

Blind Channel Estimation and Joint Symbol Detection with Data-Driven Factor Graphs

Luca Schmid, Tomer Raviv, Nir Shlezinger, *Senior Member, IEEE*, and Laurent Schmalen, *Fellow, IEEE*

Abstract—We investigate the application of the factor graph framework for blind joint channel estimation and symbol detection on time-variant linear inter-symbol interference channels. In particular, we consider the expectation maximization (EM) algorithm for maximum likelihood estimation, which typically suffers from high complexity as it requires the computation of the symbol-wise posterior distributions in every iteration. We address this issue by efficiently approximating the posteriors using the belief propagation (BP) algorithm on a suitable factor graph. By interweaving the iterations of BP and EM, the detection complexity can be further reduced to a single BP iteration per EM step. In addition, we propose a data-driven version of our algorithm that introduces momentum in the BP updates and learns a suitable EM parameter update schedule, thereby significantly improving the performance-complexity tradeoff with a few offline training samples. Our numerical experiments demonstrate the excellent performance of the proposed blind detector and show that it even outperforms coherent BP detection in high signal-to-noise scenarios.

Index Terms—Factor graphs, expectation maximization, belief propagation, joint detection, model-based machine learning, 6G.

I. INTRODUCTION

WE study the fundamental problem of symbol detection in digital communications, and particularly the inference of transmitted symbols at the receiver impaired by a linear channel with memory and additive white Gaussian noise (AWGN). While there exists a plethora of detection algorithms ranging from low-complexity linear detectors [2] to high-performance solutions that are optimal concerning the symbol error probability [3], the majority of works in the literature assume the availability of channel state information (CSI) at the receiver. In most practical systems, this knowledge of the underlying channel is gained by sending predesigned *training* or *pilot* symbols from which the receiver can estimate the current channel state. The transmission of such training symbols reduces the data rate which leads, especially for rapidly time-variant channels, to a substantial reduction of

the communication throughput [4]. For example, in many emerging communication scenarios with short block length transmission, like low-latency communications in 5G/6G or Internet of Things systems, the training signals occupy a significant part of the short transmission blocks. Hence, there is considerable interest in efficient blind detection schemes that do not rely on the availability of CSI.

State-of-the-art blind detectors like the constant modulus algorithm (CMA) [5] or variational methods [6], [7] are still consistently outperformed by coherent detection schemes, i.e., detectors with fully available CSI. On the other hand, there exist blind estimation and detection algorithms based on the maximum likelihood (ML) criterion [8], [9], [4] that are competitive with coherent detectors in terms of detection performance. For instance, a joint estimation and detection approach based on the Baum-Welch algorithm is proposed in [9], which is an earlier version of the expectation maximization (EM) algorithm. The EM algorithm aims to iteratively converge to a local ML solution by alternately performing likelihood-increasing detection and estimation steps that mutually depend on each other [10]. However, such iterative estimation and detection schemes typically induce a notable computational burden, e.g., the method proposed in [9] requires applying a forward-backward algorithm for the maximum a posteriori (MAP) detection in every iteration.

An alternative approach to design receivers that can operate without CSI is based on machine learning. The ability of deep neural networks (DNNs) to efficiently approximate high-dimensional non-linear mappings has led to a new class of so-called “deep receivers” that do not (only) rely on the underlying channel model but operate in a data-driven fashion [11]. For instance, the authors in [12] use a Turbo-autoencoder neural network to learn a holistic transmission scheme with superimposed pilots for joint detection and decoding, enabling efficient short-packet communications. On the other hand, the approach in [13] is completely channel-agnostic, i.e., the communication channel is assumed to be completely unknown and a recurrent neural network is trained to perform the detection task. In comparison to “classical” receivers that use training symbols to estimate certain parameters of the underlying channel model, a deep receiver instead uses the available data to directly adapt the DNN-based detector to the channel.

To reduce the required number of training symbols, model-based deep learning [14], [15] combines the advantages of both methods—model-based domain knowledge as well as data-driven optimization—by either integrating knowledge of the underlying model in the DNN design, or vice versa,

This work has received funding in part from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101001899), in part from the German Federal Ministry of Education and Research (BMBF) within the project Open6GHub (grant agreement 16KISK010), and in part from the Israeli Ministry of Science and Technology.

This work has been presented at the Asilomar Conference on Signals, Systems, and Computers 2024, Pacific Grove, CA, USA [1].

