

Data-driven symbol detection via model-based machine learning*

NARIMAN FARSAAD, NIR SHLEZINGER,
ANDREA J. GOLDSMITH, AND YONINA C. ELДАР

The design of symbol detectors in digital communication systems has traditionally relied on statistical channel models that describe the relation between the transmitted symbols and the observed signal at the receiver. Here we review a data-driven framework to symbol detection design which combines machine learning (ML) and model-based algorithms. In this hybrid approach, well-known channel-model-based algorithms such as the Viterbi method, BCJR detection, and multiple-input multiple-output (MIMO) soft interference cancellation (SIC) are augmented with ML-based algorithms to remove their channel-model-dependence, allowing the receiver to learn to implement these algorithms solely from data. The resulting data-driven receivers are most suitable for systems where the underlying channel models are poorly understood, highly complex, or do not well-capture the underlying physics. Our approach is unique in that it only replaces the channel-model-based computations with dedicated neural networks that can be trained from a small amount of data, while keeping the general algorithm intact. Our results demonstrate that these techniques can yield near-optimal performance of model-based algorithms without knowing the exact channel input-output statistical relationship and in the presence of channel state information uncertainty.

1. 85th birthday tribute to Tom Kailath

It is our great pleasure to contribute this work in honor of Tom Kailath's 85th birthday. Tom's contributions are legendary, profoundly impacting theory and practice across a breathtaking range of disciplines. In particular, he laid the mathematical foundations for solving challenging problems across

*This work was supported in part by the US - Israel Binational Science Foundation under grant No. 2026094, by the Israel Science Foundation under grant No. 0100101, by the Office of the Naval Research under grant No. 18-1-2191, and by NSERC Discovery Grant and CFI John Evans Leaders Funds.

communications, control, integrated circuits, and signal processing. Moreover, he illuminated threads and connections across these diverse fields, leading to much better fundamental understanding and insights within them. Not satisfied with just developing theory, Tom also created technologies and successful companies on the basis of his mathematical foundations, demonstrating the importance of linking theory and practice while having transformational impact in both. Tom has also been a spectacular mentor and role model for generations of students and colleagues. This work is contributed in the spirit of Tom, coupling mathematical foundations within communication system design with the new tool of machine learning to improve overall performance. Thanks Tom for all you have contributed to our profession, and we wish you a very happy 85th birthday.

2. Introduction

In digital communication systems, the receiver is required 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 [1, 2]. Consequently, in some cases, conventional methods cannot be applied, particularly when the channel model is highly complex, poorly understood, or does not well-capture the underlying physics of the system. Furthermore, when the channel models are known, many detection algorithms rely on channel state information (CSI), i.e., the instantaneous parameters of the channel model, for detection. Therefore, conventional channel-model-based techniques require the instantaneous CSI to be estimated. This process entails overhead, which decreases the data rate. Moreover, inaccurate CSI estimation typically degrades the detection performance.

An alternative data-driven approach to model-based algorithms uses 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]. ML schemes have several benefits over traditional model-based approaches: First, ML methods are independent of the underlying stochastic model, and 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 meaningful semantic information from

the observed data [4], a task which is very difficult to carry out using model-based methods. Finally, ML techniques often lead to faster inference 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; see detailed surveys in [7, 8, 9, 10, 11]. Unlike model-based receivers, which implement a specified detection rule, ML-based receivers learn how to map the channel outputs into the transmitted symbols from training, namely, they operate in a data-driven manner. Broadly speaking, previously proposed ML-based receivers can be divided into two main categories: Conventional DNNs and unfolded networks [6]. The first group replaces the receiver processing with a DNN architecture from established methods in the ML literature. A-priori knowledge of the channel model is accounted for in the selection of the network type, which is typically treated as a black box. For example, recurrent neural networks (RNNs) were applied for decoding sequential codes in [12]; the work [13] used sliding bi-directional RNNs for intersymbol interference (ISI) channels; and the work [14] used variational autoencoders for unsupervised equalization. Such DNNs, which use conventional network architectures that are ignorant of the underlying channel model, can typically operate reliably in various scenarios with or without CSI and channel model knowledge, assuming that they were properly trained for the specific setup. Nonetheless, block box DNNs tend to have a large number of parameters, and thus require large data sets to train [11], limiting their application in dynamic environments, which are commonly encountered in communications.

Unlike conventional DNNs, which utilize established architectures, in unfolded receivers the network structure is designed following a model-based algorithm. In particular, deep unfolding is a method for converting an iterative algorithm into a DNN by designing each layer of the network to resemble a single iteration [15, 6, 16]. The resulting DNN tends to demonstrate improved convergence speed and robustness compared to the model-based algorithm [5, 6]. In the context of symbol detection, the works [17, 18] designed deep receivers by unfolding the projected gradient descent algorithm for recovering the MAP solution, while [19] unfolded the approximate message passing optimization scheme, and [20] proposed to recover signals obtained from one-bit quantized measurements by unfolding gradient descent optimization. Compared to conventional DNNs, unfolded networks are typically interpretable, and tend to have a smaller number of parameters, and can thus be trained quicker [11]. However, these previously proposed re-

ceivers all assume linear channel models with Gaussian noise, in which CSI is available [17, 18, 19, 20]. Consequently, these methods do not capture the potential of ML in being independent of the model, and thus are applicable only under specific setups.

In this work, we design and study ML-based symbol detection, focusing and presenting some of our recent work in this area in a tutorial fashion [21, 22, 23]. The main distinction between our approach and prior work is using ML in conjunction with well-known detection algorithms in a hybrid fashion. Specifically, our design is based on integrating ML methods into established detection algorithms, allowing them to be learned from training and operate in a data-driven fashion. Our approach to symbol detection implements well-known channel-model-based algorithms such as the Viterbi algorithm [24], the BCJR method [25], and multiple-input multiple-output (MIMO) soft interference cancellation (SIC) [26], while only removing their channel-model-dependence by replacing the CSI-based computations with dedicated DNNs, i.e., we *integrate ML into these channel-model-based algorithms to learn CSI*.

We begin by presenting in detail our approach for combining ML and model-based algorithms, and its relationship with deep unfolding. Then we present two channel models to which we apply our approach for symbol detection: First, we focus on finite-memory channels and present ViterbiNet [21] and BCJRNet [22], which are data-driven implementations of the maximum-likelihood sequence detector and the MAP symbol detector for this channel, respectively. We then focus on the multiuser MIMO channel and present DeepSIC [23], which learns to implement the iterative SIC algorithm under this channel model. Our resulting receivers are shown to approach the performance of their model-based counterparts without requiring CSI and using relatively small training sets. We also observe that, in the presence of CSI uncertainty, the performance of the model-based algorithms is significantly degraded, while the combined ML and model-based approach is capable of reliably detecting the symbols. Our results demonstrate the potential gains of combining ML and model-based algorithms over both pure model-based and pure ML-based methods.

The rest of this paper is organized as follows: In Section 3 we describe the general hybrid ML approach. Section 4 focuses on finite-memory channels and presents the ViterbiNet and the BCJRNet for these channels. Section 5 details the data-driven SIC algorithm, and Section 6 provides concluding remarks.




Figure 1: Channel-model-based symbol detection illustration.

3. Model-based machine learning

In this section we present our proposed methodology for converting model-based algorithms to data-driven techniques by integrating ML into them. We focus here on the high-level rationale, while concrete examples for symbol detection methods are presented in Sections 4–5. In particular, we first motivate the fusion of ML and channel-model-based symbol detection schemes in Section 3.1. Then, in Section 3.2 we present how ML methods can be combined with model-based algorithms.

3.1. Machine learning for symbol detection

Symbol detection is arguably the most basic problem in digital communications. Symbol detection refers to how the receiver recovers a set of transmitted symbols denoted by \mathbf{S} , which belong to some discrete constellation, from its observed channel output denoted by \mathbf{Y} . This problem has been widely studied, and various algorithms with optimality guarantees and controllable complexity measures have been proposed over decades of research, depending on the statistical model relating \mathbf{Y} and \mathbf{S} , i.e., their conditional distribution $p_{\mathbf{Y}|\mathbf{S}}$ [1, Ch. 5]. The vast majority of these previously proposed schemes are *channel-model-based*, namely, they require accurate prior knowledge of $p_{\mathbf{Y}|\mathbf{S}}$, and typically assume it takes some simplified form, e.g., that of a linear channel with additive white Gaussian noise (AWGN). An illustration of conventional channel-model-based symbol detection is depicted in Fig. 1. As discussed in the introduction, these requirements limit the application of conventional model-based schemes in complex environments, and may entail substantial overhead due to the need to constantly estimate CSI at the receiver.

Recent years have witnessed a dramatically growing interest in ML methods, and particularly, in DNNs. These data-driven trainable structures have demonstrated unprecedented empirical success in various applications, including computer vision and speech processing [3]. ML-driven techniques

have several key benefits 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 [27, Ch. 1]. Second, when the underlying model is extremely complex, ML algorithms have demonstrated the ability to extract and disentangle the meaningful semantic information from the observed data [4], a task which is very difficult to carry out using traditional model-based approaches, even when the model is perfectly known. Specifically, ML tools have empirically demonstrated the ability to disentangle the relevant semantic information under complex analytically intractable statistical relationships. Furthermore, DNN techniques tend to generalize well, despite being over parameterized, and remain resilient under minor non-adversarial parameter perturbations after training [28]. Finally, ML algorithms often lead to faster inference compared to iterative model-based approaches, even when the model is known [5].

