proposed approach utilizes channel coding to generate meta-training from each decoded block. In particular, we use the fact that forward error correction (FEC) codes can compensate for decision errors as a method for generating accurate training from the decoded block. This approach exploits both the inherent structure of coded communication signals, as well as the ability of ViterbiNet to train with the relatively small number of samples obtained from decoding a single block.

Our numerical evaluations demonstrate that, when the training data obeys the same statistical model as the test data, ViterbiNet achieves roughly the same performance as the CSI-based Viterbi algorithm. Furthermore, when ViterbiNet is trained for a variety of different channels, it notably outperforms the Viterbi algorithm operating with the same level of CSI uncertainty. It is also illustrated that ViterbiNet performs well in complex channel models, where the Viterbi algorithm is extremely difficult to implement even when CSI is available. In the presence of block-fading channels, we show that by using our proposed online training scheme, ViterbiNet is capable of tracking the varying channel conditions and approaches the performance achievable with the Viterbi detector, which requires accurate instantaneous CSI for each block. Our results demonstrate that reliable, efficient, and robust ML-based communication systems can be realized by integrating ML into existing CSI-based techniques and accounting for the inherent structure of digital communication signals.

The rest of this paper is organized as follows: In Section II we present the system model and review the Viterbi algorithm. Section III proposes ViterbiNet: A receiver architecture that implements Viterbi detection using DNNs. Section IV presents a method for training ViterbiNet online, allowing it to track block-fading channels. Section V details numerical training and performance results of ViterbiNet, and Section VI provides concluding remarks.

Throughout the paper, we use upper case letters for random variables (RVs), e.g.  $X$ . Boldface lower-case letters denote vectors, e.g.,  $\mathbf{x}$  is a deterministic vector, and  $\mathbf{X}$  is a random vector, and the  $i$ th element of  $\mathbf{x}$  is written as  $(\mathbf{x})_i$ . The probability density function (PDF) of an RV  $X$  evaluated at  $x$  is denoted  $p_X(x)$ ,  $\mathcal{Z}$  is the set of integers,  $\mathcal{N}$  is the set of positive integers,  $\mathcal{R}$  is the set of real numbers, and  $(\cdot)^T$  is the transpose operator. All logarithms are taken to base 2. Finally, for any sequence, possibly multivariate,  $\mathbf{y}[i]$ ,  $i \in \mathcal{Z}$ , and integers  $b_1 < b_2$ ,  $\mathbf{y}_{b_1}^{b_2}$  is the column vector  $[\mathbf{y}^T[b_1], \dots, \mathbf{y}^T[b_2]]^T$  and  $\mathbf{y}^{b_2} \equiv \mathbf{y}_1^{b_2}$ .

## II. SYSTEM MODEL AND PRELIMINARIES

### A. System Model

We consider the problem of recovering a block of  $t$  symbols transmitted over a finite-memory stationary causal channel. Let  $S[i] \in \mathcal{S}$  be the symbol transmitted at time index  $i \in \{1, 2, \dots, t\} \triangleq \mathcal{T}$ , where each symbol is uniformly distributed over a set of  $m$  constellation points, thus  $|\mathcal{S}| = m$ . We use  $Y[i] \in \mathcal{Y}$  to denote the channel output at time index  $i$ .

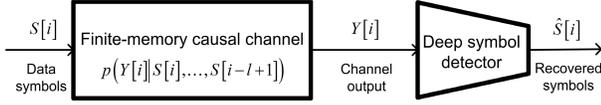


Fig. 1. System model.

Since the channel is causal and has finite memory,  $Y[i]$  is given by some stochastic mapping of  $\mathcal{S}_{i-l+1}^i$ , where  $l$  denotes the memory of the channel, assumed to be smaller than the blocklength, i.e.,  $l < t$ . Consequently, the conditional PDF of the channel output given the channel input satisfies

$$p_{\mathbf{Y}_{k_2}^{k_2} | \mathcal{S}^t}(\mathbf{y}_{k_1}^{k_2} | \mathbf{s}^t) = \prod_{i=k_1}^{k_2} p_{Y[i] | \mathcal{S}_{i-l+1}^i}(y[i] | \mathbf{s}_{i-l+1}^i), \quad (1)$$

for all  $k_1, k_2 \in \mathcal{T}$  such that  $k_1 \leq k_2$ . The fact that the channel is stationary implies that for each  $y \in \mathcal{Y}$ ,  $\mathbf{s} \in \mathcal{S}^l$ , the conditional PDF  $p_{Y[i] | \mathcal{S}_{i-l+1}^i}(y | \mathbf{s})$  does not depend on the index  $i$ . An illustration of the system is depicted in Fig. 1.

Our goal is to design a DNN architecture for recovering  $\mathcal{S}^t$  from the channel output  $\mathbf{Y}^t$ . In particular, in our model the receiver assumes that the channel is stationary, causal, and has finite memory  $l$ , namely, that the input-output statistical relationship is of the form (1). The receiver also knows the constellation  $\mathcal{S}$ . We do not assume that the receiver knows the conditional PDF  $p_{Y[i] | \mathcal{S}_{i-l+1}^i}(y | \mathbf{s})$ , i.e., the receiver does not have this CSI. The optimal detector for finite-memory channels, based on which we design our network, is the Viterbi algorithm [1]. Therefore, as a preliminary step to designing the DNN, we review Viterbi detection in the following subsection.

### B. The Viterbi Detection Algorithm

The following description of the Viterbi algorithm is based on [27, Ch. 3.4]. Since the constellation points are equiprobable, the optimal decision rule in the sense of minimal probability of error is the maximum likelihood decision rule, namely, for a given channel output  $\mathbf{y}^t$ , the estimated output is given by

$$\begin{aligned} \hat{\mathbf{s}}^t(\mathbf{y}^t) &\triangleq \arg \max_{\mathbf{s}^t \in \mathcal{S}^t} p_{\mathbf{Y}^t | \mathcal{S}^t}(\mathbf{y}^t | \mathbf{s}^t) \\ &= \arg \min_{\mathbf{s}^t \in \mathcal{S}^t} -\log p_{\mathbf{Y}^t | \mathcal{S}^t}(\mathbf{y}^t | \mathbf{s}^t). \end{aligned} \quad (2)$$

By defining

$$c_i(\mathbf{s}) \triangleq -\log p_{Y[i] | \mathcal{S}_{i-l+1}^i}(y[i] | \mathbf{s}), \quad \mathbf{s} \in \mathcal{S}^l, \quad (3)$$

it follows from (1) that the log-likelihood function in (2) can be written as

$$\log p_{\mathbf{Y}^t | \mathcal{S}^t}(\mathbf{y}^t | \mathbf{s}^t) = \sum_{i=1}^t c_i(\mathbf{s}_{i-l+1}^i), \quad (4)$$

and the optimization problem (2) becomes

$$\hat{\mathbf{s}}^t(\mathbf{y}^t) = \arg \min_{\mathbf{s}^t \in \mathcal{S}^t} \sum_{i=1}^t c_i(\mathbf{s}_{i-l+1}^i). \quad (5)$$

### Algorithm 1 The Viterbi Algorithm

- 1: **Input:** Block of channel outputs  $\mathbf{y}^t$ , where  $t > l$ .
- 2: **Initialization:** Set  $k = 1$ , and fix the initial path cost  $\tilde{c}_0(\tilde{\mathbf{s}}) = 0$ , for each state  $\tilde{\mathbf{s}} \in \mathcal{S}^l$ .
- 3: For each state  $\tilde{\mathbf{s}} \in \mathcal{S}^l$ , compute the path cost via  $\tilde{c}_k(\tilde{\mathbf{s}}) = \min_{\mathbf{u} \in \mathcal{S}^l: \mathbf{u}_2^l = \tilde{\mathbf{s}}^{l-1}} (\tilde{c}_{k-1}(\mathbf{u}) + c_k(\tilde{\mathbf{s}}))$ .
- 4: If  $k \geq l$ , set  $(\hat{\mathbf{s}})_{k-l+1} := (\tilde{\mathbf{s}}_k^o)_1$ , where  $\tilde{\mathbf{s}}_k^o = \arg \min_{\tilde{\mathbf{s}} \in \mathcal{S}^l} \tilde{c}_k(\tilde{\mathbf{s}})$ .
- 5: Set  $k := k + 1$ . If  $k \leq t$  go to Step 3.
- 6: **Output:** decoded output  $\hat{\mathbf{s}}^t$ , where  $\hat{\mathbf{s}}_{t-l+1}^t := \tilde{\mathbf{s}}_t^o$ .

The optimization problem (5) can be solved recursively using dynamic programming, by treating the possible combinations of transmitted symbols at each time instance as *states* and iteratively updating a *path cost* value for each state. The resulting recursive solution, which is known as the Viterbi algorithm, is given below as Algorithm 1.

In addition to its ability to achieve the minimal error probability [2], the Viterbi algorithm has several major advantages:

- A1 The algorithm solves (2) at a computational complexity that is linear in the blocklength  $t$ . For comparison, the computational complexity of solving (2) directly grows exponentially with  $t$ .
- A2 The algorithm produces estimates sequentially on runtime. In particular, while in (2) the estimated output  $\hat{\mathbf{s}}^t$  is computed using the entire received block  $\mathbf{y}^t$ , Algorithm 1 computes  $\hat{s}[i]$  once  $y[i+l-1]$  is received.

In order to implement Algorithm 1, one must be able to compute  $c_i(\mathbf{s})$  of (3) for all  $i \in \mathcal{T}$  and for each  $\mathbf{s} \in \mathcal{S}^l$ . Consequently, the conditional PDF of the channel, which we refer to as full CSI, must be explicitly known. As discussed in the introduction, obtaining full CSI may be extremely difficult in rapidly changing channels and may also require a large training overhead. In the following section we propose ViterbiNet, an ML-based symbol decoder based on the Viterbi algorithm, that does not require CSI.

