Data-Driven Factor Graphs for Deep Symbol Detection

Nir Shlezinger, Nariman Farsad, Yonina C. Eldar, and Andrea J. Goldsmith

Abstract

Many important schemes in signal processing and communications, ranging from the BCJR algorithm to the Kalman filter, are instances of factor graph methods. This family of algorithms is based on recursive message passing-based computations carried out over graphical models, representing a factorization of the underlying statistics. Consequently, in order to implement these algorithms, one must have accurate knowledge of the statistical model of the considered signals. In this work we propose to implement factor graph methods in a data-driven manner. In particular, we propose to use machine learning (ML) tools to learn the factor graph, instead of the overall system task, which in turn is used for inference by message passing over the learned graph. We apply the proposed approach to learn the factor graph representing a finite-memory channel, demonstrating the resulting ability to implement BCJR detection in a data-driven fashion. We demonstrate that the proposed system, referred to as BCJRNet, learns to implement the BCJR algorithm from a small training set, and that the resulting receiver exhibits improved robustness to inaccurate training compared to the conventional channel-model-based receiver operating under the same level of uncertainty. Our results indicate that by utilizing ML tools to learn factor graphs from labeled data, one can implement a broad range of model-based algorithms, which traditionally require full knowledge of the underlying statistics, in a data-driven fashion.

I. INTRODUCTION

A broad range of algorithms in communications, signal processing, statistics, and machine learning are obtained by message passing over graphical models [1]. In particular, the combination of factor graphs, and specifically, the factor graph model proposed by Forney in [2], with the sum-product message passing scheme [3], was shown to provide a unified framework which specializes many important algorithms [4].

Broadly speaking, factor graphs provide a graphical model for a function of multiple variables, commonly a joint distribution measure. The resulting model allows prohibitive computations to be carried out in a recursive manner at controllable complexity by message passing along the graph [4]. Some important instances of factor graph methods include the BCJR symbol detector [5]; the Kalman filter [6]; and hidden Markov model (HMM) predictors [7]. Furthermore, the application of message passing over factor graphs was shown to give rise to efficient and accurate methods for, e.g., equalization [8], [9], multiple-input multiple-output detection [10], and joint decoding and channel estimation [11], [12].

The aforementioned algorithms require prior knowledge on the underlying statistical model. For example, in order to apply the BCJR symbol detector, one must know or have an accurate estimate of the channel model and its parameters, i.e., full channel state information (CSI). This limits the application of these methods in setups where the statistical model is complex, difficult to estimate, or poorly understood. Furthermore, a potentially large overhead is required in order to estimate the model parameters, from which the factor graph is obtained, and the reliability of message-passing methods may be substantially degraded in the presence of model inaccuracy.

In our previous works [13] and [14] we proposed data-driven receivers which learn to implement two important symbol detection algorithms: the Viterbi algorithm [15], and the iterative soft interference cancellation scheme [16]. By using dedicated machine learning (ML) tools to learn the graphical model used by the channel-model-based symbol detector algorithm, e.g., the trellis diagram for the Viterbi algorithm, we obtained systems which implement the symbol detection scheme in a data-driven fashion using a relatively small amount of training samples, while exhibiting improved robustness to CSI uncertainty compared to their model-based counterpart. Our success in implementing a model-based algorithm in a data-driven manner by learning its underlying graphical model, combined with fact that the Viterbi algorithm is a factor graph method, motivates the study of a data-driven implementation for a broader family of factor graph methods.

In this paper we take a first step in examining the ability to implement data-driven algorithms by utilizing ML tools to learn the underlying factor graph. We focus on problems which can be solved using factor graph methods when the model is known, proposing to utilize ML tools to learn the factor graph, and to carry out the task using conventional message passing over the learned graph. This is opposed to the common approach in ML, which is based on training deep neural networks (DNNs) to carry out the task in an end-to-end manner [17]. This approach which

combines ML and model-based algorithms in a hybrid manner, having the potential of achieving the performance guarantees and controllable complexity of model-based methods while operating in a data-driven manner and requiring relatively small training sets.

We show how the proposed approach of learning factor graphs can be applied to symbol detection in finite-memory channels. In particular, we detail how this strategy can be used to implement the BCJR algorithm, which is the maximum a-posteriori probability (MAP) symbol detector for such channels, in a data-driven channel-model-independent manner. The resulting data-driven receiver, referred to as *BCJRNet*, is capable of learning to carry out MAP symbol detection from a relatively small training set, without requiring prior knowledge of the channel model and its parameters. Our numerical studies also show that the BCJRNet exhibits improved resiliency to inaccurate training compared to the model-based BCJR algorithm operating under the same level of uncertainty. Our results demonstrate the potential gains of properly combining ML and model-based algorithms in realizing optimal-approaching data-driven methods in communications and signal processing.