Nonetheless, not every problem can and should be solved using DNNs. While networks with a large set of parameters can realize a broad set of functions, making them suitable for a wide range of tasks, a large data set reflecting the task under which the system must operate is required to train these networks. For example, symbol detection over wireless communication channels is typically dynamic in nature. This implies that DNN-based symbol detectors must be either trained anew for each channel realization, a task which may be infeasible when large labeled sets are required for training, or alternatively, that the data-driven receiver should be trained in advance using samples corresponding to a large set of expected channel conditions. In the latter, typically the resulting deep symbol detector is likely to achieve reasonable but non-optimal performance when inferring using samples corresponding to each of channels condition.

Our goal is to allow model-based symbol detection schemes to be applied in scenarios for which, due to either a complex statistical model or lack of its knowledge, these methods cannot be applied directly. This is achieved by combining ML and model-aware algorithms via model-based ML, as discussed in the sequel.

3.2. Combining ML and model-based algorithms

Here, we present our rationale for combining ML, and in particular DNNs, with channel-model-based symbol detection algorithms. Consider a symbol detection algorithm which produces an output vector $\hat{\mathbf{S}}$ based on an observed input vector \mathbf{Y} . Particularly, we focus on channels in which, given prior




Figure 2: Generic iterative algorithm illustration.

knowledge of $p_{\mathbf{Y}|\mathbf{S}}$, $\hat{\mathbf{S}}$ can be recovered from \mathbf{Y} with provable performance and controllable complexity in an iterative fashion, e.g., via the Viterbi detector [24] or the BCJR algorithm [25] for finite memory channels, or via interference cancellation methods for MIMO channels [26]. This generic family of iterative algorithms consists of some input and output processing stages, with an intermediate iterative procedure. The latter can in turn be divided into a channel-model-based computation, namely, a procedure that is determined by $p_{\mathbf{Y}|\mathbf{S}}$; and a set of generic mathematical manipulations. An illustration of this generic algorithmic procedure is depicted in Fig. 2.

Here, we want to implement the aforementioned algorithm in a data-driven manner. Possible motivations for doing so are that $p_{\mathbf{S},\mathbf{Y}}$ is unknown, costly to estimate accurately, or too complex to express analytically. A powerful tool for learning to implement tasks from data samples are DNNs, with their dramatic empirical success in applications such as computer vision, where the underlying model is very complex or poorly understood. Modern ML suggests two main approaches for obtaining $\hat{\mathbf{S}}$ from \mathbf{Y} , as done in the above model-based algorithm, using DNNs:

End-to-end networks Arguably the most common framework is to replace the algorithm with some standard DNN architecture. A-priori knowledge of the underlying model is accounted for in the selection of the network type, which is typically treated as a black box. By providing the network with sufficient samples from the joint distribution of the input and label, it should be able to tune its parameters to minimize the loss function. If the network is sufficiently parameterized, i.e., deep, its feasible loss-minimizing structure is within close proximity of the optimal mapping from \mathbf{Y} to $\hat{\mathbf{S}}$, which means that a properly trained network can approach the performance of the model-based algorithm. The main drawback of using end-to-end networks is that learning a large number of parameters requires a large data set to train. Even when a sufficiently large data set is available, the resulting long training period limits the application of this approach in dynamic environments, where the system has to be retrained often.




Figure 3: Deep unfolding illustration.




Figure 4: Integrating ML into a model-based algorithm illustration.

Deep unfolding A possible method to design a DNN based on an iterative model-based algorithm is to replace each iteration with a layer in the network, where the design of each layer is inspired by the iterative procedure [15, 6], as illustrated in Fig. 3. The main benefit of deep unfolding over using end-to-end networks stems from the reduced complexity, as unfolded networks typically have less parameters compared to conventional end-to-end networks [11]. Furthermore, even when the model-based algorithm is feasible, processing \mathbf{Y} through a trained unfolded DNN is typically faster than applying the iterative algorithm [5].

Here, we propose a new approach for realizing model-based algorithms of the aforementioned iterative structure in a data-driven fashion. In particular, we first identify the specific computations which depend on the underlying statistical model $p_{\mathbf{S}, \mathbf{Y}}$. Then we replace these computations with ML-based techniques such as dedicated DNNs that can be trained from data. In various algorithms, as shown in the following sections, this specific computation often requires accurate knowledge of $p_{\mathbf{S}, \mathbf{Y}}$ to, e.g., compute a conditional probability measure, or to estimate some parameter, both tasks which can be accurately learned using relatively simple neural networks. An illustration of this strategy is depicted in Fig. 4.

Despite their similarity in combining DNNs with model-based algorithms, there is a fundamental difference between the proposed approach and 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*. We propose to implement the algorithm, while only removing its model-dependence by replacing the model-based computations with dedicated DNNs, *thus integrating ML into the algorithm*.

4. Example 1: finite-memory channels

In this section, we present how the rationale for combining ML and model-based algorithms detailed in the previous section can be applied for data-driven symbol detection over finite-memory causal channels. We first present the channel model in Section 4.1. Then, we propose the integration of DNNs into two established detection methods for such channels: We begin with the Viterbi algorithm [29], which implements the maximum likelihood sequence detector, proposing a receiver architecture referred to as *ViterbiNet* in Section 4.2. ViterbiNet learns to implement the Viterbi algorithm from labeled data. Then, we extend these ideas to propose the *BCJRNet* receiver in Section 4.3. BCJRNet learns the factor graph representing the underlying statistical model of the channel, which it uses to implement the symbol-level MAP detector via the BCJR algorithm [25]. We numerically demonstrate the performance of these ML-based receivers in Section 4.4.

4.1. System model

We consider the recovery of 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 . Since the channel is causal and has a finite memory, $Y[i]$ is given by a stochastic mapping of \mathbf{S}_{i-l+1}^i , where l is the memory of the channel, assumed to be smaller than the blocklength, i.e., $l < t$. The conditional probability density function (PDF) of the channel output given its input thus satisfies

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

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 $p_{Y[i] | \mathbf{S}_{i-l+1}^i}(y | \mathbf{s})$




Figure 5: System model.

does not depend on the index i . An illustration of the system is depicted in Fig. 5.

We focus on two main symbol detector algorithms designed for the finite state channel model in (1): the Viterbi detector [24], and the BCJR detector [25]. The Viterbi scheme implements maximum likelihood sequence detection given by

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

The BCJR receiver implements the MAP symbol detector given by

$$(3) \quad \hat{S}[k] = \arg \max_{s \in \mathcal{S}} p_{S[k] | \mathbf{Y}^t}(s | \mathbf{y}^t) = \arg \max_{s \in \mathcal{S}} p_{S[k], \mathbf{Y}^t}(s, \mathbf{y}^t), \quad k \in \mathcal{T}.$$

Over the next two sections we present ViterbiNet and BCJRNet, which are two algorithms that learn to carry out Viterbi and BCJR detection, respectively, in a data-driven manner, using ML tools.

4.2. ViterbiNet

First we present a data-driven framework for recovering \mathbf{S}^t from the channel output \mathbf{Y}^t based on maximum-likelihood sequence detection. 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 $p_{Y[i] | \mathcal{S}_{i-l+1}^i}(\mathbf{y} | \mathbf{s})$, i.e., the receiver does not have this CSI. As a preliminary step to designing ViterbiNet, we review Viterbi detection.

4.2.1. The Viterbi Algorithm The following description of the Viterbi algorithm is based on [2, 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 in (2). By defining

$$(4) \quad c_i(\mathbf{s}) \triangleq -\log p_{Y[i] | \mathcal{S}_{i-l+1}^i}(\mathbf{y}[i] | \mathbf{s}), \quad \mathbf{s} \in \mathcal{S}^l,$$

it follows from (1) that the log-likelihood in (2) can be written as $\log p_{\mathbf{Y}^t|\mathcal{S}^t}(\mathbf{y}^t|\mathbf{s}^t) = \sum_{i=1}^t c_i(\mathbf{s}_{i-l+1}^i)$, and the optimization problem (2) becomes

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

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* for each state. The resulting scheme, known as the Viterbi algorithm, is given below as Algorithm 1.

Algorithm 1 The Viterbi Algorithm [24]

- 1: Input: Block of channel outputs \mathbf{y}^t , where $t > l$.
 - 2: Initialization: Set $k=1$, and fix an initial path cost $\tilde{c}_0(\tilde{\mathbf{s}})=0$, for each $\tilde{\mathbf{s}} \in \mathcal{S}^l$.
 - 3: For each state $\tilde{\mathbf{s}} \in \mathcal{S}^l$, compute $\tilde{c}_k(\tilde{\mathbf{s}}) = \min_{\mathbf{u} \in \mathcal{S}^l: \mathbf{u}_2^l = \tilde{\mathbf{s}}^{t-1}} (\tilde{c}_{k-1}(\mathbf{u}) + c_k(\tilde{\mathbf{s}}))$.
 - 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 Viterbi algorithm has two major advantages: 1) 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 ; 2) The algorithm produces estimates sequentially during run-time. 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 compute $c_i(\mathbf{s})$ of (4) 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 we propose ViterbiNet, an ML-based detector based on the Viterbi algorithm, that does not require CSI.