## III. ViterbiNet

### A. Integrating ML into the Viterbi Algorithm

In order to integrate ML into the Viterbi Algorithm, we note that CSI is required in Algorithm 1 only in Step 3 to compute the log-likelihood function  $c_i(\mathbf{s})$ . Once  $c_i(\mathbf{s})$  is computed for each  $\mathbf{s} \in \mathcal{S}^l$ , the Viterbi algorithm only requires knowledge of the memory length  $l$  of the channel. This requirement is much easier to satisfy compared to knowledge of the exact channel input-output statistical relationship, i.e., full CSI. In fact, a common practice when using the model-based Viterbi algorithm is to limit the channel memory by filtering the channel output prior to symbol detection [28] in order to reduce computational complexity. The approach in [28], which assumes that the channel is modeled as a linear time-invariant (LTI) filter, is an example of how the channel memory used by the algorithm can be fixed in advance, regardless of the true memory length of the underlying channels. Furthermore, estimating the channel memory in a data-driven fashion without a-priori knowledge of the underlying model is a much

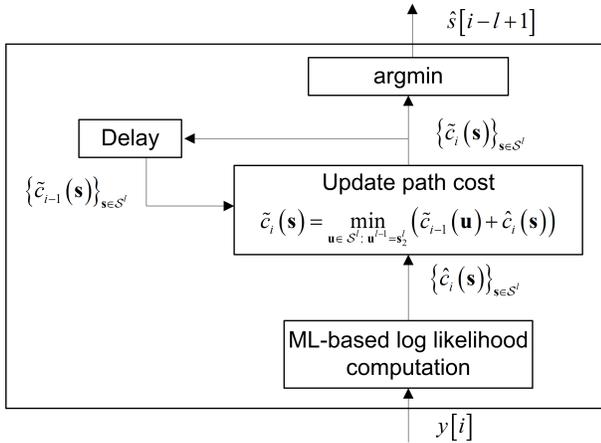


Fig. 2. Proposed DNN-based Viterbi decoder.

simpler task compared to symbol detection, and can be reliably implemented using standard correlation peak based estimators, see, e.g., [29]. We henceforth assume that the receiver has either accurate knowledge or a reliable upper bound on  $l$ , and leave the extension of the decoder to account for unknown memory length to future exploration.

Since the channel is stationary, it holds by (3) that if  $y[i] = y[k]$  then  $c_i(\mathbf{s}) = c_k(\mathbf{s})$ , for each  $\mathbf{s} \in \mathcal{S}^l$ , and  $i, k \in \mathcal{T}$ . Consequently, the log-likelihood function  $c_i(\mathbf{s})$  depends only on the values of  $y[i]$  and of  $\mathbf{s}$ , and not on the time index  $i$ . Therefore, in our proposed structure we replace the explicit computation of the log-likelihood (3) with an ML-based system that learns to evaluate the cost function from the training data. In this case, the input of the system is  $y[i]$  and the output is an estimate of  $c_i(\mathbf{s})$ , denoted  $\hat{c}_i(\mathbf{s})$ , for each  $\mathbf{s} \in \mathcal{S}^l$ . The rest of the Viterbi algorithm remains intact, namely, the detector implements Algorithm 1 while using ML techniques to compute the log-likelihood function  $c_i(\mathbf{s})$ . The proposed architecture is illustrated in Fig. 2.

A major challenge in implementing a network capable of computing  $c_i(\mathbf{s})$  from  $y[i]$  stems from the fact that, by (3),  $c_i(\mathbf{s})$  represents the log-likelihood of  $Y[i] = y[i]$  given  $\mathcal{S}_{i-l+1}^i = \mathbf{s}$ . However, DNNs trained to minimize the cross entropy loss typically output the conditional distribution of  $\mathcal{S}_{i-l+1}^i = \mathbf{s}$  given  $Y[i] = y[i]$ , i.e.,  $p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$ . Specifically, classification DNNs with input  $Y[i] = y[i]$ , which typically takes values in a discrete set, or alternatively, is discretized by binning a continuous quantity as in [30], [31], output the distribution of the label  $\mathcal{S}_{i-l+1}^i$  conditioned on that input  $Y[i] = y[i]$ , and not the distribution of the input conditioned on all possible values of the label  $\mathcal{S}_{i-l+1}^i$ . For this reason, the previous work [15] used a DNN to approximate the MAP detector by considering  $p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$ . It is emphasized that the quantity which is needed for the Viterbi algorithm is the conditional PDF  $p_{Y[i]|\mathcal{S}_{i-l+1}^i}(\mathbf{s}|y[i])$ , and not the conditional distribution  $p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$ . Note, for example, that while  $\sum_{\mathbf{s} \in \mathcal{S}^l} p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i]) = 1$ , for the desired conditional PDF  $\sum_{\mathbf{s} \in \mathcal{S}^l} p_{Y[i]|\mathcal{S}_{i-l+1}^i}(\mathbf{s}|y[i]) \neq 1$  in general. Therefore, outputs generated using conventional DNNs with a softmax output layer are not applicable. The

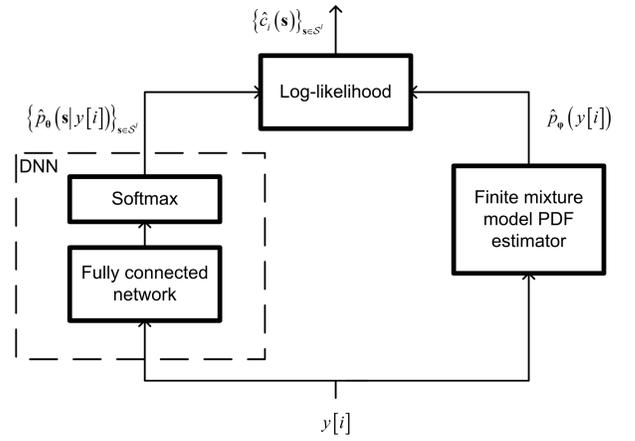


Fig. 3. ML-based log-likelihood computation.

fact that Algorithm 1 specifically uses the conditional PDF  $p_{Y[i]|\mathcal{S}_{i-l+1}^i}(y[i]|\mathbf{s})$  instead of  $p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$  allows it to exploit the Markovian nature of the channel, induced by the finite memory in (1), resulting in the advantages A1-A2 discussed in Subsection II-B.

In order to tackle this difficulty, we recall that by Bayes' theorem, as the channel inputs are equiprobable, the desired conditional PDF  $p_{Y[i]|\mathcal{S}_{i-l+1}^i}(y|\mathbf{s})$  can be written as

$$p_{Y[i]|\mathcal{S}_{i-l+1}^i}(y|\mathbf{s}) = \frac{p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y) \cdot p_{Y[i]}(y)}{m^{-l}}. \quad (6)$$

Therefore, given estimates of  $p_{Y[i]}(y[i])$  and of  $p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$  for each  $\mathbf{s} \in \mathcal{S}^l$ , the log-likelihood function  $c_i(\mathbf{s})$  can be recovered using (6) and (3).

A parametric estimate of  $p_{\mathcal{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$ , denoted  $\hat{p}_\theta(\mathbf{s}|y[i])$ , can be reliably obtained from training data using standard classification DNNs with a softmax output layer. The marginal PDF of  $Y[i]$  may be estimated from the training data using conventional kernel density estimation methods [32]. Furthermore, the fact that  $Y[i]$  is a stochastic mapping of  $\mathcal{S}_{i-l+1}^i$  implies that its distribution can be approximated as a mixture model of  $m^l$  kernel functions [33]. Consequently, a parametric estimate of  $p_{Y[i]}(y[i])$ , denoted  $\hat{p}_\phi(y[i])$ , can be obtained from the training data using mixture density networks [35], expectation minimization (EM)-based algorithms [33, Ch. 2], or any other finite mixture model fitting method. The resulting ML-based log-likelihood computation is illustrated in Fig. 3.

The ML-based log-likelihood computation depicted in Fig. 3 is applicable for both real-valued channels, where  $\mathcal{Y}$  and  $\mathcal{S}$  are subsets of  $\mathcal{R}$ , as well as complex channels, in which the elements of  $\mathcal{Y}$  and  $\mathcal{S}$  take complex values. In particular, for complex-valued channels, the input to the classification DNN should be a real vector consisting of the real and imaginary parts of  $y[i]$ , as DNNs typically operate on real-valued quantities. Furthermore, the finite mixture model PDF estimator should use kernel functions for complex distributions, e.g., complex Gaussian mixture. Except for these modifications, the architecture of ViterbiNet is invariant as to whether the channel is real or complex-valued.

## B. Discussion

The architecture of ViterbiNet is based on the Viterbi algorithm, which achieves the minimal probability of error for communications over finite-memory causal channels. Since ViterbiNet is data-driven, it learns the log-likelihood function from training data without requiring full CSI. When properly trained, the proposed DNN is therefore expected to approach the performance achievable with the conventional CSI-based Viterbi algorithm, as numerically demonstrated in Section V. Unlike direct applications of classification DNNs, which produce a parametric estimate of the conditional distribution of the transmitted symbols given the channel output, ViterbiNet is designed to estimate the conditional distribution of the channel output given the transmitted symbols. This conditional distribution, which encapsulates the Markovian structure of finite-memory channels, is then used to obtain the desired log-likelihoods, as discussed in the previous subsection. The estimation of these log-likelihoods, which are not obtained by simply applying a classification DNN to the channel output, is then exploited to recover the transmitted symbols using the same operations carried out by the conventional Viterbi algorithm.