The rest of this paper is organized as follows: In Section II we briefly review factor graph methods, and propose the concept of data-driven factor graphs. Section III presents the application of this approach to realize a data-driven BCJR detector. Section IV provides numerical performance results of BCJRNet, while Section V 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., $\boldsymbol{x}$ is a deterministic vector, and $\boldsymbol{X}$ is a random vector; the ith element of $\boldsymbol{x}$ is written as $(\boldsymbol{x})_i$. The probability measure of an RV X evaluated at x is denoted $P_X(x)$. We use caligraphic letters for sets, e.g., $\mathcal{X}$, where $|\mathcal{X}|$ is the cardinality of a finite set $\mathcal{X}$, and $\mathcal{R}$ is the set of real numbers.

II. DATA-DRIVEN FACTOR GRAPHS

In this section we present the concept of data-driven factor graphs. We begin with a brief review of conventional model-based factor graph methods in Subsection II-A, after which we discuss how these graphs can be learned from labeled data using ML tools in Subsection II-B.

A. Model-Based Factor Graphs

In the following we provide a brief introduction to factor graphs, focusing on the model proposed by Forney in [2], known as *Forney-style factor graphs*. We then review the sum-product method for computing marginal distributions using factor graphs [3]. The main example

considered in the sequel is that of the factor graph representation of finite-state channels, which is utilized in Section III for formulating the data-driven BCJRNet MAP symbol detector.

A factor graph is a graphical representation of the factorization of a function of several variables [4], commonly a joint distribution measure. Consider an $n \times 1$ random vector $\boldsymbol{X} \in \mathcal{X}^n$ where $\mathcal{X}$ is a finite set, i.e., the entries of $\boldsymbol{X}$, denoted $\{X_i\}$, are discrete RVs. The joint distribution of $\boldsymbol{X}$, $P_{\boldsymbol{X}}(\boldsymbol{x})$, can be factorized if it can be represented as the product of m functions $\{f_i(\cdot)\}_{i=1}^m$, i.e., there exists some partition variables $\{\mathcal{V}_i\}_{i=1}^m$ where $\mathcal{V}_i \subset \{x_1, \ldots, x_n\}$ such that

$$P_{\boldsymbol{X}}(\boldsymbol{x}) = \prod_{i=1}^m f_i(\mathcal{V}_i). \tag{1}$$

In order to represent (1) using a factor graph, the functions $\{f_i(\cdot)\}_{i=1}^m$ should be set such that each variable x_k appears in no more than two partitions[1] $\{\mathcal{V}_i\}_{i=1}^m$. Subject to this assumption, the joint distribution $P_{\boldsymbol{X}}(\boldsymbol{x})$ can be represented using a factor graph with m factor nodes, which are the functions $\{f_i(\cdot)\}_{i=1}^m$, and n edges $\{x_k\}_{k=1}^n$, where variables appearing only in a single partition are treated as half-edges.

Example 1. *A* finite-memory channel *models a statistical relationship between a channel input $X_i \in \mathcal{X}$ and a channel output Y_i, such that Y_i depends only on the last $L > 0$ channel inputs $X_i, \ldots X_{i-L+1}$. For a block of n i.i.d. input symbols $\boldsymbol{X}$, the channel input-output relationship satisfies*

$$P_{\boldsymbol{Y}|\boldsymbol{X}}(\boldsymbol{y}|\boldsymbol{x}) = \prod_{i=1}^n P_{Y_i|X_i,\ldots,X_{i=L+1}}(y_i|x_i, \ldots, x_{i-L+1}), \tag{2}$$

where $X_i \equiv 0$ for $i \leq 0$. By defining the $L \times 1$ state vector $\boldsymbol{S}_i \triangleq [X_i, \ldots X_{i-L+1}]^T$ and stacking the states over the entire block into a vector $\boldsymbol{S} \in \mathcal{X}^{Ln}$, we can write the joint state-output distribution as

$$
\begin{aligned}
P_{\boldsymbol{Y},\boldsymbol{S}}(\boldsymbol{y}, \boldsymbol{s}) &= P_{\boldsymbol{Y}|\boldsymbol{S}}(\boldsymbol{y}|\boldsymbol{s}) P_{\boldsymbol{S}}(\boldsymbol{s}) \\
&= \prod_{i=1}^n P_{Y_i|\boldsymbol{S}_i}(y_i|\boldsymbol{s}_i) P_{\boldsymbol{S}_i|\boldsymbol{S}_{i-1}}(\boldsymbol{s}_i|\boldsymbol{s}_{i-1}). \tag{3}
\end{aligned}
$$

---

[1]A factorization in which a variable appears in more than two factors can always be modified to meet the above constraint by introducing additional variables and identity factors, see [4].
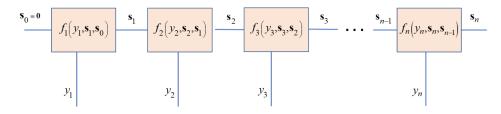
Fig. 1. Factor graph of finite-memory channel.