4.2.2. 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 $c_i(\mathbf{s})$.




Figure 6: ViterbiNet Architecture.

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 full CSI (e.g., by using an upper bound on memory length).

Since the channel is stationary, it holds by (4) 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, to implement Algorithm 1 in a data-driven fashion, we replace the explicit computation of the log-likelihood (4) 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, and the detector implements Algorithm 1 using the learned $c_i(\mathbf{s})$. The proposed architecture is illustrated in Fig. 6.

A major challenge in implementing a network capable of computing $c_i(\mathbf{s})$ from $y[i]$ stems from the fact that, by (4), $c_i(\mathbf{s})$ represents the log-likelihood of $Y[i] = y[i]$ given $\mathbf{S}_{i-l+1}^i = \mathbf{s}$. However, DNNs trained with the cross entropy loss typically output the conditional distribution of $\mathbf{S}_{i-l+1}^i = \mathbf{s}$ given $Y[i] = y[i]$, i.e., $p_{\mathbf{S}_{i-l+1}^i | Y[i]}(\mathbf{s} | y[i])$. 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], output the distribution of the label \mathbf{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 \mathbf{S}_{i-l+1}^i . For this reason, the previous work [13] used a DNN to approximate the MAP detector by considering $p_{S[i] | \mathbf{Y}_{i-l+1}^i}(s | \mathbf{y}_{i-l+1}^i)$. The quantity needed for the Viterbi algorithm is the conditional PDF $p_{Y[i] | \mathbf{S}_{i-l+1}^i}(y[i] | \mathbf{s})$,




Figure 7: ML-based likelihood computation.

and not the conditional distribution $p_{\mathbf{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$. Note that while $\sum_{\mathbf{s} \in \mathcal{S}^l} p_{\mathbf{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i]) = 1$, for the desired conditional PDF $\sum_{\mathbf{s} \in \mathcal{S}^l} p_{Y[i]|\mathbf{S}_{i-l+1}^i}(y[i]|\mathbf{s}) \neq 1$ in general. Therefore, outputs generated using conventional DNNs with a softmax output layer are not applicable. The fact that Algorithm 1 uses $p_{Y[i]|\mathbf{S}_{i-l+1}^i}(y[i]|\mathbf{s})$ instead of $p_{\mathbf{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$ allows it to exploit the Markovian nature of the channel, induced by the finite memory in (1), resulting in the aforementioned advantages of the Viterbi algorithm.

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

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

Therefore, given estimates of $p_{Y[i]}(y[i])$ and of $p_{\mathbf{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$ for each $\mathbf{s} \in \mathcal{S}^l$, the log-likelihood function $c_i(\mathbf{s})$ can be recovered using (6) and (4).

A parametric estimate of $p_{\mathbf{S}_{i-l+1}^i|Y[i]}(\mathbf{s}|y[i])$, 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. Furthermore, the fact that $Y[i]$ is a stochastic mapping of \mathbf{S}_{i-l+1}^i implies that its distribution can be approximated as a mixture model of m^l kernel functions [31]. Consequently, a parametric estimate of $p_{Y[i]}(y[i])$, denoted $\hat{p}_\varphi(y[i])$, can be obtained from the training data using mixture density networks [32], expectation minimization (EM)-based algorithms [31, Ch. 2], or any other finite mixture model fitting method. The resulting ML-based log-likelihood computation is illustrated in Fig. 7.

4.3. BCJRNet

In this section, we expand the approach presented in the previous section to factor graphs and detail how this strategy can be used to implement BCJRNet, which is a data-driven channel-model-independent MAP symbol detector for finite-memory channels. We first discuss how ML can be incorporated into factor graphs, and then present the BCJRNet detector.

4.3.1. Learning factor graphs Factor graphs are graphical models that are used to represent factorization of multivariate functions (e.g. joint probability distributions) [33]. In factor graphs, each factor is represented by a *factor or function node*, and each variable is represented by an edge or half-edge in the graph.¹ Factor graph methods, such as the sum-product algorithm [33], exploit the factorization of a joint distribution to efficiently compute a desired quantity by passing messages along the edges of the graph. For example, the application of the sum-product algorithm in the finite-memory channel in (1) exploits its factorization to compute marginal distributions, an operation whose burden typically grows exponentially with the block size, with complexity that only grows linearly with t . In fact, the sum-product algorithm is exactly the recursive computation carried out in the BCJR symbol detector [25]. More generally, the sum-product algorithm specializes a multitude of common signal processing techniques, including the Kalman filter and hidden Markov model (HMM) prediction [35].

The factor graph representing finite-memory channels (1) is presented in Fig. 8, where we use the simplifying notation $\mathbf{s}_i = \mathbf{s}_{i-l+1}^i$, $y_i = y[i]$, and the fact that

$$P_{\mathbf{Y}^t, \mathbf{S}^t}(\mathbf{y}^t, \mathbf{s}^t) = \prod_{i=1}^t P_{Y[i]|\mathbf{S}_{i-l+1}^i}(y_i|\mathbf{s}_i) P_{\mathbf{S}_{i-l+1}^i|\mathbf{S}_{i-l}^{i-1}}(\mathbf{s}_i|\mathbf{s}_{i-1}) = \prod_{i=1}^t f_i(y_i, \mathbf{s}_i, \mathbf{s}_{i-1}).$$

In order to implement the sum-product scheme, one must be able to specify the factor graph encapsulating the underlying distribution, and in particular, the function nodes $\{f_j\}_{j=1}^t$. This implies that the BCJR detector requires full CSI.

We next generalize the approach we took in the previous section with the ViterbiNet to realize a broad family of data-driven factor graph methods by learning the mappings carried out at the function nodes from a relatively

¹In this work we use the style proposed in [34], known as *Forney-style factor graphs*.




Figure 8: Factor graph of finite-memory channel.

small set of labeled data using ML tools. By doing so, one can train a system to learn an underlying factor graph, which can then be utilized for inference using conventional factor graph methods, such as the sum-product algorithm.

The proposed approach requires prior knowledge of the graph structure, but not its nodes. For example, a finite-memory channel with memory length not larger than l can be represented using the structure in Fig. 8 while its specific input-output relationship dictate the function nodes. Consequently, in order to learn such a factor graph from samples, one must learn its function nodes. As these mappings often represent conditional distribution measures, they can be naturally learned using classification networks, e.g., fully-connected DNNs with a softmax output layer and cross-entropy objective, which are known to reliably learn conditional distributions in complex environments [4]. Furthermore, in many scenarios of interest, e.g., stationary finite-memory channels and time-invariant state-space models, the mapping implemented by the factor nodes $f_i(\cdot)$ does not depend on the index i . In such cases, only a fixed set of mappings whose size does not grow with the dimensionality t has to be tuned in order to learn the complete factor graph. An example of how this concept of learned factor graphs can be applied to the BCJR detector is presented next.

4.3.2. BCJRNet: data-driven MAP recovery In this section, we first present the application of the sum-product method for MAP symbol detection in finite-memory channels, also known as the BCJR algorithm [25]. We then describe how the approach proposed in Section 3 can be incorporated into the BCJR algorithm.

Consider a stationary finite-memory channel, namely, a channel obeying the model in Section 4.1. The detector which minimizes the symbol error probability is the MAP rule given in (3). Using the formulation of the state vectors $\{\mathbf{S}_{k-l+1}^k\}$, the desired joint probability can be written as [36, Ch. 9.3]

$$(7) \quad p_{S[k], \mathbf{Y}^t}(s_k, \mathbf{y}^t) = \sum_{\mathbf{s} \in \mathcal{S}^L} p_{\mathbf{S}_{k-l}^{k-1}, \mathbf{S}_{k-l+1}^k, \mathbf{Y}^t}(\mathbf{s}, [s_k, (\mathbf{s})_1 \dots (\mathbf{s})_{l-1}]^T, \mathbf{y}^t).$$

The summands in (7) are the joint distributions evaluated recursively from the channel factor graph. Thus, when the factor graph is known, the MAP rule (3) can be computed efficiently using the sum-product algorithm.

Now, we show how the rationale presented in Section 3 yields a method for learning the factor graphs of finite-memory channels, from which the MAP detector can be obtained. We assume that the channel memory length, l , is known. However, the channel model, i.e., the conditional distribution $p_{Y[i]|\mathbf{S}_{i-l+1}^i}(\cdot)$ is unknown, and only a set of labeled input-output pairs is available.