ViterbiNet replaces the CSI-based parts of the Viterbi algorithm with an ML-based scheme. Since this aspect of the channel model is relatively straightforward to learn, ViterbiNet can utilize a simple and standard DNN architecture. The simple DNN structure implies that the network can be trained quickly with a relatively small number of training samples, as also observed in the numerical study in Section V. This indicates the potential of the proposed architecture to adapt on runtime to channel variations with minimal overhead, possibly using pilot sequences periodically embedded in the transmitted frame for online training.

The approach used in designing ViterbiNet, namely, maintaining the detection scheme while replacing its channel-model-based computations with dedicated ML methods, can be utilized for implementing additional algorithms in communications and signal processing in a data-driven manner. For instance, in [36], we showed how this approach can allow a MIMO receiver to learn to implement the soft iterative interference cancellation scheme of [37] from a small training set. In particular, the ML-based computation of the conditional probability used in ViterbiNet to carry out Viterbi detection can be utilized to realize other trellis or factor graph based detection schemes, such as the BCJR algorithm [38], by properly modifying the processing of these estimates. We leave the study of these alternative data-driven schemes for future investigation, and initial results can be found in [39].

The fact that ViterbiNet is based on the Viterbi algorithm implies that it also suffers from some of its drawbacks. For example, when the constellation size  $m$  and the channel memory  $l$  grow large, the Viterbi algorithm becomes computationally complex due to the need to compute the log-likelihood for each of the possible  $m^l$  different values of  $s$ . Consequently, the complexity of ViterbiNet is expected to grow exponentially as  $m$  and  $l$  grow, since the label space of the DNN grows exponentially. It is noted that greedy schemes for reducing the

complexity of the Viterbi algorithm, such as beam search [40] and reduced-state equalization [41], were shown to result in minimal performance degradation, and we thus expect that these methods can inspire similar modifications to ViterbiNet, facilitating its application with large  $m$  and  $l$ . We leave the research into reducing the complexity of ViterbiNet through such methods to future investigation.

## IV. EXTENSION TO BLOCK-FADING CHANNELS

A major challenge associated with using ML algorithms in communications stems from the fact that, in order to exploit the ability of DNNs to learn an underlying model, the network should operate under the same (or a closely resembling) statistical relationship for which it was trained. Many communication channels, and particularly wireless channels, are dynamic by nature, and are commonly modeled as *block-fading* [42]. In such channels, each transmitted block may undergo a different statistical transformation. In Section V, we show that ViterbiNet is capable of achieving relatively good performance when trained using samples taken from a variety of channels rather than from the specific channel for which it is tested. Consequently, a possible approach to use ViterbiNet in block-fading channels is to train the system using samples acquired from a large variety of channels. However, such an approach requires a large amount of training data in order to cover a multitude of different channel conditions, and is also expected to result in degraded performance, especially when tested under a channel model that deviates significantly from the channels used during training. This motivates us to extend ViterbiNet by allowing the receiver to track and adjust to time-varying channel conditions in real-time, by exploiting the inherent structure of digital communication signals and, particularly, of coded communications. Broadly speaking, we utilize the fact that in coded systems, the receiver is capable of identifying and correcting symbol detection errors. The identification and correction of these errors can be used to retrain ViterbiNet online using meta-learning principles [25], allowing it to track channel variations. To present this approach, we first extend the channel model of Subsection IV-A to account for coded communications and block-fading channels. Then, we detail in Subsection IV-B how ViterbiNet can exploit the use of channel coding for online training.

### A. Coded Communications over Block-Fading Channels

In coded communications, each block of  $t$  transmitted channel symbols represents a channel codeword conveying  $b$  bits. During the  $j$ th block,  $j \in \mathcal{N}$ , a vector of bits, denoted  $\mathbf{B}_{(j)}^b \in \{0, 1\}^b$ , is encoded and modulated into the symbol vector  $\mathbf{S}_{(j)}^t$  which is transmitted over the channel. The purpose of such channel coding is to facilitate the recovery of the information bits at the receiver. A channel code can consist of both FEC codes, such as Reed-Solomon (RS) codes, as well as error detection codes, such as checksums and cyclic redundancy checks. The receiver uses the recovered symbols  $\hat{\mathbf{S}}_{(j)}^t$  to decode the information bits  $\hat{\mathbf{B}}_{(j)}^b \in \{0, 1\}^b$ . The presence of FEC codes implies that even when some of the symbols are

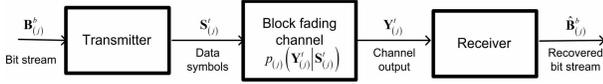


Fig. 4. Coded communications over block-fading channels.

not correctly recovered, i.e.,  $\mathbf{S}_{(j)}^t \neq \hat{\mathbf{S}}_{(j)}^t$ , the information bits can still be perfectly decoded. Error detection codes allow the receiver to detect whether the information bits were correctly recovered and estimate the number of bit errors, i.e., the Hamming distance between  $\mathbf{B}_{(j)}^b$  and  $\hat{\mathbf{B}}_{(j)}^b$ , denoted  $\epsilon_{(j)}$ .

In block-fading channels, each block of transmitted channel symbols undergoes a different channel. In particular, for the  $j$ th block,  $j \in \mathcal{N}$ , the conditional distribution of the channel output  $\mathbf{Y}_{(j)}^t$  given the input  $\mathbf{S}_{(j)}^t$  represents a finite-memory causal stationary channel as described in Subsection II-A. However, this conditional distribution can change between different blocks, namely, it depends on the block index  $j$ . An illustration of this channel model is depicted in Fig. 4. Block-fading channels model dynamic environments in which the transmitted bursts are shorter than the coherence duration of the channel, faithfully representing many scenarios of interest in wireless communications [42].

Since each block undergoes a different channel input-output relationship, the receiver must be able to track and adapt to the varying channel conditions in order to optimize its performance. In the next subsection, we show how the fact that ViterbiNet requires relatively small training sets can be combined with the presence of coded communications such that the receiver can track channel variations online.

### B. ViterbiNet with Online Training

We next discuss how coded communications can allow ViterbiNet to train and track channel variations online. Our proposed approach is inspired by concepts from meta-learning [25], which is a field of research focusing on self-teaching ML algorithms, as well as from decision-directed adaptive filter theory, which considers the blind adaptation of adjustable filters based on some decision mechanism, see, e.g., [43].

ViterbiNet uses the channel outputs over a single block, denoted  $\mathbf{Y}_{(j)}^t$ , where  $j \in \mathcal{N}$  represents the block index, to generate an estimate of the transmitted symbols, denoted  $\hat{\mathbf{S}}_{(j)}^t$ . In the presence of FEC codes, when the number of detection errors is not larger than the minimal distance of the code, the encoded bits can still be perfectly recovered [44, Ch. 8]. It therefore follows that when the symbol error rate (SER) of ViterbiNet is small enough such that the FEC code is able to compensate for these errors, the transmitted message can be recovered. The recovered message may also be re-encoded to generate new training samples, denoted  $\tilde{\mathbf{S}}_{(j)}^t$ , referred to henceforth as *meta-training*. If indeed the encoded bits were successfully recovered, then  $\tilde{\mathbf{S}}_{(j)}^t$  represents the true channel input from which  $\mathbf{Y}_{(j)}^t$  was obtained. Consequently, the pair  $\tilde{\mathbf{S}}_{(j)}^t$  and  $\mathbf{Y}_{(j)}^t$  can be used to re-train ViterbiNet. The simple DNN structure used in ViterbiNet implies that it can be efficiently and quickly retrained with a relatively

---

### Algorithm 2 Online Training of ViterbiNet

---

- 1: **Input:** Block of channel outputs  $\mathbf{y}^t$ ; ViterbiNet DNN weights  $\theta$  and mixture model parameters  $\varphi$ ; Bit error threshold  $\bar{\epsilon}$ .
  - 2: Apply ViterbiNet with parameters  $\theta, \varphi$  to recover the symbols  $\hat{\mathbf{s}}^t$  from  $\mathbf{y}^t$ .
  - 3: Decode bits vector  $\hat{\mathbf{b}}^b$  and number of error  $\epsilon$  from  $\hat{\mathbf{s}}^t$ .
  - 4: **if**  $\epsilon < \bar{\epsilon}$  **then**
  - 5:   Encode and modulate  $\hat{\mathbf{b}}^b$  into  $\tilde{\mathbf{s}}^t$ .
  - 6:   Re-train DNN weights  $\theta$  using meta-training  $(\tilde{\mathbf{s}}^t, \mathbf{y}^t)$ .
  - 7:   Update mixture model  $\varphi$  using  $\mathbf{y}^t$ .
  - 8: **end if**
  - 9: **Output:** Decoded bits  $\hat{\mathbf{b}}^b$ ; ViterbiNet DNN weights  $\theta$  and mixture model parameters  $\varphi$ .
- 

small number of training samples, such as the size of a single transmission block in a block-fading channel. An illustration of the proposed online training mechanism is depicted in Fig. 5.

A natural drawback of decision-directed approaches is their sensitivity to decision errors. For example, if the FEC decoder fails to successfully recover the encoded bits, then the meta-training samples  $\tilde{\mathbf{S}}_{(j)}^t$  do not accurately represent the channel input resulting in  $\mathbf{Y}_{(j)}^t$ . In such cases, the inaccurate training sequence may gradually deteriorate the recovery accuracy of the re-trained ViterbiNet, making the proposed approach unreliable, particularly in low signal-to-noise ratios (SNRs) where decoding errors occur frequently. Nonetheless, when error detection codes are present in addition to FEC codes, the effect of decision errors can be mitigated. In particular, when the receiver has a reliable estimate of  $\epsilon_{(j)}$ , namely, the number of bit errors in decoding the  $j$ th block, it can decide to generate meta-training and re-train the symbol detector only when the number of errors is smaller than some threshold  $\bar{\epsilon}$ . Using this approach, only accurate meta-training is used, and the effect of decision errors can be mitigated and controlled. The proposed online training mechanism is summarized below as Algorithm 2.