By defining the functions

$$f_i(y_i, \boldsymbol{s}_i, \boldsymbol{s}_{i-1}) \triangleq P_{Y_i|\boldsymbol{S}_i}(y_i|\boldsymbol{s}_i) P_{\boldsymbol{S}_i|\boldsymbol{S}_{i-1}}(\boldsymbol{s}_i|\boldsymbol{s}_{i-1}), \tag{4}$$

for each $i \in \{1, \dots, n\} \triangleq \mathcal{N}$, it holds that the channel input-output relationship (3) obeys the form of (1). The resulting factor graph, in which $\{\boldsymbol{s}_i\}_{i=1}^{n-1}$ are the edges while the remaining variables are half-edges, is depicted in Fig. 1.

A major advantage for representing joint distributions via factor graphs is that they allow some desired computations to be carried out with reduced complexity. One of the most common methods is the sum-product algorithm for evaluating a marginal distribution from a factor graph representation of a joint probability measure [3]. To formulate this method, consider the factorized distribution (1), where the ordering of the partitions $\{\mathcal{V}_i\}$ corresponds to the order of the variables $\{x_k\}$, e.g., $\mathcal{V}_1 = \{x_1, x_2\}$, $\mathcal{V}_2 = \{x_2, x_3, x_4\}$, $\mathcal{V}_3 = \{x_4, x_5\}$, etc. Furthermore, assume that the factor graph of (1) does not contain any cycles[2]. In this case, the marginal distribution of a single RV X_k whose corresponding variable appears in the partitions $\mathcal{V}_j$ and $\mathcal{V}_{j+1}$ can be computed from the joint distribution of $\boldsymbol{X}$ via

$$P_{X_k}(x_k) = \sum_{\{x_1, \dots, x_{k-1}, x_{k+1}, \dots, x_n\}} P_{\boldsymbol{X}}(x_1, \dots, x_n)$$

$$= \underbrace{\left(\sum_{\{x_1 \dots x_{k-1}\}} \prod_{i=1}^{j} f_i(\mathcal{V}_i) \right)}_{\triangleq \mu_{f_j \to x_k}(x_k)} \underbrace{\left(\sum_{\{x_{k+1} \dots x_n\}} \prod_{i=j+1}^{m} f_i(\mathcal{V}_i) \right)}_{\triangleq \mu_{f_{j+1} \to x_k}(x_k)}. \tag{5}$$

The factorization of the joint distribution in (1) implies that the marginal distribution, whose computation typically requires summation over $|\mathcal{X}|^{n-1}$ variables, can now be evaluated as the product of two terms, $\mu_{f_j \to x_k}(x_k)$ and $\mu_{f_{j+1} \to x_k}(x_k)$. These terms can be viewed as messages

[2]While there the sum-product algorithm can be extended for handling cycles, e.g., [18], we focus here on the standard application for cycle-free graphs [1], [4].

propagating forward and backward along the factor graph, e.g., $\mu_{f_j \to x_k}(x_k)$ represents a forward message conveyed from function node f_j to edge x_k. In particular, these messages can be computed recursively, e.g., by writing $\mathcal{V}_j/x_k = \{x_{k-l}, \ldots, x_{k-1}\}$ for some $l \geq 1$, then the sum-product rule [4] implies that

$$
\mu_{f_j \to x_k}(x_k) = \sum_{\{x_1,\ldots,x_{k-1}\}} \prod_{i=1}^{j} f_i(\mathcal{V}_i)
$$

$$
= \sum_{\{x_{k-l},\ldots,x_{k-1}\}} f_j(\mathcal{V}_j) \prod_{i=k-l}^{k-1} \mu_{\tilde{f}_{i,j} \to x_i}(x_i), \tag{6}
$$

where $\tilde{f}_{i,j}$ is the function node connected to the edge x_i other than f_j. This method of computing marginal distributions using forward and backward recursions along the factor graph is referred to as belief propagation [19], or the sum-product algorithm [3]. The application of this scheme in finite-memory channels is detailed in the following example:

Example 2. *Consider again the setup discussed in Example 1 representing a channel with finite memory of length $L > 0$. We are interested in computing the joint distribution of two consecutive state vectors, $\boldsymbol{S}_k$ and $\boldsymbol{S}_{k+1}$, given a realization of the channel output $\boldsymbol{Y} = \boldsymbol{y}$. Using the sum-product method, one can compute this joint distribution by recursive message passing along its factor graph. In particular,*

$$
P_{\boldsymbol{S}_k, \boldsymbol{S}_{k+1}, \boldsymbol{Y}}(\boldsymbol{s}_k, \boldsymbol{s}_{k+1}, \boldsymbol{y}) = \mu_{f_k \to \boldsymbol{s}_k}(\boldsymbol{s}_k) f_{k+1}(y_{k+1}, \boldsymbol{s}_{k+1}, \boldsymbol{s}_k)
$$

$$
\times \mu_{f_{k+2} \to \boldsymbol{s}_{k+1}}(\boldsymbol{s}_{k+1}), \tag{7}
$$

where the forward path messages are computed recursively via

$$
\mu_{f_i \to \boldsymbol{s}_i}(\boldsymbol{s}_i) = \sum_{\boldsymbol{s}_{i-1}} f_i(y_i, \boldsymbol{s}_i, \boldsymbol{s}_{i-1}) \mu_{f_{i-1} \to \boldsymbol{s}_{i-1}}(\boldsymbol{s}_{i-1}), \tag{8}
$$

for $i = 1, 2, \ldots, k$. Similarly, the backward messages are

$$
\mu_{f_{i+1} \to \boldsymbol{s}_i}(\boldsymbol{s}_i) = \sum_{\boldsymbol{s}_{i+1}} f_{i+1}(y_{i+1}, \boldsymbol{s}_{i+1}, \boldsymbol{s}_i) \mu_{f_{i+2} \to \boldsymbol{s}_{i+1}}(\boldsymbol{s}_{i+1}), \tag{9}
$$

for $i = n-1, n-2, \ldots, k+1$. An illustration of this message passing along the factor graph for $k = 2$ is depicted in Fig. 2.
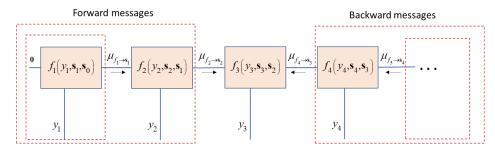
Fig. 2. Message passing over the factor graph of a finite-memory channel.

B. Learning Factor Graphs

Factor graph methods, such as the sum-product algorithm, exploit the factorization of a joint distribution to efficiently compute a desired quantity. For example, the application of the sum-product algorithm in a finite-memory channel detailed in Example 2 exploits its factorization to compute marginal distributions, an operation whose burden typically grows exponentially with th block size, with complexity that only grows linearly with n. As we show in Section III, the instance of the sum-product algorithm discussed in Example 2 is exactly the recursive computation carried out in the BCJR symbol detector [5]. In fact, the sum-product algorithm specializes a multitude of common signal processing techniques, including the Kalman filter and HMM prediction [4].

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 $\{f_j\}_{j=1}^m$. This implies that, for example, the BCJR detector, which builds upon the recursive computation in Example 2, requires prior knowledge of the channel model and its parameters, i.e., full CSI. To realize the Kalman filter, one must first know the underlying state-space equations. In some case, such accurate prior knowledge may not be available or costly to acquire.

In our previous works [13] we utilized ML tools to implement the Viterbi algorithm, which can also be considered as an instance of a factor graph methods [20], in a data-driven fashion. We used a similar approach in [14] to realize a data-driven implementation of the soft interference cancellation scheme [16], which is also a message passing-based symbol detection method. This was achieved by utilizing dedicated compact DNNs to learn the channel-model-based computations of the underlying graphical model, e.g., the log-likelihood values used for constructing the trellis diagram in the Viterbi detector, while maintaining the overall flow of the algorithm. We propose to generalize this concept to realize a more general family of data-driven factor graph methods by learning the mappings carried out at the function nodes from a relatively small set

of labeled data using ML tools. By doing so, one can train a system to learn an underlying factor graph, which can be then utilizing 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. 1 while its specific input-output relationship dictate the function nodes. Consequently, in order to learn such a factor graph from samples, one must only 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 softmax output layer and cross-entropy objective, which are known to reliably learn conditional distributions in complex environments [21]. 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 n 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 is presented in the following section.