Since the channel memory is known, the structure of the factor graph is fixed to that depicted in Fig. 8. Consequently, in order to learn the factor graph, one must only adapt the function nodes $f_i(\cdot)$, which are given by

$$(8) \quad \begin{aligned} f_i(y_i, \mathbf{s}_i, \mathbf{s}_{i-1}) &= P_{Y[i]|\mathbf{S}_{i-l+1}^i}(y_i|\mathbf{s}_i) P_{\mathbf{S}_{i-l+1}^i|\mathbf{S}_{i-l}^{i-1}}(\mathbf{s}_i|\mathbf{s}_{i-1}) \\ &\stackrel{(a)}{=} \begin{cases} \frac{1}{m} p_{Y|\mathbf{S}}(y_i|\mathbf{s}_i) & (\mathbf{s}_i)_j = (\mathbf{s}_{i-1})_{j-1}, \quad \forall j \in \{2, \dots, l\} \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

where (a) follows from the definition of the state vectors $\{\mathbf{S}_{k-l+1}^k\}$ and the channel stationarity, which implies that $p_{Y[i]|\mathbf{S}_{i-l+1}^i}(\cdot)$ does not depend on i and can be written as $p_{Y|\mathbf{S}}(\cdot)$. The formulation of the function nodes in (8) implies that they can be estimated by training an ML-based system to evaluate $p_{Y|\mathbf{S}}(\cdot)$ from which the corresponding function node value is obtained via (8). Once the factor graph representing the channel is learned, symbol recovery is carried out using the sum-product method detailed earlier. The resulting receiver, referred to as *BCJRNet*, thus implements BCJR detection in a data-driven manner, and is expected to approach MAP-performance when the function nodes are accurately estimated, as demonstrated in our numerical study in the next section.

Taking the same approach as ViterbiNet, since $Y[i] = y_i$ is given and may take continuous values while \mathbf{S}_{i-l+1}^i , representing the label, takes discrete values, a natural approach to evaluate $p_{Y[i]|\mathbf{S}_{i-l+1}^i}(y_i|\mathbf{s}_i)$ for each $\mathbf{s}_i \in \mathcal{S}^l$ using ML tools is to estimate $p_{\mathbf{S}_{i-l+1}^i|Y[i]}(\mathbf{s}_i|y_i)$, from which the desired $p_{Y[i]|\mathbf{S}_{i-l+1}^i}(y_i|\mathbf{s}_i)$ can be obtained using Bayes rule via (6). In particular, BCJRNet utilizes two parametric models: one for evaluating the conditional $p_{\mathbf{S}_{i-l+1}^i|Y[i]}(\mathbf{s}_i|y_i)$, and another for computing the marginal PDF $p_{Y[i]}(y_i)$, using a similar system as that used for computing the likelihood functions in ViterbiNet, i.e., via the architecture illustrated in Fig. 7. The estimates are then combined into the learned function nodes using (8) and (6), which are used by BCJRNet to carry out BCJR detection in a data-driven manner.

4.4. Numerical evaluations

In this section, we numerically compare the performance of the proposed data-driven ViterbiNet and BCJRNet to the conventional model-based Viterbi algorithm and BCJR detector as well as to previously proposed deep symbol detectors. Throughout this numerical study we implement the fully-connected network in Fig. 7 using three layers: a 1×100 layer followed by a 100×50 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 [31, Ch. 2]. The network is trained using 5000 training samples to minimize the cross-entropy loss via the Adam optimizer [37] 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. 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.

We consider two finite-memory channels: An ISI channel with AWGN, and a Poisson channel. In both channels we set the memory length to $l = 4$. For the AWGN channel, 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 obeying an exponentially decaying profile $(\mathbf{h})_\tau \triangleq e^{-\gamma(\tau-1)}$ for $\gamma > 0$. The input-output relationship is given by

$$(9) \quad Y[i] = \sqrt{\rho} \cdot \sum_{\tau=1}^l (\mathbf{h}(\gamma))_\tau S[i - \tau + 1] + W[i],$$

where $\rho > 0$ represents the signal-to-noise ratio (SNR). The channel input is randomized from a binary phase shift keying (BPSK) constellation, i.e., $\mathcal{S} = \{-1, 1\}$. For the Poisson channel, the channel input represents on-off keying, namely, $\mathcal{S} = \{0, 1\}$, and the channel output $Y[i]$ is generated from the input via

$$(10) \quad Y[i] | \mathbf{S}^t \sim \mathbb{P} \left(\sqrt{\rho} \cdot \sum_{\tau=1}^l (\mathbf{h}(\gamma))_\tau S[i - \tau + 1] + 1 \right),$$

where $\mathbb{P}(\lambda)$ is the Poisson distribution with parameter $\lambda > 0$, and $X \sim f(X)$ indicates that the random variable X is distributed according to $f(X)$.

For each channel, we numerically compute the symbol error rate (SER) of ViterbiNet and BCJRNet for different values of the SNR parameter ρ . In the following study, the DNN in Fig. 7, which produces a parametric estimate of the log-likelihoods in ViterbiNet and produces the learned function nodes in BCJRNet, is trained anew for each value of ρ . For each SNR ρ , the SER values are averaged over 20 different channel vectors $\mathbf{h}(\gamma)$, obtained by letting γ vary in the range $[0.1, 2]$. For comparison, we numerically compute the SER of the Viterbi and BCJR algorithms, as well as that of the sliding bidirectional RNN (SBRNN) deep symbol decoder proposed in [13]. In order to study the resiliency of the data-driven detectors 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 σ_e^2 . In particular, we use $\sigma_e^2 = 0.1$ for the Gaussian channel (9), and $\sigma_e^2 = 0.08$ for the Poisson channel (10). We consider two cases: *Perfect CSI*, in which the channel-model-based detectors have accurate knowledge of $\mathbf{h}(\gamma)$, while the data-driven receivers are trained using labeled samples generated with the same $\mathbf{h}(\gamma)$ used for generating the test data; and *CSI uncertainty*, where the model-based algorithms are implemented with the log-likelihoods (for Viterbi algorithm) and function nodes (for BCJR detection) computed using the noisy version of $\mathbf{h}(\gamma)$, while the data used for training ViterbiNet and BCJRNet is generated with the noisy version of $\mathbf{h}(\gamma)$ instead of the true one. 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 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. 9, while the corresponding performance versus $\rho \in [10, 30]$ dB for the Poisson channel are depicted in Fig. 10. Observing Figs. 9–10, we note that the performance of the data-driven receivers approaches that of their corresponding CSI-based counterparts. 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 the SER of BCJRNet is within a very small gap of the CSI-based BCJR receiver. For the Poisson channel, both data-driven receivers achieve performance within a small gap of the model-based ones, which is more notable at high SNRs. This gap stems from the model mismatch induced by approximating the distribution of $Y[i]$ as a Gaussian mixture. We also observe that the SBRNN receiver, which was shown in [13] to approach the performance of the CSI-based Viterbi algorithm when sufficient training is

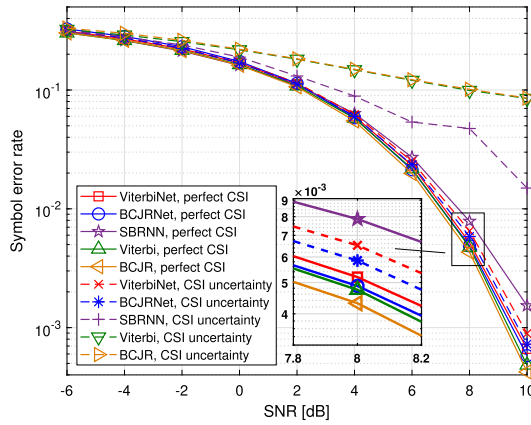


Figure 9: ISI channel with AWGN.

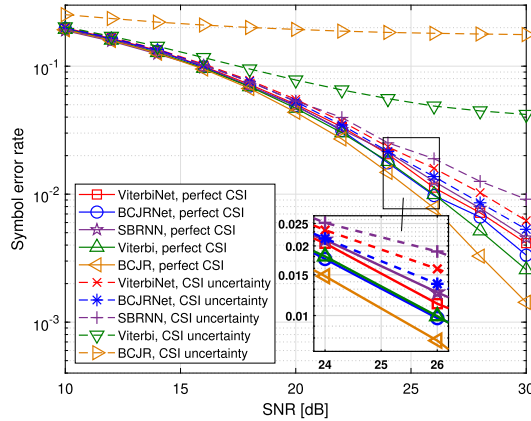


Figure 10: Poisson channel.

provided, is outperformed by ViterbiNet and BCJRNet here due to the small training set size. These results demonstrate that our proposed data-driven detectors, which use simple DNN structures embedded into an established detection algorithms, require significantly less training compared to previously proposed ML-based receivers.

In the presence of CSI uncertainty, it is observed in Figs. 9–10 that both ViterbiNet and BCJRNet significantly outperform the model-based algorithms from which they originate. In particular, when ViterbiNet and BCJRNet are trained with a variety of different channel conditions, they are 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 and BCJR algorithms 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 and BCJR algorithms, as was also observed in [13], it is outperformed by ViterbiNet and BCJRNet with the same level of uncertainty, and the performance gap is more notable in the AWGN channel. The reduced gain of ViterbiNet and BCJRNet over the SBRNN receiver for the Poisson channel stems from the fact that the ML-based module of Fig. 7 uses a Gaussian mixture density estimator for the PDF of $Y[i]$, which obeys a Poisson mixture distribution for the channel (10).

The results presented in this section demonstrate that ML-based receivers can be trained to carry out accurate and robust symbol detection using a relatively small amount of labeled samples (compared to amount of data typically needed to train DNNs). This is achieved by utilizing dedicated ML tools to learn only the model-based computations of symbol detection algorithms, such as the Viterbi and BCJR methods, while combining these learned computations with the generic operations of the algorithms, e.g., the dynamic programming processing of the Viterbi algorithm or the sum-product recursion of the BCJR receiver.