Setting the value of  $\bar{\epsilon}$  should be based on the nature of the expected variations between blocks as well as on the FEC scheme. When operating with gradual variations, one can set  $\bar{\epsilon}$  to represent zero errors, guaranteeing that the meta-training symbols used for re-training ViterbiNet are accurate. However, this setting may limit the ability of ViterbiNet to track moderate channel variations. Using larger values of  $\bar{\epsilon}$  should also account for the FEC method used, as for some codes, encoding a bit stream with a small number of errors may result in a distant codeword compared to the one transmitted, thus yielding inaccurate meta-training labels. In Subsection V-B we numerically evaluate the effect of different values of  $\bar{\epsilon}$  on the ability of ViterbiNet to reliably track block-fading channels.

The learning rate used in re-training the DNN in Step 6 of Algorithm 2 should be set based on the expected magnitude of the channel variations. Specifically, the value of the learning rate balances the contributions of the initial DNN weights, which were tuned based on a past training set corresponding to previous channel realizations, and that of the typically much

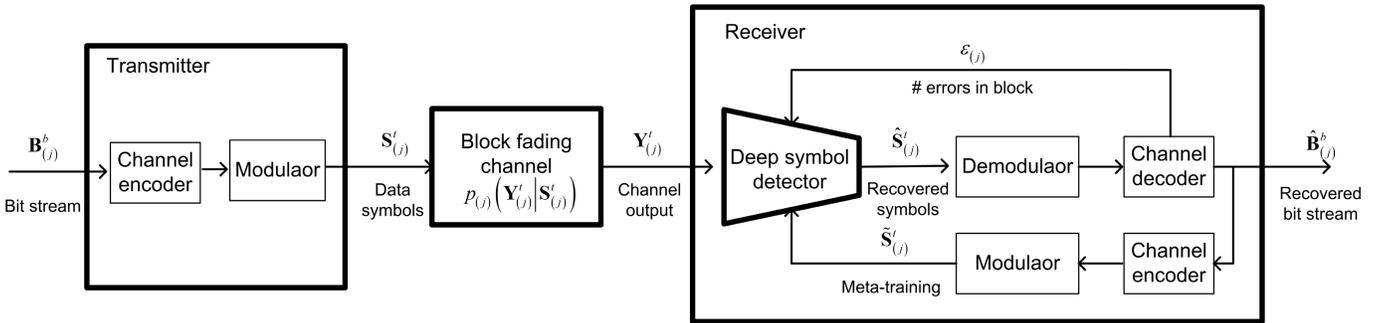


Fig. 5. Online training model.

smaller meta-training set, representing the current channel. Using larger learning rates implies that the network can adapt to nontrivial channel variations, but also increases the sensitivity of the algorithm to decision errors, and should thus be used with smaller values of the threshold  $\bar{\epsilon}$ .

The method for updating the mixture model parameters  $\varphi$  in Step 7 of Algorithm 2 depends on which of the possible finite mixture model estimators discussed in Subsection III-A is used. For example, when using EM-based model fitting, the previous model parameters  $\varphi$  can be used as the initial guess used for fitting the PDF based on the current channel outputs.

Finally, we note that our meta-learning based online training scheme exploits only the structure induced by channel codes. Modern communication protocols dictate a multitude of additional structures that can be exploited to generate meta-training. For example, many communication protocols use a-priori known preamble sequences to facilitate synchronization between the communicating nodes [44, Ch. 14.2]. The receiver can thus use the a-priori knowledge of the preamble to generate meta-training from the corresponding channel outputs, which can be used for re-training ViterbiNet online. Additional structures induced by protocols that can be exploited to re-train ViterbiNet in a similar manner include specific header formats, periodic pilots, and network management frames. Consequently, while the proposed approach focuses on using channel codes to generate meta-training, it can be extended to utilize various other inherent structures of digital communication signals.

## V. NUMERICAL STUDY

In this section we numerically evaluate the performance of ViterbiNet. In particular, we first consider time-invariant channels in Subsection V-A. For such setups, we numerically compare the performance of ViterbiNet to the conventional model-based Viterbi algorithm as well as to previously proposed deep symbol detectors, and evaluate the resiliency of ViterbiNet to inaccurate training. Then, we consider block-fading channels, and evaluate the performance of ViterbiNet with online training in Subsection V-B. In the following we assume that the channel memory  $l$  is a-priori known. It is emphasized that for all the simulated setups we were able to accurately detect the memory length using a standard

correlation peak based estimator,<sup>1</sup> see, e.g., [29], hence this assumption is not essential. Throughout this numerical study we implement the fully connected network in Fig. 3 using three layers: a  $1 \times 100$  layer followed by a  $100 \times 50$  layer and a  $50 \times 16 (= |m|^l)$  layer, using intermediate sigmoid and ReLU activation functions, respectively. The mixture model estimator approximates the distribution as a Gaussian mixture using EM-based fitting [33, Ch. 2]. The network is trained using 5000 training samples to minimize the cross-entropy loss via the Adam optimizer [45] with learning rate 0.01, using up to 100 epochs with mini-batch size of 27 observations. We note that the number of training samples is of the same order and even smaller compared to typical preamble sequences in wireless communications.<sup>2</sup> Due to the small number of training samples and the simple architecture of the DNN, only a few minutes are required to train the network on a standard CPU.

### A. Time-Invariant Channels

We first consider two finite-memory causal scalar channels: An intersymbol interference (ISI) channel with additive white Gaussian noise (AWGN), and a Poisson channel. In both channels we set the memory length to  $l = 4$ . For the ISI channel with AWGN, we let  $W[i]$  be a zero-mean unit variance AWGN independent of  $S[i]$ , and let  $\mathbf{h}(\gamma) \in \mathcal{R}^l$  be the channel vector representing an exponentially decaying profile given by  $(\mathbf{h})_\tau \triangleq e^{-\gamma(\tau-1)}$  for  $\gamma > 0$ . The input-output relationship is given by

$$Y[i] = \sqrt{\rho} \cdot \sum_{\tau=1}^l (\mathbf{h}(\gamma))_\tau S[i - \tau + 1] + W[i], \quad (7)$$

where  $\rho > 0$  represents the SNR. Note that even though the effect of the ISI becomes less significant as  $\gamma$  grows, the channel (7) has memory length of  $l = 4$  for all values of  $\gamma$  considered in this study. The channel input is randomized from a binary phase shift keying constellation, i.e.,  $\mathcal{S} = \{-1, 1\}$ . For the Poisson channel, the channel input represents on-off

<sup>1</sup>We note that correlation based channel length estimators assume that the transmitter and the receiver are synchronized in time.

<sup>2</sup>For example, preamble sequences in fourth-generation cellular LTE systems consist of up to 10 subframes of  $2.4 \cdot 10^4$  samples, each embedding 839 symbols [34, Ch. 17].

keying, namely,  $\mathcal{S} = \{0, 1\}$ , and the channel output  $Y[i]$  is generated from the input via

$$Y[i]|S^t \sim \mathbb{P} \left( \sqrt{\rho} \cdot \sum_{\tau=1}^l (\mathbf{h}(\gamma))_{\tau} S[i - \tau + 1] + 1 \right), \quad (8)$$

where  $\mathbb{P}(\lambda)$  is the Poisson distribution with parameter  $\lambda > 0$ .

For each channel, we numerically compute the SER of ViterbiNet for different values of the SNR parameter  $\rho$ . In the following study, the DNN is trained anew for each value of  $\rho$ . For each SNR  $\rho$ , the SER values are averaged over 20 different channel vectors  $\mathbf{h}(\gamma)$ , obtained by letting  $\gamma$  vary in the range  $[0.1, 2]$ . For comparison, we numerically compute the SER of the Viterbi algorithm, as well as that of the sliding bidirectional RNN (SBRNN) deep symbol decoder proposed in [15]. In order to study the resiliency of ViterbiNet to inaccurate training, we also compute the performance when the receiver only has access to a noisy estimate of  $\mathbf{h}(\gamma)$ , and specifically, to a copy of  $\mathbf{h}(\gamma)$  whose entries are corrupted by i.i.d. zero-mean Gaussian noise with variance  $\sigma_e^2$ . In particular, we use  $\sigma_e^2 = 0.1$  for the Gaussian channel (7), and  $\sigma_e^2 = 0.08$  for the Poisson channel (8). We consider two cases: *Perfect CSI*, in which the model-based Viterbi detector has accurate knowledge of  $\mathbf{h}(\gamma)$ , while ViterbiNet is trained using labeled samples generated with the same  $\mathbf{h}(\gamma)$  used for generating the test data; and *CSI uncertainty*, where the Viterbi algorithm implements Algorithm 1 with the log-likelihoods computed using the noisy version of  $\mathbf{h}(\gamma)$ , while the labeled data used for training ViterbiNet is generated with the noisy version of  $\mathbf{h}(\gamma)$  instead of the true one. In particular, for CSI uncertainty, the 5000 samples used for training ViterbiNet are divided into 10 subsets, each generated from a channel with a different realization of the noise in  $\mathbf{h}(\gamma)$ . In all cases, the information symbols are uniformly randomized in an i.i.d. fashion from  $\mathcal{S}$ , and the test samples are generated from their corresponding channel, i.e., (7) and (8) for the Gaussian and Poisson channels, respectively, with the true channel vector  $\mathbf{h}(\gamma)$ .