III. APPLICATION: SYMBOL DETECTION

In this section we demonstrate how the concept of data-driven factor graphs can be applied for symbol detection in finite-memory channels. We begin with reviewing the BCJR algorithm, which is the model-based application of the sum-product algorithm for symbol detection, in Subsection III-A. Then, we show how the factor graph can be learned from training in Subsection III-B, and discuss the pros and cons of the resulting architecture in Subsection III-C.

A. Sum-Product for Symbol Detection

We next present the application of the sum-product method for symbol detection in finite-memory channels, also known as the BCJR algorithm [5]. Consider a stationary finite-memory channel, namely, a channel obeying the model in Example 1 in which the conditional distribution $P_{Y_k|\boldsymbol{S}_k}(y|\boldsymbol{s})$ does not depend on the time index $k \in \mathcal{N}$. Symbol detection refers to the task of recovering the i.i.d. transmitted symbols $\{X_k\}_{k=1}^n$, each uniformly distributed over the constel-

lation set $\mathcal{X}$, from a realization of the channel output $\boldsymbol{Y} = \boldsymbol{y}$. The detector which minimizes the symbol error probability is the MAP rule

$$\hat{X}_k = \arg\max_{x \in \mathcal{X}} P_{X_k|\boldsymbol{Y}}(x|\boldsymbol{y})$$

$$= \arg\max_{x \in \mathcal{X}} P_{X_k,\boldsymbol{Y}}(x, \boldsymbol{y}), \qquad k \in \mathcal{N}. \tag{10}$$

Using the formulation of the state vectors $\{\boldsymbol{S}_k\}$, the desired joint probability can be written as [22, Ch. 9.3]

$$P_{X_k,\boldsymbol{Y}}(x, \boldsymbol{y}) = \sum_{\boldsymbol{s} \in \mathcal{X}^L} P_{\boldsymbol{S}_{k-1}, \boldsymbol{S}_k, \boldsymbol{Y}}(\boldsymbol{s}, [x, (\boldsymbol{s})_1 \dots (\boldsymbol{s})_{L-1}]^T, \boldsymbol{y}). \tag{11}$$

The summands in (11) are the joint distributions evaluated recursively from the channel factor graph in Example 2. Thus, when the factor graph is known, the MAP rule (10) can be computed efficiently using the sum-product algorithm.

B. BCJRNet: Data-Driven MAP Recovery

Here, we show how the rationale presented in Section II yields a method for learning the factor graphs of finite-memory channels, using 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 $P_{Y|\boldsymbol{S}}(\cdot)$ is unknown, and only a set of labeled input-output pairs, denoted $\{x_i, y_i\}$ is available.

Since the channel memory is known, the structure of the factor graph is fixed to that depicted in Fig. 1. Consequently, in order to learn the factor graph, one must only adapt the function nodes $f_i(\cdot)$, as discussed in Subsection II-B. Based on (3)-(4), the function nodes are given by

$$f_i(y_i, \boldsymbol{s}_i, \boldsymbol{s}_{i-1}) = P_{Y|\boldsymbol{S}}(y_i|\boldsymbol{s}_i) P_{\boldsymbol{S}_i|\boldsymbol{S}_{i-1}}(\boldsymbol{s}_i|\boldsymbol{s}_{i-1})$$

$$\stackrel{(a)}{=} \begin{cases} \frac{1}{|\mathcal{X}|} P_{Y|\boldsymbol{S}}(y_i|\boldsymbol{s}_i) & (\boldsymbol{s}_i)_j = (\boldsymbol{s}_{i-1})_{j-1}, j = 2 \dots L \\ 0 & \text{otherwise,} \end{cases} \tag{12}$$

where (a) follows from the definition of the state vectors $\boldsymbol{S}_i$ and the channel stationarity. The formulation of the function nodes in (12) implies that they can be estimated by training an ML-based system to evaluate $P_{Y|\boldsymbol{S}}(\cdot)$ from which the corresponding function node value is obtained via (12). Once the factor graph representing the channel is learned, symbol recovery is carried out using the sum-product method detailed in Subsection III-A. The resulting receiver, referred
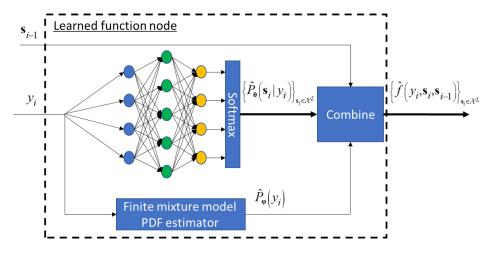
Fig. 3. Learned function node of BCJRNet illustration.

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 Section IV.