5. Example 2: memoryless MIMO channels

In the previous section we utilized the rationale detailed in Section 3 for combining ML and model-based algorithms in order to implement two symbol detection schemes – the Viterbi detector and the BCJR algorithm – in a data-driven fashion. Both these methods are based on computing the likelihood for each possible combination of channel inputs which could have lead to the observed output, i.e., for each state, followed by some generic recursive operations. Such methods, both the original channel-model-based algorithms as well as their data-driven ML-based implementations, grow computationally infeasible when multiple symbols are transmitted simultaneously, as in, e.g., multiuser MIMO systems, since the possible number of states grows rapidly with the number of transmitted symbols. In this section we demonstrate how the proposed model-based ML rationale can be exploited to realize symbol detection methods suitable for such MIMO systems in a data-driven manner. We begin by presenting the channel model in Section 5.1. Then, in Section 5.2 we show how the iterative SIC symbol detection mapping proposed in [26], which is a MIMO receiver scheme capable of approaching MAP performance at affordable complexity, can be learned from training, resulting in an ML-based receiver referred to as *DeepSIC*. The achievable performance of DeepSIC is demonstrated in a numerical study presented in Section 5.3.




Figure 11: System model.

5.1. System model

We consider a multiuser uplink MIMO system in which K single antenna users communicate with a receiver equipped with n_r antennas over a memoryless stationary channel. At each time instance i , the k th user, $k \in \{1, 2, \dots, K\} \triangleq \mathcal{K}$, transmits a symbol $S_k[i]$ drawn from a constellation \mathcal{S} of size m . Each symbol is uniformly distributed over \mathcal{S} , and the symbols transmitted by different users are mutually independent. We use $\mathbf{Y}[i] \in \mathcal{R}^{n_r}$ to denote the channel output at time index i . While we focus on real-valued channels, the system model can be adapted to complex-valued channels, as complex vectors can be equivalently represented using real vectors of extended dimensions. Since the channel is memoryless, $\mathbf{Y}[i]$ is given by some stochastic mapping of $\mathbf{S}[i] \triangleq [S_1[i], S_2[i], \dots, S_K[i]]^T$, represented by the conditional distribution $p_{\mathbf{Y}[i]|\mathbf{S}[i]}$. Since the channel is stationary, this conditional distribution does not depend on the index i , and is thus denoted henceforth by $p_{\mathbf{Y}|\mathbf{S}}$. An illustration of the system is depicted in Fig. 11.

We focus on the problem of recovering the transmitted symbols $\mathbf{S}[i]$ from the channel output $\mathbf{Y}[i]$. The optimal detection rule which minimizes the probability of error given a channel output realization $\mathbf{Y}[i] = \mathbf{y}$ is the MAP detector. Letting $p_{\mathbf{S}|\mathbf{Y}}$ be the conditional distribution of $\mathbf{S}[i]$ given $\mathbf{Y}[i]$, the MAP rule is given by

$$(11) \quad \hat{\mathbf{s}}_{\text{MAP}}[i] \triangleq \arg \max_{\mathbf{s} \in \mathcal{S}^K} p_{\mathbf{S}|\mathbf{Y}}(\mathbf{s}|\mathbf{y}).$$

The MAP detector jointly recovers the symbols of all users by searching over a set of m^K different possible input combinations, and thus becomes infeasible when the number of users K grows. For example, when binary constellations are used, i.e., $m = 2$, the number of different channel inputs is larger than 10^6 for merely $K = 20$ users. Furthermore, the MAP decoder requires accurate knowledge of the channel model, i.e., the conditional distribution $p_{\mathbf{Y}|\mathbf{S}}$ must be fully known. A common strategy to implement joint

decoding with affordable computational complexity, suitable for channels in which $\mathbf{Y}[i]$ is given by a linear transformation of $\mathbf{S}[i]$ corrupted by additive noise, is interference cancellation [38]. Interference cancellation refers to a family of algorithms which implement joint decoding in an iterative fashion by recovering a subset of $\mathbf{S}[i]$ based on the channel output as well as an estimate of the remaining interfering symbols. These algorithms facilitate the recovery of the subset of $\mathbf{S}[i]$ from the channel output by canceling the effect of the estimated interference using knowledge of the channel parameters, and specifically, how each interfering symbol contributes to the channel output. In the sequel we present how interference cancellation, and specifically iterative SIC [26], can be learned from data without prior knowledge of the channel model $p_{\mathbf{Y}|\mathbf{S}}$.

5.2. Symbol detector derivation

Here, we design a data-driven method for recovering $\mathbf{S}[i]$ from the channel output $\mathbf{Y}[i]$. In particular, in our model the receiver knows the constellation \mathcal{S} , and that the channel is stationary and memoryless. We do not assume that the channel is linear nor that the receiver knows the conditional probability measure $p_{\mathbf{Y}|\mathbf{S}}$. Following the approach detailed in Section 3, we design our network to implement interference cancellation in a data-driven fashion. In particular, our proposed receiver is based on the iterative SIC algorithm proposed in [26]. Therefore, as a preliminary step to designing the data-driven detector, we first review iterative SIC, after which we present DeepSIC, which is an ML-based implementation of iterative SIC.

5.2.1. Iterative soft interference cancellation The iterative SIC algorithm proposed in [26] is a multiuser detection method that combines multi-stage interference cancellation with soft decisions. Broadly speaking, the detector operates in an iterative fashion where, in each iteration, an estimate of the conditional distribution of $S_k[i]$ given the observed $\mathbf{Y}[i] = \mathbf{y}$ is generated for every user $k \in \mathcal{K}$ using the corresponding estimates of the interfering symbols $\{S_l[i]\}_{l \neq k}$ obtained in the previous iteration. Iteratively repeating this procedure refines the conditional distribution estimates, allowing the detector to accurately recover each symbol from the output of the last iteration. This iterative procedure is illustrated in Fig. 12.

To formulate the algorithm, we consider a channel whose output is obtained as a linear transformation of its input corrupted by AWGN, i.e.,

$$(12) \quad \mathbf{Y}[i] = \mathbf{H}\mathbf{S}[i] + \mathbf{W}[i],$$




Figure 12: Soft iterative interference cancellation illustration.

where $\mathbf{H} \in \mathcal{R}^{n_r \times K}$ is an a-priori known channel matrix, and $\mathbf{W}[i] \in \mathcal{R}^K$ is a zero-mean Gaussian vector with covariance $\sigma_w^2 \mathbf{I}_K$, independent of $\mathbf{S}[i]$.

Iterative SIC consists of Q iterations. Each iteration indexed $q \in \{1, 2, \dots, Q\} \triangleq \mathcal{Q}$ generates K distribution vectors $\hat{\mathbf{p}}_k^{(q)} \in \mathcal{R}^m$, $k \in \mathcal{K}$. These vectors are computed from the channel output \mathbf{y} as well as the distribution vectors obtained at the previous iteration, $\{\hat{\mathbf{p}}_k^{(q-1)}\}_{k=1}^K$. The entries of $\hat{\mathbf{p}}_k^{(q)}$ are estimates of the distribution of $S_k[i]$ for each possible symbol in \mathcal{S} , given the channel output $\mathbf{Y}[i] = \mathbf{y}$ and assuming that the interfering symbols $\{S_l[i]\}_{l \neq k}$ are distributed via $\{\hat{\mathbf{p}}_l^{(q-1)}\}_{l \neq k}$. Every iteration consists of two steps, carried out in parallel for each user: *Interference cancellation*, and *soft decoding*. Focusing on the k th user and the q th iteration, the interference cancellation stage first computes the expected values and variances of $\{S_l[i]\}_{l \neq k}$ based on $\{\hat{\mathbf{p}}_l^{(q-1)}\}_{l \neq k}$. Letting $\{\alpha_j\}_{j=1}^m$ be the indexed elements of the constellation set \mathcal{S} , the expected values and variances are computed via $e_l^{(q-1)} = \sum_{\alpha_j \in \mathcal{S}} \alpha_j (\hat{\mathbf{p}}_l^{(q-1)})_j$, and $v_l^{(q-1)} = \sum_{\alpha_j \in \mathcal{S}} (\alpha_j - e_l^{(q-1)})^2 (\hat{\mathbf{p}}_l^{(q-1)})_j$, respectively. The contribution of the interfering symbols from \mathbf{y} is then canceled by replacing them with $\{e_l^{(q-1)}\}$ and subtracting their resulting term. Letting \mathbf{h}_l be the l th column of \mathbf{H} , the interference canceled channel output is given by

$$(13) \quad \mathbf{Z}_k^{(q)}[i] = \mathbf{Y}[i] - \sum_{l \neq k} \mathbf{h}_l e_l^{(q-1)} = \mathbf{h}_k S_k[i] + \sum_{l \neq k} \mathbf{h}_l (S_l[i] - e_l^{(q-1)}) + \mathbf{W}[i].$$

Substituting the channel output \mathbf{y} into (13), the realization of the interference canceled $\mathbf{Z}_k^{(q)}[i]$, denoted $\mathbf{z}_k^{(q)}$, is obtained.