The numerically computed SER values, averaged over 50000 Monte Carlo simulations, versus  $\rho \in [-6, 10]$  dB for the ISI channel with AWGN are depicted in Fig. 6, while the corresponding performance versus  $\rho \in [10, 34]$  dB for the Poisson channel are depicted in Fig. 7. Observing Figs. 6-7, we note that the performance of ViterbiNet approaches that of the conventional CSI-based Viterbi algorithm. In particular, for the AWGN case, in which the channel output obeys a Gaussian mixture distribution, the performance of ViterbiNet coincides with that of the Viterbi algorithm, while for the Poisson channel a very small gap is observed at high SNRs due to the model mismatch induced by approximating the distribution of  $Y[i]$  as a Gaussian mixture. We expect the performance of ViterbiNet to improve in that scenario when using a Poisson kernel density estimator for  $\hat{p}_{\varphi}(y[i])$ , however, this requires some basic level of knowledge of the input-output relationship. It is also observed in Figs. 6-7 that the SBRNN receiver, which was shown in [15] to approach the performance of the CSI-based Viterbi algorithm when sufficient training is provided, is outperformed by ViterbiNet here due to the small amount of training data provided. For example, for the

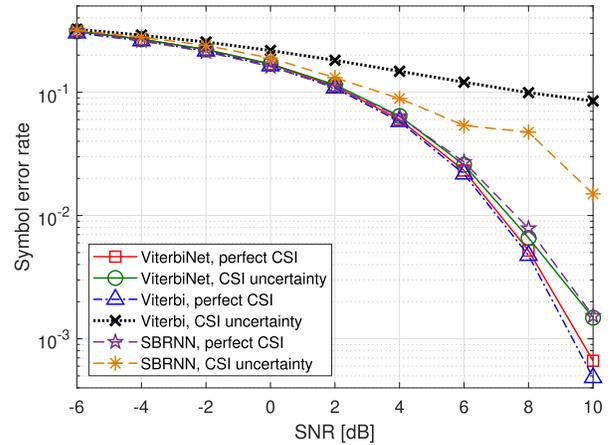


Fig. 6. SER versus SNR, ISI channel with AWGN.

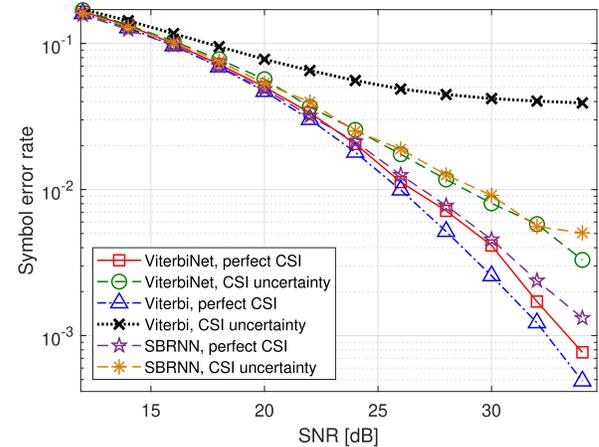


Fig. 7. SER versus SNR, Poisson channel.

ISI channel with AWGN, both the Viterbi detector as well as ViterbiNet achieve an SER of  $4.7 \cdot 10^{-3}$  for SNR of 8 dB, while the SBRNN detector achieves an SER of  $8.5 \cdot 10^{-3}$  for the same SNR value. For the Poisson channel at an SNR of 28 dB, the Viterbi algorithm achieves an SER of  $5.1 \cdot 10^{-3}$ , while ViterbiNet and the SBRNN detector achieve SER values of  $6.6 \cdot 10^{-3}$  and  $8.3 \cdot 10^{-3}$ , respectively. These results demonstrate that ViterbiNet, which uses simple DNN structures embedded into the Viterbi algorithm, requires significantly less training compared to previously proposed ML-based receivers.

In the presence of CSI uncertainty, it is observed in Figs. 6-7 that ViterbiNet significantly outperforms the Viterbi algorithm. In particular, when ViterbiNet is trained with a variety of different channel conditions, it is still capable of achieving relatively good SER performance under each of the channel conditions for which it is trained, while the performance of the conventional Viterbi algorithm is significantly degraded in the presence of imperfect CSI. While the SBRNN receiver is shown to be more resilient to inaccurate CSI compared to the Viterbi algorithm, as was also observed in [15], it is outperformed by ViterbiNet with the same level of uncertainty, and the performance gap is more notable in the AWGN channel. The reduced gain of ViterbiNet over the SBRNN receiver for

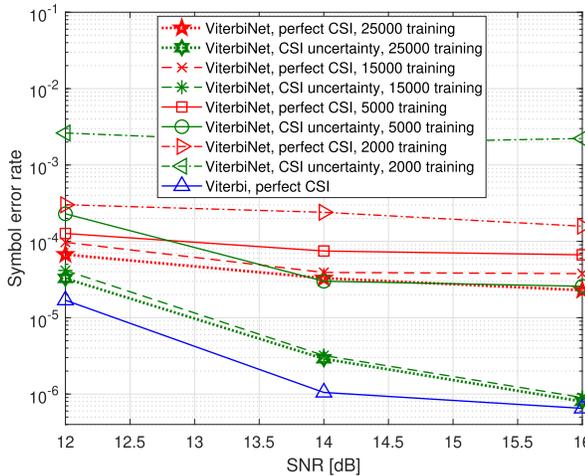


Fig. 8. High-SNR performance for different training sizes, ISI channel with AWGN.

the Poisson channel stems from the fact that ViterbiNet here uses a Gaussian mixture kernel density estimator for the PDF of  $Y[i]$ , which obeys a Poisson mixture distribution for the Poisson channel (8).

The simulation results depicted in Figs. 6-7 focus on SNRs for which the SER values are of the orders of  $10^{-2}$  and  $10^{-3}$ . While it is observed in Figs. 6-7 that ViterbiNet approaches the performance of the Viterbi algorithm with perfect CSI, it is also noted that there is a small SER gap between these symbol detector in higher SNRs. This gap indicates that training is more difficult in higher SNRs, for which there is a smaller diversity in the training samples compared to lower SNR values, and thus larger training sets are required to improve performance. To demonstrate this, we depict in Fig. 8 the SER values achieved by ViterbiNet when trained using training sets of sizes  $\{2, 5, 15, 25\} \cdot 10^3$  for the ISI channel with AWGN with SNRs of 12–16 dB, i.e., higher SNRs than those depicted in Fig. 6. The numerically evaluated performance, averaged here over  $10^6$  Monte Carlo simulations to faithfully capture SER values in the order of  $10^{-5}$ , is compared to the model-based Viterbi algorithm with perfect CSI. Observing Fig. 8, we note that in high SNRs, ViterbiNet requires larger training sets in order to approach the performance of the Viterbi algorithm, when trained using samples from the same channel under which it is tested, i.e., perfect CSI. However, it is also noted that, except for very small training sets, ViterbiNet with CSI uncertainty outperforms ViterbiNet with perfect CSI, and that its performance is within a small gap from that of the Viterbi detector. These results indicate that, in high SNRs, the additional diversity induced when training under CSI uncertainty is in fact beneficial, and allows ViterbiNet to achieve smaller SER values compared to training using samples from the same channel realization.

The main benefit of ViterbiNet is its ability to accurately implement the Viterbi algorithm in causal finite-memory channels while requiring only knowledge of the memory length, i.e., no CSI is required. The simulation study presented in Figs. 6-7 demonstrates an additional gain of ViterbiNet over the Viterbi algorithm, which is its improved resiliency to

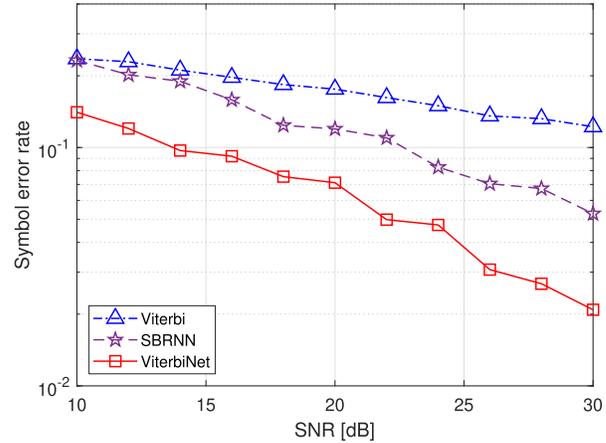


Fig. 9. SER versus SNR, ISI channel with alpha-stable noise.

inaccurate CSI. Another benefit of ViterbiNet is its ability to reliably operate in setups where, even when full instantaneous CSI is available, the Viterbi algorithm is extremely complicated to implement. To demonstrate this benefit, we again consider the ISI channel in (7) with an exponential decay parameter  $\gamma = 0.2$ , however, here we set the distribution of the additive noise signal  $W[i]$  to an *alpha-stable* distribution [46, Ch. 1]. Alpha-stable distributions, which are obtained as a sum of heavy tailed independent RVs via the generalized central limit theorem, are used to model the noise in various communications scenarios, including molecular channels [47], underwater acoustic channels [48], and impulsive noise channels [49]. In particular, we simulate an alpha-stable noise with stability parameter  $\alpha = 0.5$ , skewness parameter  $\beta = 0.75$ , scale parameter  $c = 1$ , and location parameter  $\mu = 0$ , following [47, Ch. III]. Recall that the PDF of alpha-stable RVs is in general not analytically expressible [46, Ch. 1.4]. Thus, implementing the Viterbi algorithm for such channels is extremely difficult even when full CSI is available, as it requires accurate numerical approximations of the conditional PDFs to compute the log likelihoods. Consequently, in order to implement the Viterbi algorithm here, we numerically approximated the PDF of such alpha-stable RVs using the direct integration method [46, Ch. 3.3] over a grid of 50 equally spaced points in the range  $[-5, 5]$ , and used the PDF of the closest grid point when computing the log-likelihood terms in the Viterbi detector.