As noted in [13], since y_i is given and may take continuous values while $\boldsymbol{s}$, representing the label, takes discrete values, a natural approach to evaluate $P_{Y|\boldsymbol{S}}(y_i|\boldsymbol{s})$ for each $\boldsymbol{s} \in \mathcal{X}^L$ using ML tools is to estimate $P_{\boldsymbol{S}|Y}(\boldsymbol{s}|y_i)$, from which the desired $P_{Y|\boldsymbol{S}}(y_i|\boldsymbol{s})$ can be obtained using Bayes rule via

$$P_{Y|\boldsymbol{S}}(y_i|\boldsymbol{s}) = |\mathcal{X}|^L P_{\boldsymbol{S}|Y}(\boldsymbol{s}|y_i) P_Y(y_i). \tag{13}$$

In particular, BCJRNet utilizes two parametric models: one for evaluating the conditional $P_{\boldsymbol{S}|Y}(\boldsymbol{s}|y_i)$, and another for computing the marginal probability density function (PDF) $P_Y(y_i)$. A reliable parametric estimate of $P_{\boldsymbol{S}|Y}(\boldsymbol{s}|y_i)$, denoted $P_{\boldsymbol{\theta}}(\boldsymbol{s}|y_i)$, can be obtained for each $\boldsymbol{s} \in \mathcal{X}^L$ by training a relatively compact classification network with a softmax output layer. For example, in our numerical study in Section IV we use a three-layer network which can be trained with merely 10000 training samples. In order to estimate the marginal PDF of Y_i, we note that since it is given by a stochastic mapping of $\boldsymbol{S}_i$, its distribution can be approximated as a mixture model of $|\mathcal{X}|^L$ kernel functions [23]. Consequently, a parametric estimate of $P_Y(y_i)$, denoted $P_{\boldsymbol{\varphi}}(y_i)$, can be obtained from the training data using mixture density estimation via, e.g., expectation maximization (EM) [23, Ch. 2], or any other finite mixture model fitting method. The resulting structure in which the parameteric $P_{\boldsymbol{\theta}}(\boldsymbol{s}|y_i)$ and $P_{\boldsymbol{\varphi}}(y_i)$ are combined into a learned function node using (12)-(13), is illustrated in Fig. 3.

C. Discussion

Learned factor graph combine ML and signal processing algorithms, allowing the latter to be carried out in a data-driven manner. Compared to the conventional ML approach of utilizing a DNN to carry out the desired end-to-end mapping, DNNs are utilized here for the simpler task of uncovering the function nodes of a factor graphs. Consequently, networks with a relatively small number of parameters can be utilized, allowing the system to train quickly from small training sets. In particular, in our numerical study detailed in Section IV we used a three layer fully-connected network for learning the factor nodes in BCJRNet, allowing the system to adapt using merely 10000 training bits, e.g., not much more than a typical 4G preamble [24, Ch. 17]. For comparison, conventional DNN architectures used for carrying out the complete symbol detection task in finite-memory channels typically involve deep networks and require much larger training sets, e.g., [25], [26]. Furthermore, the ability to implement a data-driven receiver which is capable of adapting its mapping using a small data sets paves the way to the possibility of an on-line training receiver, utilizing the inherent structure of coded communications to train without additional overhead, see, e.g., [13] and [27]. We leave the exploration of such on-line trained symbol detectors for future work.

The proposed BCJRNet is one example of how learned factor graphs give rise to a data-driven implementation of model-based algorithms. We envision learned factor graphs to be a technique for realizing additional systems capable of learning to carry out important methods in communications, signal processing, and statistics, from relatively small labeled data sets. Furthermore, we believe that this concept unifies additional previously proposed systems combining ML and signal processing algorithms, such as the hybrid extended Kalman filter proposed in [28]. We leave this extension of our study for future work.

IV. NUMERICAL EVALUATIONS

In this section we numerically evaluate the performance of BCJRNet, which is obtained by learning the underlying factor graph of a finite-memory channel as detailed in Section III, and compare it to the conventional model-based BCJR algorithm. In the following, the DNN used for computing the conditional distributions in the learned function node in Fig. 3 consists of three fully-connected layers: a 1×100 layer, a 100×50 layer, and a 50×4 layer, with intermediate sigmoid and ReLU activations, respectively, and a softmax output layer. The network is trained using the Adam optimizer [29] with learning rate of 0.01. The finite mixture model estimator

approximates the distribution as a Gaussian mixture using EM-based fitting [23, Ch. 2]. The function node is learned from 10000 training samples, and is tested over 50000 Monte Carlo simulations. Due to the small number of training samples and the simple DNN architecture, only a few minutes are required to train BCJRNet on a standard CPU.