To implement soft decoding, it is assumed that $\tilde{\mathbf{W}}_k^{(q)}[i] \triangleq \sum_{l \neq k} \mathbf{h}_l (S_l[i] - e_l^{(q-1)}) + \mathbf{W}[i]$ obeys a zero-mean Gaussian distribution, independent of $S_k[i]$,

and that its covariance is given by $\Sigma_{\tilde{W}_k^{(q)}} = \sigma_w^2 \mathbf{I}_K + \sum_{l \neq k} v_l^{(q-1)} \mathbf{h}_l \mathbf{h}_l^T$. Combining this assumption with (13), the conditional distribution of $\mathbf{Z}_k^{(q)}$ given $S_k[i] = \alpha_j$ is multivariate Gaussian with mean value $\mathbf{h}_k \alpha_j$ and covariance $\Sigma_{\tilde{W}_k^{(q)}}$. Since $\mathbf{Z}_k^{(q)}[i]$ is given by a bijective transformation of $\mathbf{Y}[i]$, it holds that $p_{S_k|\mathbf{Y}}(\alpha_j|\mathbf{y}) = p_{S_k|\mathbf{Z}_k^{(q)}}(\alpha_j|\mathbf{z}_k^{(q)})$ for each $\alpha_j \in \mathcal{S}$ under the above assumptions. Consequently, the conditional distribution of $S_k[i]$ given $\mathbf{Y}[i]$ is approximated from the conditional distribution of $\mathbf{Z}_k^{(q)}$ given $S_k[i]$ via Bayes theorem. Since the symbols are equiprobable, this estimated conditional distribution is computed as

$$(14) \quad \left(\hat{\mathbf{p}}_k^{(q)}\right)_j = \frac{\exp\left\{-\frac{1}{2}\left(\mathbf{z}_k^{(q)} - \mathbf{h}_k \alpha_j\right)^T \Sigma_{\tilde{W}_k^{(q)}}^{-1}\left(\mathbf{z}_k^{(q)} - \mathbf{h}_k \alpha_j\right)\right\}}{\sum_{\alpha_{j'} \in \mathcal{S}} \exp\left\{-\frac{1}{2}\left(\mathbf{z}_k^{(q)} - \mathbf{h}_k \alpha_{j'}\right)^T \Sigma_{\tilde{W}_k^{(q)}}^{-1}\left(\mathbf{z}_k^{(q)} - \mathbf{h}_k \alpha_{j'}\right)\right\}}.$$

After the final iteration, the symbols are decoded by taking the symbol that maximizes the estimated conditional distribution for each user, i.e.,

$$(15) \quad \hat{s}_k = \arg \max_{j \in \{1, \dots, m\}} \left(\hat{\mathbf{p}}_k^{(Q)}\right)_j.$$

The overall joint detection scheme is summarized below as Algorithm 2. The initial estimates $\{\hat{\mathbf{p}}_k^{(0)}\}_{k=1}^K$ can be arbitrarily set. For example, these may be chosen based on a linear separate estimation of each symbol for \mathbf{y} , as proposed in [26].

Algorithm 2 Iterative Soft Interference Cancellation Algorithm [26]

- 1: Input: Channel output \mathbf{y} .
 - 2: Initialization: Set $q = 1$, and generate an initial guess of $\{\hat{\mathbf{p}}_k^{(0)}\}_{k=1}^K$.
 - 3: Compute the expected values $\{e_k^{(q-1)}\}$ and variances $\{v_k^{q-1}\}$.
 - 4: *Interference cancellation*: For each $k \in \mathcal{K}$ compute $\mathbf{z}_k^{(q)}$ via (13).
 - 5: *Soft decoding*: For each $k \in \mathcal{K}$, estimate $\hat{\mathbf{p}}_k^{(q)}$ via (14).
 - 6: Set $q := q + 1$. If $q \leq Q$ go to Step 3.
 - 7: Output: Hard decoded output $\hat{\mathbf{s}}$, obtained via (15).
-

Iterative SIC has several notable advantages as a joint decoding method: In terms of computational complexity, it replaces the joint exhaustive search over all different channel input combinations, required by the MAP decoder (11), with a set of computations carried out separately for each user.

Hence, its computational complexity only grows linearly with the number of users [38], making it feasible also with large values of K . Unlike conventional separate decoding, in which the symbol of each user is recovered individually while treating the interference as noise, the iterative procedure refines the separate estimates sequentially, and the usage of soft values mitigates the effect of error propagation. Consequently, Algorithm 2 is capable of achieving performance approaching that of the MAP detector, which is only feasible for small values of K .

Iterative SIC is specifically designed for linear channels of the form (12). In particular, the interference cancellation in Step 4 of Algorithm 2 requires the contribution of the interfering symbols to be additive. This limits the application of the algorithm in non-linear channels. Additionally, the fact that the distribution of the interference canceled channel output $\mathbf{Z}_k^{(g)}$ given $S_k[i]$ is approximated as Gaussian in Algorithm 2 degrades the performance in channels which do not obey the linear Gaussian model (12). Furthermore, even when the channel obeys the linear model of (12), iterative SIC requires full CSI, i.e., knowledge of the channel matrix \mathbf{H} and the noise variance σ_w^2 , which may entail substantial overhead. The dependence on accurate CSI and the assumption of linear channels are not unique to iterative SIC, and are in fact common to most interference cancellation based joint detection algorithms [38]. These limitations motivate the design of a joint detector which exploits the computational feasibility of interference cancellation methods while operating in a data-driven fashion. We specifically select iterative SIC since it is capable of achieving MAP-comparable performance, with a structure that can be readily converted to be data-driven. This is a result of the fact that its specific model-based computations, i.e., Steps 4–5 in Algorithm 2, can be naturally implemented using relatively simple ML methods. The resulting receiver, detailed in the following, integrates ML methods into Algorithm 2, allowing it to be implemented for arbitrary memoryless stationary channels without requiring a-priori knowledge of the channel model and its parameters.

5.2.2. Data-driven receiver architecture Here, we present a receiver architecture that implements iterative SIC in a data-driven fashion. Following the approach of Section 3, we keep the overall structure of the iterative SIC algorithm, depicted in Fig. 12, while replacing the channel-model-based computations with dedicated suitable DNNs. To that aim, we note that iterative SIC can be viewed as a set of interconnected basic building blocks, each implementing the two stages of interference cancellation and soft decoding, i.e., Steps 4–5 of Algorithm 2. While the high level architecture of Fig. 12

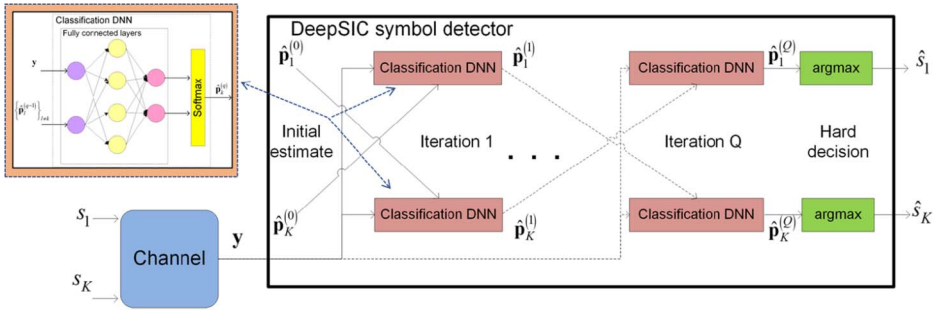


Figure 13: DeepSIC illustration.

is ignorant of the underlying channel model, its basic building blocks are channel-model-dependent. In particular, interference cancellation requires the contribution of the interference to be additive, i.e., a linear model channel as in (12), as well as full CSI, in order to cancel the contribution of the interference. Soft decoding requires complete knowledge of the channel input-output relationship in order to estimate the conditional probabilities via (14).

Although each of these basic building blocks consists of two sequential procedures which are completely channel-model-based, we note that the purpose of these computations is to carry out a classification task. In particular, the k th building block of the q th iteration, $k \in \mathcal{K}$, $q \in \mathcal{Q}$, produces $\hat{\mathbf{p}}_k^{(q)}$, which is an estimate of the conditional distribution of $S_k[i]$ given $\mathbf{Y}[i] = \mathbf{y}$ based on $\{\hat{\mathbf{p}}_l^{(q-1)}\}_{l \neq k}$. Such computations are naturally implemented by classification DNNs, e.g., fully-connected networks with a softmax output layer. Embedding these ML-based conditional distribution computations into the iterative SIC block diagram in Fig. 12 yields the overall receiver architecture depicted in Fig. 13. We set the initial estimates $\{\hat{\mathbf{p}}_k^{(0)}\}_{k=1}^K$ to represent a uniform distribution, i.e., $(\hat{\mathbf{p}}_k^{(0)})_j = \frac{1}{m}$ for each $j \in \{1, 2, \dots, m\}$ and $k \in \mathcal{K}$. The resulting data-driven implementation of Algorithm 2 is repeated below as Algorithm 3. Note that the model-based Steps 3–5 of Algorithm 2, which estimate the conditional distributions, are replaced with the ML-based conditional distribution estimation Step 3 in Algorithm 3.