In Fig. 9 we depict the SER performance of ViterbiNet compared to the SBRNN detector and the Viterbi algorithm for this channel versus  $\rho \in [10, 30]$  dB. Observing Fig. 9, we note that the CSI-based Viterbi detector, which has to use a numerical approximation of the PDF, is notably outperformed by the data-driven symbol detectors, which are ignorant of the exact statistical model and do not require an analytical expression of the conditional PDF to operate. Furthermore, we note that ViterbiNet achieves a substantial SNR gap from the SBRNN detector. For example, ViterbiNet achieves an SER smaller than  $5 \cdot 10^{-2}$  for SNR values larger than 22 dB, while the SBRNN detector requires an SNR of at least 30 dB to achieve a similar SER, namely, an SNR gap of 8 dB. Finally, the performance of ViterbiNet scales with

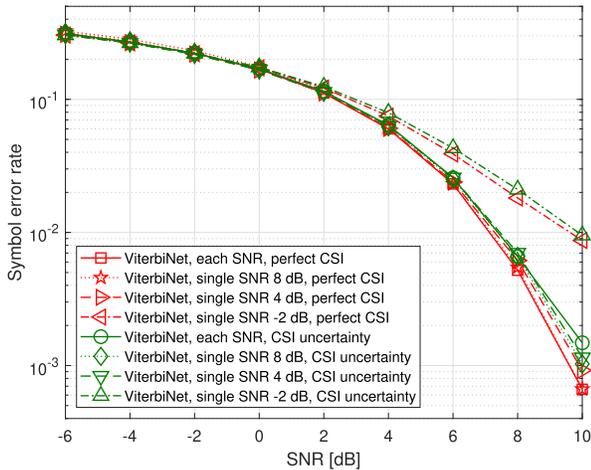


Fig. 10. Training for each SNR compared to training over a single SNR, ISI channel with AWGN.

respect to SNR in a similar manner as in Figs. 6-7. The results in Fig. 9 demonstrate the ability of ViterbiNet to reliably operate under non-conventional channel models for which the Viterbi algorithm is difficult to accurately implement, even when full CSI is available.

Finally, we study the robustness of ViterbiNet to different noise powers than those used during training. Recall that so far we have trained ViterbiNet for each value of  $\rho$ , i.e., for each SNR. To evaluate its resiliency to an inaccurate noise level in training, we next numerically compute the SER of ViterbiNet for different values of the SNR parameter  $\rho$ . In particular, we repeat the simulation study of the ISI channel with AWGN, whose results are depicted in Fig. 6, while training ViterbiNet using 5000 samples drawn from a channel with a single SNR, i.e., ViterbiNet is trained once instead of being re-trained for each SNR. The numerically computed SER values are depicted in Fig. 10, where the curves denoted ‘ViterbiNet, single SNR’ represent the performance achieved when training only with samples corresponding to a single SNR, which is in the set  $\{-2, 4, 8\}$  dB; The curves denoted ‘ViterbiNet, each SNR’ stand for the performance achieved when ViterbiNet is trained anew for each SNR level, i.e., the same curves as those of Fig. 6 of the revised manuscript. Observing Fig. 10, we note that ViterbiNet achieves approximately the same performance when trained using a single SNR value of either 4 dB or 8 dB as when it is trained anew for each SNR. A small gap is observed when ViterbiNet is trained using samples corresponding to a low SNR level of  $-2$  dB. These results demonstrate the ability of ViterbiNet to generalize to multiple SNR levels.

### B. Block-Fading Channels

Next, we numerically evaluate the online training method for tracking block-fading channels detailed in Section IV. To that aim, we consider two block-fading channels: an ISI channel with AWGN and a Poisson channel. For each transmitted block we use the channel models in (7) and (8), respectively, with a fixed exponential decay parameter

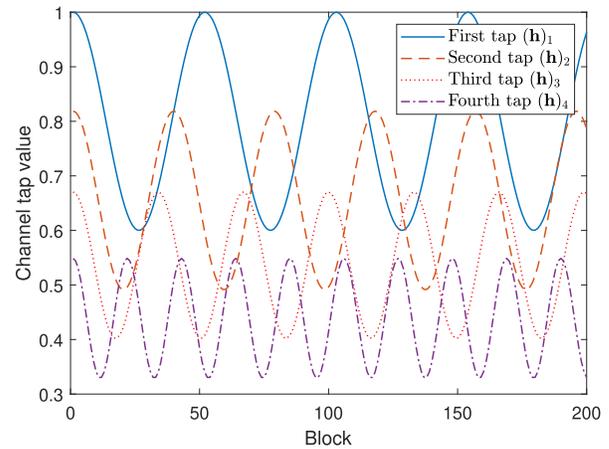


Fig. 11. Channel coefficient variations between different blocks.

$\gamma = 0.2$ , and thus we omit the notation  $\gamma$  from the vector  $\mathbf{h}$  in this subsection. However, unlike the scenario studied in the previous subsection, here the channel coefficient vector  $\mathbf{h}$  varies between blocks. In particular, for the  $j$ th block we use  $(\mathbf{h}_{(j)})_{\tau} \triangleq e^{-0.2(\tau-1)} \cdot \left(0.8 + 0.2 \cos\left(\frac{2\pi \cdot j}{p}\right)\right)$ , with  $\mathbf{p} = [51, 39, 33, 21]^T$ . An illustration of the variations in the channel coefficients with respect to the block index  $j$  is depicted in Fig. 11.

In each channel block a single codeword of an RS [255, 223] FEC code is transmitted, namely, during each block  $b = 1784$  bits are conveyed using  $t = 2040$  binary symbols. Before the first block is transmitted, ViterbiNet is trained using 5000 training samples taken using the initial channel coefficients, i.e., using the initial channel coefficients vector  $\mathbf{h}_{(1)}$ . For online training, we use a learning rate of 0.002, i.e., five times smaller than that used for initial training, and a bit error threshold of  $\bar{\epsilon} = 2\%$ . In addition to ViterbiNet with online training, we also compute the coded bit error rate (BER) performance of ViterbiNet when trained only once using the 5000 training samples representing the initial channel  $\mathbf{h}_{(1)}$ , referred to as *initial training*, as well as ViterbiNet trained once using 5000 training samples corresponding to 10 different channels  $\{\mathbf{h}_{(3,k)}\}_{k=1}^{10}$ , referred to as *composite training*. The performance of ViterbiNet is compared to the Viterbi algorithm with full instantaneous CSI as well as to the Viterbi detector which assumes that the channel is the time-invariant initial channel  $\mathbf{h}_{(1)}$ . The coded BER results, averaged over 200 consecutive blocks, are depicted in Figs. 12-13 for the ISI channel with AWGN and for the Poisson channel, respectively.

Observing Figs. 12-13, we note that for both channels, as the SNR increases, ViterbiNet with online training approaches the performance of the CSI-based Viterbi algorithm which requires accurate knowledge of the complete input-output statistical relationship for each block, while for low SNRs, its performance is only slightly improved compared to that of the Viterbi algorithm using only the initial channel. This can be explained by noting that for high SNR values, the BER at each block is smaller than the threshold  $\bar{\epsilon}$ , and thus our proposed online training scheme is capable of generating reliable

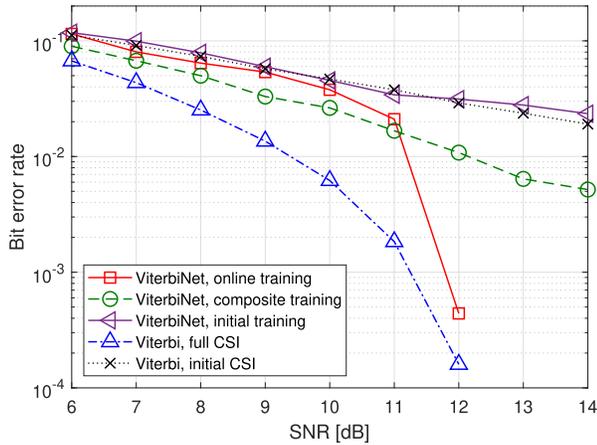


Fig. 12. Coded BER versus SNR, block-fading ISI channel with AWGN.

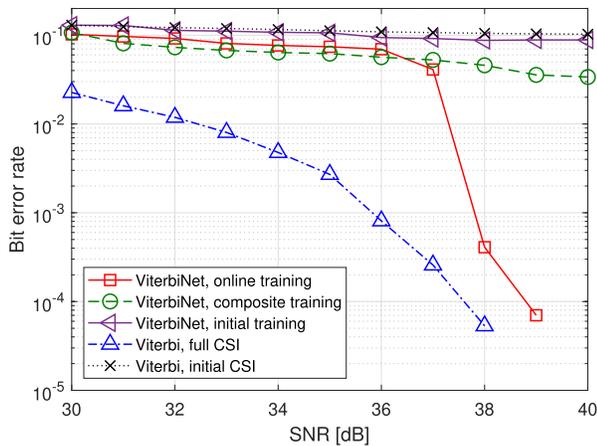


Fig. 13. Coded BER versus SNR, block-fading Poisson channel.

meta-training, which ViterbiNet uses to accurately track the channel, thus approaching the optimal performance. However, for low SNR values, the BER is typically larger than  $\bar{\epsilon}$ , thus ViterbiNet does not frequently update its coefficients in real-time, thus achieving only a minor improvement over using only the initial training data set. It is emphasized that increasing the BER threshold  $\bar{\epsilon}$  can severely deteriorate the performance of ViterbiNet at low SNRs, as it may be trained online using inaccurate meta-training.