We simulate two finite-memory channels with memory $L = 2$: An intersymbol interference (ISI) channel with additive white Gaussian noise (AWGN), and a Poisson channel. To formulate the input-output relationships of these channels, we let $\{h_1(\gamma), h_2(\gamma)\}$ be L coefficients representing an exponentially decaying profile, given by $h_\tau(\gamma) \triangleq e^{-\gamma(\tau-1)}$ for $\gamma > 0$. For the ISI channel with AWGN, we consider i.i.d. binary phase shift keying inputs, i.e., $\mathcal{X} = \{-1, 1\}$, and the channel output Y_i is related to the input via

$$Y_i = \sqrt{\rho} \cdot \sum_{\tau=1}^{L} h_\tau(\gamma) X_{i-\tau+1} + W_i, \tag{14}$$

where W_i is a unit variance AWGN independent of X_i, and $\rho > 0$ represents the signal-to-noise ratio (SNR). For the Poisson channel, the input represents on-off keying, namely, $\mathcal{X} = \{0, 1\}$, and the statistical input-output relationship is

$$Y_i | \boldsymbol{X} \sim \mathbb{P}\left(\sqrt{\rho} \cdot \sum_{\tau=1}^{L} h_\tau(\gamma) X_{i-\tau+1} + 1 \right), \tag{15}$$

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

BCJRNet is trained for each SNR value ρ, and the symbol error rate (SER) values are averaged over 20 different channel coefficients $\{h_\tau(\gamma)\}$ with γ taking values in $[0.1, 2]$. In addition, we evaluate resiliency to inaccurate training by computing the SER when the function nodes in Fig. 3 are learned from samples taken from a channel with a noisy estimate of $\{h_\tau(\gamma)\}$. The estimation noise is randomized as an i.i.d. zero-mean Gaussian process with variance σ_e^2, where we use $\sigma_e^2 = 0.1$ for the Gaussian channel (14), and $\sigma_e^2 = 0.08$ for the Poisson channel (15). We refer to this scenario as *CSI uncertainty*. Under such uncertainty, the model-based BCJR computes the messages detailed in Example 2 using the noisy channel estimate, while BCJRNet is trained using samples taken from different realizations of the noisy $\{h_\tau(\gamma)\}$.

The numerically computed SER values are depicted in Figs. 4-5 for the Gaussian channel and the Poisson channel, respectively. We first note that for both channels, the performance of BCJRNet approaches that of the channel-model-based algorithm, and that the curves effectively coincide for most simulated SNR values. A small gap is noted for the Poisson channel at high
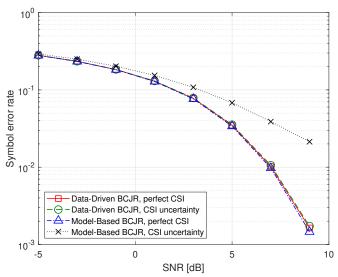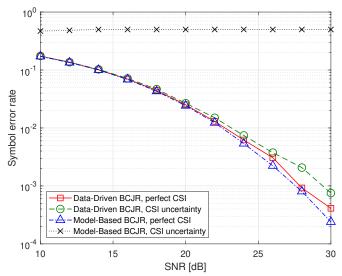
Fig. 4. SER versus SNR, ISI channel with AWGN.

Fig. 5. SER versus SNR, Poisson channel.

SNRs, which stems from the model mismatch induced by approximating the marginal PDF of Y_i as a Gaussian mixture in the parameteric estimate $P_{\boldsymbol{\varphi}}(y_i)$.

We also observe in Figs. 4-5 that the data-driven implementation of the BCJR algorithm is substantially more robust to CSI uncertainty compared to the channel-model-based detector. For example, in Fig. 4 we note that BCJRNet trained with a variety of different channel conditions achieves roughly the same performance as it does when trained and tested using samples from the same statistical model. However, the channel-model-based BCJR algorithm achieves significantly degraded SER performance due to imperfect CSI.

The results presented in this section demonstrate that an ML-based receiver can be trained to carry out accurate and robust MAP symbol detection using a relatively small amount of labeled samples. This is achieved by learning the function nodes of the underlying factor graphs, and carrying out the sum-product recursions over the learned graph.