A major advantage of using classification DNNs as the basic building blocks in Fig. 13 stems from the fact that such ML-based methods are capable of accurately computing conditional distributions in complex non-linear setups without requiring a-priori knowledge of the channel model and its parameters. Consequently, when these building blocks are trained to properly implement their classification task, the receiver essentially realizes iterative

Algorithm 3 Deep Soft Interference Cancellation (DeepSIC)

- 1: **Input:** Channel output \mathbf{y} .
 - 2: **Initialization:** Set $q = 1$, and generate an initial guess of the conditional distributions $\{\hat{\mathbf{p}}_k^{(0)}\}_{k=1}^K$.
 - 3: *ML-based conditional distribution estimation:* For each $k \in \mathcal{K}$, estimate the conditional distribution $\hat{\mathbf{p}}_k^{(q)}$ from \mathbf{y} and $\{\hat{\mathbf{p}}_l^{(q-1)}\}_{l \neq k}$ using the (q, k) th classification DNN.
 - 4: Set $q := q + 1$. If $q \leq Q$ go to Step 3.
 - 5: **Output:** Hard decoded output $\hat{\mathbf{s}}$, obtained via (15).
-

soft interference cancellation for arbitrary channel models in a data-driven fashion.

5.2.3. Training the DNNs In order for the ML-based receiver structure of Fig. 13 to reliably implement joint decoding, its building block classification DNNs must be properly trained. Here, we consider two possible approaches to train the receiver based on a set of n_t pairs of channel inputs and their corresponding outputs, denoted $\{\mathbf{s}_i, \mathbf{y}_i\}_{i=1}^{n_t}$: *End-to-end training*, and *sequential training*.

End-to-end training The first approach jointly trains the entire network, i.e., all the building block DNNs. Since the output of the deep network is the set of conditional distributions $\{\hat{\mathbf{p}}_k^{(Q)}\}_{k=1}^K$, where each $\hat{\mathbf{p}}_k^{(Q)}$ is used to estimate $S_k[i]$, we use the sum cross entropy as the training objective. Let $\boldsymbol{\theta}$ be the network parameters, and $\hat{\mathbf{p}}_k^{(Q)}(\mathbf{y}, \alpha; \boldsymbol{\theta})$ be the entry of $\hat{\mathbf{p}}_k^{(Q)}$ corresponding to $S_k[i] = \alpha$ when the input to the network parameterized by $\boldsymbol{\theta}$ is \mathbf{y} . The sum cross entropy loss is

$$(16) \quad \mathcal{L}_{\text{SumCE}}(\boldsymbol{\theta}) = \frac{1}{n_t} \sum_{i=1}^{n_t} \sum_{k=1}^K -\log \hat{\mathbf{p}}_k^{(Q)}(\mathbf{y}_i, (s_i)_k; \boldsymbol{\theta}).$$

Training the receiver in Fig. 13 in an end-to-end manner based on the loss (16) jointly updates the coefficients of all the $K \cdot Q$ building block DNNs. Since for a large number of users, training so many parameters simultaneously is expected to require a large labeled set, we next propose a sequential training approach.

Sequential training To allow the network to be trained with a reduced number of training samples, we note that the goal of each building block DNN does not depend on the iteration index: The k th building block of the

q th iteration outputs a soft estimate of $S_k[i]$ for each $q \in \mathcal{Q}$. Therefore, each building block can be trained individually to minimize the cross entropy loss. To formulate this objective, let $\boldsymbol{\theta}_k^{(q)}$ be the parameters of the k th DNN at iteration q , and write $\hat{\boldsymbol{p}}_k^{(q)}(\mathbf{y}, \{\hat{\boldsymbol{p}}_l^{(q-1)}\}_{l \neq k}, \alpha; \boldsymbol{\theta}_k^{(q)})$ as the entry of $\hat{\boldsymbol{p}}_k^{(q)}$ corresponding to $S_k[i] = \alpha$ when its inputs are \mathbf{y} and $\{\hat{\boldsymbol{p}}_l^{(q-1)}\}_{l \neq k}$. The cross entropy loss is

$$(17) \quad \mathcal{L}_{\text{CE}}(\boldsymbol{\theta}_k^{(q)}) = \frac{1}{n_t} \sum_{i=1}^{n_t} -\log \hat{\boldsymbol{p}}_k^{(q)}(\mathbf{y}_i, \{\hat{\boldsymbol{p}}_{i,l}^{(q-1)}\}_{l \neq k}, (s_i)_k; \boldsymbol{\theta}_k^{(q)}),$$

where $\{\hat{\boldsymbol{p}}_{i,l}^{(q-1)}\}$ represent the estimated probabilities associated with \mathbf{y}_i computed at the previous iteration. The problem with training each DNN individually is that the soft estimates $\{\hat{\boldsymbol{p}}_{i,l}^{(q-1)}\}$ are not provided as part of the training set. This challenge can be tackled by training the DNNs corresponding to each layer in a sequential manner, where for each layer the outputs of the trained DNN corresponding to the previous iterations are used as the soft estimates fed as training samples.

Sequential training uses the n_t input-output pairs to train each DNN individually. Compared to the end-to-end training that utilizes the training samples to learn the complete set of parameters, which can be quite large, sequential training uses the same data set to learn a significantly smaller number of parameters, reduced by a factor of $K \cdot Q$, multiple times. Consequently, this approach is expected to require much fewer training samples, at the cost of a longer learning procedure for a given training set, due to its sequential operation, and possible performance degradation as the building blocks are not jointly trained. This behavior is numerically demonstrated in the simulation study detailed in Section 5.3.

5.3. Numerical evaluations

In the following section we numerically evaluate DeepSIC. We train DeepSIC using the ADAM optimizer [37] with a relatively small training set of 5000 training samples, and tested over 20000 symbols. In the implementation of the DNN-based building blocks of DeepSIC, we used a different fully-connected network for each training method: For end-to-end training, where all the building blocks are jointly trained, we used a compact network consisting of a $(n_r + K - 1) \times 60$ layer followed by ReLU activation and a $60 \times m$ layer. For sequential training, which sequentially adapts subsets of the building blocks and can thus tune a more parameters using the same training set


(or, alternatively, requires a smaller training set) compared to end-to-end training, we used three fully-connected layers: An $(n_r + K - 1) \times 100$ first layer, a 100×50 second layer, and a $50 \times m$ third layer, with a sigmoid and a ReLU intermediate activation functions, respectively. In both iterative SIC as well as DeepSIC, we set the number of iterations to $Q = 5$.

We first consider a linear AWGN channel as in (12). Recall that iterative SIC as well as previously proposed unfolding-based data-driven MIMO receivers [11, Sec. II], are all designed for such channels. Consequently, the following study compares DeepSIC in terms of performance and robustness to competing detectors in a scenario for which these previous schemes are applicable. In particular, we evaluate the SER of the following MIMO detectors: The MAP detector, given by (11); The iterative SIC algorithm (Algorithm 2); DeepSIC with the sequential training method, referred to in the following as *Seq. DeepSIC*; DeepSIC with end-to-end training based on the sum cross entropy loss (16), referred to henceforth as *E2E DeepSIC*; and the unfolding-based *DetNet* MIMO detector [17].

The model-based MAP and iterative SIC detectors, as well as DetNet [17], all require CSI, and specifically, accurate knowledge of the channel matrix \mathbf{H} . DeepSIC operates without a-priori knowledge of the channel model and its parameters, learning the decoding mapping from a training set sampled from the considered input-output relationship. In order to compare the robustness of the detectors to CSI uncertainty, we also evaluate them when the receiver has access to an estimate of \mathbf{H} with entries corrupted by i.i.d. additive Gaussian noise whose variance is given by σ_e^2 times the magnitude of the corresponding entry, where $\sigma_e^2 > 0$ is referred to as the *error variance*. For DeepSIC, which is model-invariant, we compute the SER under CSI uncertainty by using a training set whose samples are randomized from a channel in which the true \mathbf{H} is replaced with its noisy version.

We simulate 6×6 linear Gaussian channel, i.e., $K = 6$ users and $n_r = 6$ receive antennas. The symbols are randomized from a BPSK constellation, namely, $\mathcal{S} = \{-1, 1\}$ and $m = |\mathcal{S}| = 2$. The channel matrix \mathbf{H} models spatial exponential decay, and its entries are given by $(\mathbf{H})_{i,k} = e^{-|i-j|}$, for each $i \in \{1, \dots, n_r\}$, $k \in \mathcal{K}$. For each channel, the SER of the considered receivers is evaluated for both perfect CSI, i.e., $\sigma_e^2 = 0$, as well as CSI uncertainty, for which we use $\sigma_e^2 = 0.1$. The evaluated SER versus the SNR, defined as $1/\sigma_w^2$, is depicted in Fig. 14.

Observing Fig. 14, we note that the performance of DeepSIC with end-to-end training approaches that of the model-based iterative SIC algorithm, which is within a small gap of the optimal MAP performance. This demonstrates the ability of DeepSIC to implement iterative SIC in a data-driven

Figure 14: 6×6 AWGN channel.


fashion. The sequential training method, whose purpose is to allow DeepSIC to train with smaller data sets compared to end-to-end training, also achieves an SER that is comparable to iterative SIC. In the presence of CSI uncertainty, DeepSIC is observed to substantially outperform the model-based iterative SIC and MAP receivers, as well as DetNet operating with a noisy version of \mathbf{H} and trained with a hundred times more samples. In particular, it follows from Fig. 14 that a relatively minor error of variance $\sigma_e^2 = 0.1$ severely deteriorates the performance of the model-based methods, while DeepSIC is hardly affected by the same level of CSI uncertainty.