We also observe in Figs. 12-13 that the performance ViterbiNet with initial training coincides with that of Viterbi algorithm using the initial channel, as both detectors effectively implement the same detection rule, which is learned by ViterbiNet from the initial channel training set. Furthermore, the composite training approach allows ViterbiNet to achieve improved BER performance compared to using only initial training, as the resulting decoder is capable of operating in a broader range of different channel conditions. Still, it is noted that for both the AWGN channel and the Poisson channel, ViterbiNet with composite training is notably outperformed by the optimal Viterbi detector with instantaneous CSI, whose BER performance is approached only when using online training at high SNRs.

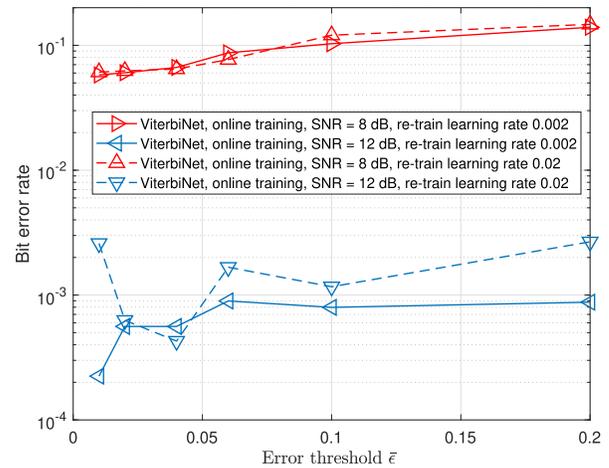


Fig. 14. Coded BER versus error threshold, block-fading ISI channel with AWGN.

As discussed in Subsection IV-B, the success of the meta-learning based method for training ViterbiNet online depends on the error threshold  $\bar{\epsilon}$  and the learning rate used for re-training the DNN. To numerically evaluate the effect of these parameters, we compare in Fig. 14 the BER of ViterbiNet versus  $\bar{\epsilon}$  for the ISI with AWGN channel under SNR values of 8 and 12 dB, and with re-training learning rates of 0.002 and 0.02. These BER values are compared to the model-based Viterbi operating with full CSI. Observing Fig. 14 we note that, for the lower SNR of 8 dB, a similar behavior, in which the BER value increases with  $\bar{\epsilon}$ , is observed for both learning rates. However, for the higher SNR of 12 dB, the BER increase with the error threshold  $\bar{\epsilon}$  for learning rate of 0.002, while for the higher learning rate of 0.02 we observe a minima around  $\bar{\epsilon} = 0.04$ . This behavior is because higher values of  $\bar{\epsilon}$  may cause the network to use inaccurate training, while values which are too small may result in the network being unable to reliably track channel variations.

The results presented in this section demonstrate that a data-driven symbol detector designed by integrating DNNs into the Viterbi algorithm is capable of approaching optimal performance without CSI. Unlike the CSI-based Viterbi algorithm, its ML-based implementation demonstrates robustness to imperfect CSI, and achieves excellent performance in channels where the Viterbi algorithm is difficult to implement. The fact that ViterbiNet requires a relatively small training set allows it to accurately track block-fading channels using meta-training generated from previous decisions on each transmission block. We thus conclude that designing a data-driven symbol detector by combining ML-based methods with the optimal model-based algorithm yields a reliable, efficient, and resilient data-driven symbol recovery mechanism.

## VI. CONCLUSION

We proposed an ML-based implementation of the Viterbi symbol detection algorithm called ViterbiNet. To design ViterbiNet, we identified the log-likelihood computation as the part of the Viterbi algorithm that requires full knowledge

of the underlying channel input-output statistical relationship. We then integrated an ML-based architecture designed to compute the log-likelihoods into the algorithm. The resulting architecture combines ML-based processing with the conventional Viterbi symbol detection scheme. In addition, we proposed a meta-learning based method that allows ViterbiNet to track time-varying channels online. Our numerical results demonstrate that ViterbiNet approaches the optimal performance of the CSI-based Viterbi algorithm and outperforms previously proposed ML-based symbol detectors using a small amount of training data. It is also illustrated that ViterbiNet is capable of reliably operating in the presence of CSI uncertainty, as well as in complex channel models where the Viterbi algorithm is extremely difficult to implement. Finally, it is shown that ViterbiNet can adapt via meta-learning in block-fading channel conditions.

## REFERENCES

- [1] A. Viterbi, "Error bounds for convolutional codes and an asymptotically optimum decoding algorithm," *IEEE Trans. Inf. Theory*, vol. 13, no. 2, pp. 260–269, Apr. 1967.
- [2] G. D. Forney, Jr., "The Viterbi algorithm," *Proc. IEEE*, vol. 61, no. 3, pp. 268–278, Mar. 1973.
- [3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [4] Y. Bengio, "Learning deep architectures for AI," *Found. Trends Mach. Learn.*, vol. 2, no. 1, pp. 1–127, 2009.
- [5] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proc. ICML*, Haifa, Israel, Jun. 2010, pp. 399–406.
- [6] S. Khobahi, N. Naimipour, M. Soltanalian, and Y. C. Eldar, "Deep signal recovery with one-bit quantization," in *Proc. ICASSP*, Brighton, U.K., May 2019, pp. 2987–2991.
- [7] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Trans. Cogn. Commun. Netw.*, vol. 3, no. 4, pp. 563–575, Dec. 2017.
- [8] O. Simeone, "A very brief introduction to machine learning with applications to communication systems," *IEEE Trans. Cogn. Commun. Netw.*, vol. 4, no. 4, pp. 648–664, Dec. 2018.
- [9] Q. Mao, F. Hu, and Q. Hao, "Deep learning for intelligent wireless networks: A comprehensive survey," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 4, pp. 2595–2621, Jun. 2018.
- [10] D. Gunduz, P. de Kerret, N. D. Sidiropoulos, D. Gesbert, C. Murthy, and M. van der Schaar, "Machine learning in the air," 2019, *arXiv:1904.12385*. [Online]. Available: <https://arxiv.org/abs/1904.12385>
- [11] E. Nachmani, E. Marciano, L. Lugosch, W. J. Gross, D. Burshtein, and Y. Be'ery, "Deep learning methods for improved decoding of linear codes," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 1, pp. 119–131, Feb. 2018.
- [12] T. Gruber, S. Cammerer, J. Hoydis, and S. T. Brink, "On deep learning-based channel decoding," in *Proc. CISS*, Baltimore, MD, USA, Mar. 2017, pp. 1–6.
- [13] N. Samuel, T. Diskin, and A. Wiesel, "Deep MIMO detection," in *Proc. SPAWC*, Sapporo, Japan, Jul. 2017, pp. 1–5.
- [14] A. Caciularu and D. Burshtein, "Blind channel equalization using variational autoencoders," 2018, *arXiv:1803.01526*. [Online]. Available: <https://arxiv.org/abs/1803.01526>
- [15] N. Farsad and A. Goldsmith, "Neural network detection of data sequences in communication systems," *IEEE Trans. Signal Process.*, vol. 66, no. 21, pp. 5663–5678, Nov. 2018.
- [16] F. Liang, C. Shen, and F. Wu, "An iterative BP-CNN architecture for channel decoding," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 1, pp. 144–159, Feb. 2018.
- [17] H. Kim, Y. Jiang, R. Rana, S. Kannan, S. Oh, and P. Viswanath, "Communication algorithms via deep learning," 2018, *arXiv:1805.09317*. [Online]. Available: <https://arxiv.org/abs/1805.09317>
- [18] X.-A. Wang and S. Wicker, "An artificial neural net Viterbi decoder," *IEEE Trans. Commun.*, vol. 44, no. 2, pp. 165–171, Feb. 1996.
- [19] M. Shohat, G. N. T. Shlezinger, and Y. C. Eldar, "Deep quantization for MIMO channel estimation," in *Proc. ICASSP*, Brighton, U.K., May 2019, pp. 3912–3916.
- [20] J. R. Hershey, J. Le Roux, and F. Weninger, "Deep unfolding: Model-based inspiration of novel deep architectures," 2014, *arXiv:1409.2574*. [Online]. Available: <https://arxiv.org/abs/1409.2574>
- [21] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," 2019, *arXiv:1912.10557*. [Online]. Available: <https://arxiv.org/abs/1912.10557>
- [22] P. Sprechmann, A. M. Bronstein, and G. Sapiro, "Supervised non-Euclidean sparse NMF via bilevel optimization with applications to speech enhancement," in *Proc. HSCMA*, Villers-les-Nancy, France, May 2014, pp. 11–15.
- [23] Y. Li, M. Tofighi, J. Geng, V. Monga, and Y. C. Eldar, "An algorithm unrolling approach to deep blind image deblurring," 2019, *arXiv:1902.03493*. [Online]. Available: <https://arxiv.org/abs/1902.03493>
- [24] O. Solomon, R. Cohen, Y. Zhang, Y. Yang, H. Qiong, J. Luo, R. J. G. van Sloun, and Y. C. Eldar, "Deep unfolded robust PCA with application to clutter suppression in ultrasound," 2018, *arXiv:1811.08252*. [Online]. Available: <https://arxiv.org/abs/1811.08252>
- [25] C. Lemke, M. Budka, and B. Gabrys, "Metalearning: A survey of trends and technologies," *Artif. Intell. Rev.*, vol. 44, no. 1, pp. 117–130, Jun. 2015.
- [26] S. Park, H. Jang, O. Simeone, and J. Kang, "Learning how to demodulate from few pilots via meta-learning," 2019, *arXiv:1903.02184*. [Online]. Available: <https://arxiv.org/abs/1903.02184>
- [27] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [28] D. D. Falconer and F. R. Magee, "Adaptive channel memory truncation for maximum likelihood sequence estimation," *Bell Syst. Tech. J.*, vol. 52, no. 9, pp. 1541–1562, Nov. 1973.
- [29] V. D. Nguyen, H.-P. Kuchenbecker, H. Haas, K. Kyamakya, and G. Gelle, "Channel impulse response length and noise variance estimation for OFDM systems with adaptive guard interval," *EURASIP J. Wireless Commun. Netw.*, vol. 2007, Dec. 2007, Art. no. 024342.
- [30] A. van den Oord *et al.*, "WaveNet: A generative model for raw audio," 2016, *arXiv:1609.03499*. [Online]. Available: <https://arxiv.org/abs/1609.03499>
- [31] A. van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, and A. Graves, "Conditional image generation with PixelCNN decoders," in *Proc. NIPS*, 2016, pp. 4790–4798.
- [32] M. Rosenblatt, "Remarks on some nonparametric estimates of a density function," *Ann. Math. Statist.*, vol. 27, no. 3, pp. 832–837, Sep. 1956.
- [33] G. McLachlan and D. Peel, *Finite Mixture Models*. Hoboken, NJ, USA: Wiley, 2004.
- [34] E. Dahlman, S. Parkvall, J. Skold, and P. Beming, *3G Evolution: HSPA and LTE for Mobile Broadband*. New York, NY, USA: Academic, 2010.
- [35] C. Bishop, "Mixture density networks," Aston Univ., Birmingham, U.K., Tech. Rep. NCRG/4288, 1994.
- [36] N. Shlezinger, R. Fu, and Y. C. Eldar, "DeepSIC: Deep soft interference cancellation for multiuser MIMO detection," 2020, *arXiv:2002.03214*. [Online]. Available: <https://arxiv.org/abs/2002.03214>
- [37] W. J. Choi, K. W. Cheong, and J. M. Cioffi, "Iterative soft interference cancellation for multiple antenna systems," in *Proc. WCNC*, Chicago, IL, USA, Sep. 2000, pp. 304–309.
- [38] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error rate," *IEEE Trans. Inf. Theory*, vol. IT-20, no. 2, pp. 284–287, Mar. 1974.
- [39] N. Shlezinger, N. Farsad, Y. C. Eldar, and A. J. Goldsmith, "Data-driven factor graphs for deep symbol detection," 2020, *arXiv:2002.00758*. [Online]. Available: <https://arxiv.org/abs/2002.00758>
- [40] X. Lingyun and D. Limin, "Efficient Viterbi beam search algorithm using dynamic pruning," in *Proc. ICSP*, Beijing, China, Aug. 2004, pp. 699–702.
- [41] W. Gerstacker and R. Schober, "Equalization concepts for EDGE," *IEEE Trans. Wireless Commun.*, vol. 1, no. 1, pp. 190–199, Aug. 2002.
- [42] B. Hassibi and B. Hochwald, "How much training is needed in multiple-antenna wireless links?" *IEEE Trans. Inf. Theory*, vol. 49, no. 4, pp. 951–963, Apr. 2003.
- [43] N. Shlezinger and R. Dabora, "Frequency-shift filtering for OFDM signal recovery in narrowband power line communications," *IEEE Trans. Commun.*, vol. 62, no. 4, pp. 1283–1295, Apr. 2014.
- [44] A. J. Goldsmith, *Wireless Communications*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [45] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980v9*. [Online]. Available: <https://arxiv.org/abs/1412.6980v9>
- [46] J. P. Nolan, *Stable Distributions: Models for Heavy-Tailed Data*. Basel, Switzerland: Birkhäuser, 2003.