V. Conclusions

We have proposed the concept of data-driven factor graphs methods, which allow a broad family of model-based algorithms to be implemented in a data-driven fashion. Our approach utilizes ML tools to adapt the function nodes of the underlying statistical model, allowing to learn the factor graph data, using which the desired scheme can be implemented by message passing over the learned graph. We applied the proposed strategy to learn the factor graph of stationary finite-memory communication channels, yielding BCJRNet, which is a data-driven implementation of the BCJR MAP symbol detector. Our numerical results demonstrate that BCJRNet learns to accurately implement the BCJR algorithm from a relatively small training set, while exhibiting improved robustness to CSI uncertainty compared to the model-based BCJR detector.

References

[1] H.-A. Loeliger, J. Dauwels, J. Hu, S. Korl, L. Ping, and F. R. Kschischang, "The factor graph approach to model-based signal processing," *Proceedings of the IEEE*, vol. 95, no. 6, pp. 1295–1322, 2007.

[2] G. D. Forney, "Codes on graphs: Normal realizations," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 520–548, 2001.

[3] 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.

[4] H.-A. Loeliger, "An introduction to factor graphs," *IEEE Signal Process. Mag.*, vol. 21, no. 1, pp. 28–41, 2004.

[5] 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.

[6] H.-A. Loeliger, "Least squares and Kalman filtering on Forney graphs," in *Codes, Graphs, and Systems*. Springer, 2002, pp. 113–135.

[7] M. I. Jordan, "Graphical models," *Statistical Science*, vol. 19, no. 1, pp. 140–155, 2004.

[8] G. Colavolpe and G. Germi, "On the application of factor graphs and the sum-product algorithm to ISI channels," *IEEE Trans. Commun.*, vol. 53, no. 5, pp. 818–825, 2005.

[9] R. J. Drost and A. C. Singer, "Factor-graph algorithms for equalization," *IEEE Trans. Signal Process.*, vol. 55, no. 5, pp. 2052–2065, 2007.

[10] P. Som, T. Datta, N. Srinidhi, A. Chockalingam, and B. S. Rajan, "Low-complexity detection in large-dimension mimo-isi channels using graphical models," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 8, pp. 1497–1511, 2011.

[11] H. Niu, M. Shen, J. A. Ritcey, and H. Liu, "A factor graph approach to iterative channel estimation and ldpc decoding over fading channels," *IEEE Trans. Wireless Commun.*, vol. 4, no. 4, pp. 1345–1350, 2005.

[12] F. Lehmann, "A factor graph approach to iterative channel estimation, detection, and decoding for two-path successive relay networks," *IEEE Trans. Wireless Commun.*, vol. 15, no. 8, pp. 5414–5429, 2016.

[13] N. Shlezinger, N. Farsad, Y. C. Eldar, and A. J. Goldsmith, "ViterbiNet: A deep learning based Viterbi algorithm for symbol detection," *IEEE Trans. Wireless Commun., early access*, 2019.

[14] N. Shlezinger, R. Fu, and Y. C. Eldar, "DeepSIC: Deep soft interference cancellation for multiuser MIMO detection," *Submitted to IEEE J. Sel. A. Inform. Theory*, 2019.

[15] G. D. Forney, "The Viterbi algorithm," *Proceedings of the IEEE*, vol. 61, no. 3, pp. 268–278, 1973.

[16] 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.

[17] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, p. 436, 2015.

[18] Y. Weiss and W. T. Freeman, "On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 736–744, 2001.

[19] J. Pearl, "Fusion, propagation, and structuring in belief networks," *Artificial intelligence*, vol. 29, no. 3, pp. 241–288, 1986.

[20] N. Wiberg, H.-A. Loeliger, and R. Kotter, "Codes and iterative decoding on general graphs," *European Transactions on telecommunications*, vol. 6, no. 5, pp. 513–525, 1995.

[21] Y. Bengio, "Learning deep architectures for AI," *Foundations and trends® in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.

[22] J. M. Cioffi, "Sequence detection," *EE379B Course notes chapter 9. Stanford University*, 2008.

[23] G. McLachlan and D. Peel, *Finite mixture models*. John Wiley & Sons, 2004.

[24] E. Dahlman, S. Parkvall, J. Skold, and P. Beming, *3G evolution: HSPA and LTE for mobile broadband*. Academic press, 2010.

[25] 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.

[26] Y. Liao, N. Farsad, N. Shlezinger, Y. C. Eldar, and A. J. Goldsmith, "Deep neural network symbol detection for millimeter wave communications," *arXiv preprint arXiv:1907.11294*, 2019.

[27] L. Lugosch and W. J. Gross, "Learning from the syndrome," in *2018 52nd Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2018, pp. 594–598.

[28] B. Laufer-Goldshtein, R. Talmon, and S. Gannot, "A hybrid approach for speaker tracking based on tdoa and data-driven models," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 26, no. 4, pp. 725–735, 2018.

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