Next, we consider a Poisson channel. We use $K = 4$ and $n_r = 4$. Here, the symbols are randomized from an on-off keying for which $\mathcal{S} = \{0, 1\}$. The entries of the channel output are related to the input via the conditional distribution

$$(18) \quad (\mathbf{Y}[i])_j | \mathcal{S}[i] \sim \mathbb{P} \left(\frac{1}{\sqrt{\sigma_w^2}} (\mathbf{H}\mathcal{S}[i])_j + 1 \right), \quad j \in \{1, \dots, n_r\},$$

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

The achievable SER of DeepSIC versus SNR under both perfect CSI as well as CSI uncertainty with error variance $\sigma_e^2 = 0.1$ is compared to the MAP and iterative SIC detectors in Fig. 15. Observing Fig. 15, we again note that the performance of DeepSIC is only within a small gap of the MAP performance with perfect CSI, and that the data-driven receiver is more robust to CSI uncertainty compared to the model-based MAP. In particular, DeepSIC with sequential training, which utilizes a deeper network architecture for each building block, outperforms here end-to-end training with

Figure 15: 4×4 Poisson channel.

basic two-layer structures for the conditional distribution estimation components. We conclude that under such non-Gaussian channels, more complex DNN models are required to learn to cancel interference and carry out soft detection accurately. This further emphasizes the gain of our proposed sequential approach for training each building block separately, thus allowing to train an overall deep architecture using a limited training set based on the understanding of the role of each of its components. Furthermore, iterative SIC, which is designed for linear Gaussian channels (12) where interference is additive, achieves very poor performance when the channel model is substantially different from (12).

Our numerical results demonstrate the ability of DeepSIC to achieve excellent performance through learning from data for statistical models where model-based interference cancellation is effectively inapplicable.

6. Discussion and conclusion

We reviewed an ML-based approach for designing symbol detection algorithms for communication systems. This approach introduces ML into well-known algorithms such as Viterbi, BCJR, and MIMO SIC, by identifying the computations that require full knowledge of the underlying channel input-output statistical relationships, and replacing them with ML-based algorithms. The resulting architecture combines ML-based processing with conventional symbol detection schemes. Our numerical results demonstrate that the performance of data-driven ViterbiNet, BCJRNet, and DeepSIC approaches the optimal performance of the equivalent CSI-based versions,

and outperform previously proposed ML-based symbol detectors using a small amount of training data. It is also illustrated that these algorithms are capable of operating in the presence of CSI uncertainty with little performance degradation. Since our approach relies on small neural networks that can be trained quickly to achieve close to optimal performance with only a few thousand training symbols, it paves the way to the possibility of online training using pilot sequences.

As part of future work, we will be extending our approach to general factor graph methods as well as using meta learning to improve the adaptability of our algorithms to changing channels. Finally, we will also investigate how different density estimation methods can be used to further improve the training process.

References

- [1] A. Goldsmith, *Wireless communications*. Cambridge University Press, 2005.
- [2] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge University Press, 2005.
- [3] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [4] Y. Bengio, “Learning deep architectures for AI,” *Foundations and trends® in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.
- [5] K. Gregor and Y. LeCun, “Learning fast approximations of sparse coding,” in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 2010, pp. 399–406.
- [6] V. Monga, Y. Li, and Y. C. Eldar, “Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing,” *arXiv preprint arXiv:1912.10557*, 2019.
- [7] T. O’Shea and J. Hoydis, “An introduction to deep learning for the physical layer,” *IEEE Trans. on Cogn. Commun. Netw.*, vol. 3, no. 4, pp. 563–575, 2017.
- [8] O. Simeone, “A very brief introduction to machine learning with applications to communication systems,” *IEEE Trans. on Cogn. Commun. Netw.*, vol. 4, no. 4, pp. 648–664, 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, 2018.

- [10] D. Gündüz, P. de Kerret, N. D. Sidiropoulos, D. Gesbert, C. R. Murthy, and M. van der Schaar, “Machine learning in the air,” *IEEE J. Sel. Areas Commun.*, vol. 37, no. 10, pp. 2184–2199, 2019.
- [11] A. Balatsoukas-Stimming and C. Studer, “Deep unfolding for communications systems: A survey and some new directions,” *arXiv preprint arXiv:1906.05774*, 2019.
- [12] H. Kim, Y. Jiang, R. Rana, S. Kannan, S. Oh, and P. Viswanath, “Communication algorithms via deep learning,” *arXiv preprint arXiv:1805.09317*, 2018.
- [13] 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, 2018. [MR3884587](#)
- [14] A. Caciularu and D. Burshtein, “Blind channel equalization using variational autoencoders,” in *Proc. IEEE ICC*, 2018.
- [15] J. R. Hershey, J. L. Roux, and F. Weninger, “Deep unfolding: Model-based inspiration of novel deep architectures,” *arXiv preprint arXiv:1409.2574*, 2014.
- [16] O. Solomon, R. Cohen, Y. Zhang, Y. Yang, Q. He, J. Luo, R. J. van Sloun, and Y. C. Eldar, “Deep unfolded robust PCA with application to clutter suppression in ultrasound,” *IEEE Trans. Med. Imag.*, 2019.
- [17] N. Samuel, T. Diskin, and A. Wiesel, “Learning to detect,” *IEEE Trans. Signal Process.*, vol. 67, no. 10, pp. 2554–2564, 2019. [MR3957688](#)
- [18] S. Takabe, M. Imanishi, T. Wadayama, and K. Hayashi, “Deep learning-aided projected gradient detector for massive overloaded MIMO channels,” in *Proc. IEEE ICC*, 2019.
- [19] H. He, C.-K. Wen, S. Jin, and G. Y. Li, “A model-driven deep learning network for MIMO detection,” in *Proc. IEEE GlobalSIP*, 2018. [MR4114651](#)
- [20] S. Khobahi, N. Naimipour, M. Soltanalian, and Y. C. Eldar, “Deep signal recovery with one-bit quantization,” in *Proc. IEEE ICASSP*, 2019.
- [21] N. Shlezinger, N. Farsad, Y. C. Eldar, and A. J. Goldsmith, “Viterbi-Net: A deep learning based Viterbi algorithm for symbol detection,” *IEEE Trans. Wireless Commun.*, vol. 19, no. 5, pp. 3319–3331, 2020.

- [22] N. Shlezinger, N. Farsad, Y. C. Eldar, and A. J. Goldsmith, “Data-driven factor graphs for deep symbol detection,” *arXiv preprint arXiv:2002.00758*, 2020.
- [23] N. Shlezinger, R. Fu, and Y. C. Eldar, “DeepSIC: Deep soft interference cancellation for multiuser MIMO detection,” *arXiv preprint arXiv:2002.03214*, 2020.
- [24] A. Viterbi, “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm,” *IEEE Trans. Inf. Theory*, vol. 13, no. 2, pp. 260–269, 1967.
- [25] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, “Optimal decoding of linear codes for minimizing symbol error rate,” *IEEE Trans. Inf. Theory*, vol. 20, no. 2, pp. 284–287, 1974. [MR0368917](#)
- [26] W.-J. Choi, K.-W. Cheong, and J. M. Cioffi, “Iterative soft interference cancellation for multiple antenna systems.” in *Proc. WCNC*, 2000, pp. 304–309.
- [27] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT Press, 2016, <http://www.deeplearningbook.org>. [MR3617773](#)
- [28] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro, “Exploring generalization in deep learning,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5947–5956.
- [29] G. D. Forney, “The Viterbi algorithm,” *Proc. IEEE*, vol. 61, no. 3, pp. 268–278, 1973. [MR0439384](#)
- [30] A. V. D. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [31] G. McLachlan and D. Peel, *Finite mixture models*. John Wiley & Sons, 2004. [MR1789474](#)
- [32] C. M. Bishop, “Mixture density networks,” 1994. [Online]. Available: <http://publications.aston.ac.uk/id/eprint/373/>
- [33] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, “Factor graphs and the sum-product algorithm,” *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, 2001. [MR1820474](#)
- [34] G. D. Forney, “Codes on graphs: Normal realizations,” *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 520–548, 2001. [MR1820475](#)

- [35] H.-A. Loeliger, “An introduction to factor graphs,” *IEEE Signal Process. Mag.*, vol. 21, no. 1, pp. 28–41, 2004.
- [36] J. M. Cioffi, “Sequence detection,” *EE379B Course notes chapter 9. Stanford University*, 2008.
- [37] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [38] J. G. Andrews, “Interference cancellation for cellular systems: a contemporary overview,” *IEEE Wireless Commun.*, vol. 12, no. 2, pp. 19–29, 2005.

NARIMAN FARSAD
DEPARTMENT OF COMPUTER SCIENCE
RYERSON UNIVERSITY
TORONTO
CANADA
E-mail address: nfarsad@ryerson.edu

NIR SHLEZINGER
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING
BEN-GURION UNIVERSITY OF THE NEGEV
BE'ER-SHEVA
ISRAEL
E-mail address: nirshl@bgu.ac.il

ANDREA J. GOLDSMITH
DEPARTMENT OF ELECTRICAL ENGINEERING
PRINCETON UNIVERSITY
PRINCETON, NJ, 08540
USA
E-mail address: goldsmith@princeton.edu

YONINA C. ELДАР
FACULTY OF MATH AND CS
WEIZMANN INSTITUTE OF SCIENCE
REHOVOT
ISRAEL
E-mail address: yonina.eldar@weizmann.ac.il

RECEIVED FEBRUARY 14, 2020