- [47] N. Farsad, Y. Murin, W. Guo, C.-B. Chae, A. W. Eckford, and A. Goldsmith, "Communication system design and analysis for asynchronous molecular timing channels," *IEEE Trans. Mol. Biol. Multi-Scale Commun.*, vol. 3, no. 4, pp. 239–253, Dec. 2017.
- [48] K. Pelekanakis and M. Chitre, "Robust equalization of mobile underwater acoustic channels," *IEEE J. Ocean. Eng.*, vol. 40, no. 4, pp. 775–784, Oct. 2015.
- [49] Z. Mei, M. Johnston, S. Le Goff, and L. Chen, "Performance analysis of LDPC-coded diversity combining on Rayleigh fading channels with impulsive noise," *IEEE Trans. Commun.*, vol. 65, no. 6, pp. 2345–2356, Jun. 2017.



**Nir Shlezinger** (Member, IEEE) received the B.Sc., M.Sc., and Ph.D. degrees from Ben-Gurion University, Israel, in 2011, 2013, and 2017, respectively, all in electrical and computer engineering. From 2009 to 2013, he worked as a Research and Development Engineer with Yitran Communications. From 2017 to 2019, he was a Post-Doctoral Researcher with the Technion—Israel Institute of Technology. He is currently a Post-Doctoral Researcher with the Signal Acquisition, Modeling, Processing and Learning Lab, Weizmann Institute of Science. His

research interests include communications, information theory, signal processing, and machine learning.



**Nariman Farsad** (Member, IEEE) is currently a Research Scholar with the Department of Electrical Engineering, Stanford University, Stanford, CA, USA. His research interests are on improving the performance of communication systems through machine learning and artificial intelligence and on emerging communication technologies such as molecular communications. He received the Natural Sciences and Engineering Research Council of Canada Postdoctoral Fellowship, the Second Prize from the 2014 IEEE ComSoc Student Competition, and

the Best Demo Award at INFOCOM2015. He was recognized as the Finalist for the 2014 Bell Labs Prize. He was a member of the technical program committees of ICC2015-2019, BICT2015, and GLOBECOM2015-2019. He was the Co-Chair of the Machine Learning Data Competition at the 2020 IEEE Communication Theory Workshop (CTW). He was an Area Associate Editor of the IEEE JOURNAL ON SELECTED AREAS OF COMMUNICATION Special Issue on Emerging Technologies in Communications and a technical reviewer for a number of journals including the IEEE TRANSACTIONS ON SIGNAL PROCESSING and the IEEE TRANSACTIONS ON COMMUNICATIONS.



**Yonina C. Eldar** (Fellow, IEEE) received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering from Tel-Aviv University (TAU), Tel-Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2002.

She was a Professor with the Department of Electrical Engineering, Technion—Israel Institute of Technology, where she was the Edwards Chair in engineering. She was a Visiting Professor with

Stanford University. She is currently a Professor with the Department of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel. She is also a Visiting Professor with MIT, also a Visiting Scientist with the Broad Institute, and also an Adjunct Professor with Duke University. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, learning and optimization methods, and their applications to biology and optics.

Dr. Eldar is a member of the Israel Academy of Sciences and Humanities (elected 2017) and a EURASIP fellow. She was a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. She is a member of the IEEE Sensor Array and Multichannel Technical Committee and serves on several other IEEE committees. She was a Horev Fellow of the Leaders in Science and Technology program at the Technion—Israel Institute of Technology and an Alon Fellow. She has received many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award in 2013, the IEEE/AESS Fred Nathanson Memorial Radar Award in 2014, and the IEEE Kiyu Tomiyasu Award in 2016. She received the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion—Israel Institute of Technology, the Muriel and David Jacknow Award for Excellence in Teaching, and the Technion's Award for Excellence in Teaching (two times). She received several best paper awards and best demo awards together with her research students and colleagues including the SIAM Outstanding Paper Prize, the UFFC Outstanding Paper Award, the Signal Processing Society Best Paper Award, and the IET Circuits, Devices and Systems Premium Award. She was selected as one of the 50 most influential women in Israel. She was a co-chair and a technical co-chair of several international conferences and workshops. She served as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*. She is the Editor-in-Chief of *Foundations and Trends in Signal Processing*. In the past, she was a Signal Processing Society Distinguished Lecturer, a member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing Technical Committees.



**Andrea J. Goldsmith** (Fellow, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical engineering from University of California at Berkeley. She founded and served as a Chief Technical Officer for Plume WiFi (formerly Accelera, Inc.) and for Quantenna (QTNA), Inc. She currently serves on the Board of Directors for Medtronic (MDT) and Crown Castle, Inc. (CCI). She is currently the Stephen Harris Professor with the School of Engineering and a Professor of electrical engineering with Stanford University. She has authored the book *Wireless*

*Communications* and coauthored the books *MIMO Wireless Communications* and *Principles of Cognitive Radio* (Cambridge University Press), and an inventor on 29 patents. Her research interests are in information theory, communication theory, and signal processing, and their application to wireless communications, interconnected systems, and neuroscience.

Dr. Goldsmith is a member of the National Academy of Engineering and the American Academy of Arts and Sciences, and a fellow of Stanford. She has received several awards for her work, including the IEEE Sumner Technical Field Award, the ACM Athena Lecturer Award, the ComSoc Armstrong Technical Achievement Award, the WICE Mentoring Award, and the Silicon Valley/San Jose Business Journal's Women of Influence Award. She is currently the Founding Chair of the IEEE Board of Directors Committee on Diversity, Inclusion, and Ethics. She served as the President for the IEEE Information Theory Society in 2009, as the Founding Chair for its student committee, and as the Founding Editor-in-Chief for the IEEE JOURNAL ON SELECTED AREAS OF INFORMATION THEORY. She has also served on the Board of Governors for both the IEEE Information Theory and Communications Societies. At Stanford, she has served as a Chair for Stanford's Faculty Senate and for multiple terms as a Senator, and on its Academic Council Advisory Board, Budget Group, Committee on Research, Planning and Policy Board, Commissions on Graduate and on Undergraduate Education, Faculty Women's Forum Steering Committee, and Task Force on Women and Leadership.