L. Schmid and L. Schmalen are with the Communications Engineering Lab (CEL), Karlsruhe Institute of Technology (KIT), Hertzstr. 16, 76187 Karlsruhe, Germany (e-mail: first.last@kit.edu). T. Raviv and N. Shlezinger are with the School of ECE, Ben-Gurion University of the Negev, Be’er-Sheva, Israel (e-mail: tomerraviv95@gmail.com, nirshl@bgu.ac.il).

by replacing certain building blocks of an otherwise model-based algorithm with DNNs [16]. Thereby, an algorithm can be significantly improved without sacrificing the benefits of being model-based. In [17], DNNs are integrated into the well-known Viterbi algorithm to replace all CSI-dependent computations. Another example is BCJRNet [18], which builds upon the factor graph model of the BCJR algorithm [3] and replaces the factor nodes in the graph with a mapping that is learned from a small training set. While these data-aided receivers do not require CSI, they still rely on labeled data obtained from, e.g., pilots. Moreover, their reliance on DNNs makes their online adaptation to rapidly time-variant channels challenging in terms of efficient learning [19] and data accumulation [20], see also [15], [21].

In this work we propose a blind receiver that is particularly designed to efficiently operate in rapidly time-variant linear inter-symbol interference (ISI) channels without requiring pilots. We design our receiver in two stages. We first propose a novel joint channel estimation and symbol detection scheme based on a combination of the EM algorithm and belief propagation (BP) on factor graphs that alleviates the necessity of pilots of coherent detectors and the extensive computational complexity of other non-coherent methods. Based on the so-called Ungerboeck observation model [22], we use a dedicated factor graph on which the BP algorithm efficiently implements symbol detection with particularly low complexity, namely, with linear complexity in the number of symbols *and* in the memory of the channel [23]. Dealing with continuous variables like channel parameters in factor graphs is typically challenging because the application of BP usually leads to intractable expressions that must be approximated, e.g., using Gaussian approximation or quantization methods [24]. To overcome this issue, Eckford proposed the integration of the EM algorithm into the factor graph framework and showed its appealing property of breaking cycles in the graph [25]. The idea was later formalized as a local message passing algorithm in [26] and applied to practical communication systems such as [27]. Inspired by this idea, we integrate our BP-based detector into the EM algorithm to form a joint estimation and detection framework. In contrast to [26] and [27], we do not use the EM algorithm as a local message update rule, but instead perform global EM parameter updates, into which the BP algorithm is integrated. This turns out to be more efficient in terms of complexity, since we can derive closed-form update equations for each individual EM parameter, respectively. Moreover, to avoid running the full detection algorithm after each EM step, we leverage the iterative nature of factor graph-based message passing and interweave the EM estimation steps with the BP message passing iterations. Thereby, the total number of required BP iterations for blind joint estimation and detection does not substantially increase compared to coherent detection. This leads to a significant complexity reduction compared to other approaches in literature, e.g., [27], where each EM iteration requires multiple BP iterations.

A known weakness of the EM algorithm is its sensitivity to initialization [10], [28]. While most works in the literature circumvent this problem by using pilot symbols to obtain an initial channel estimate, we address this challenge for the

blind estimation scenario and propose a novel initialization method that leverages a lightweight version of the variational autoencoder (VAE).

Next, we follow the concept of model-based machine learning and extend our blind factor graph-based estimation and detection scheme into a data-driven receiver. In line with other works in literature, like [18], we call the optimized version of the factor graph-based EMBP algorithm a *data-driven factor graph*. However note that, compared to the approaches in [17] and [18], our method is already adaptive with respect to the channel without the requirement of training symbols, thus alleviating the possibly computationally complex frequent re-training that channel agnostic deep receivers typically require to be adaptive [15]. Therefore, we can avoid online training and only need to perform one generic offline training prior to transmission in order to optimize the detection performance as well as the computational efficiency of our proposed algorithm for a broad range of channel states. To this end, we unfold the iterative BP algorithm and replace the BP message updates with a convex combination of old and new messages. While introducing only one trainable scalar parameter per iteration, this modification significantly improves the convergence behavior and thus the detection performance. Moreover, we propose a simple, yet effective method to learn a suitable update schedule for the EM algorithm and show that the optimized schedule can substantially reduce the number of necessary EM steps by allowing parallel parameter updates while preserving the stability of the EM algorithm.

Our main contributions are summarized as follows:

- We derive a novel algorithm for joint channel estimation and detection on linear ISI channels that offers a high precision at low computational complexity based on the integration of BP into the EM algorithm.
- We significantly improve the performance-complexity tradeoff by extending our algorithm into a trainable machine learning model, that leverages offline data to learn an effective update schedule for the EM algorithm and to improve the BP-based detection performance via learned weighting factors.
- We extensively evaluate the estimation and detection performance of the proposed EMBP scheme in numerical simulations with various transmission scenarios. In particular, we analyze the sensitivity of the EM algorithm with respect to different initialization methods and compare the detection performance with other coherent and blind detection schemes from literature, including both model-based and data-driven algorithms. As a remarkable feature, we show that the EMBP algorithm outperforms a comparable coherent BP-based detector by inherently estimating a “surrogate” channel which is better suited for the suboptimal BP algorithm. Our simulations also show that the low-complexity blind EMBP scheme performs comparably to non-blind pilot-based MAP detectors. Furthermore, we show that the EMBP algorithm can even outperform the pilot-based detectors when compared in the context of a forward error correction (FEC) scheme, by leveraging the saved pilot overhead for an enhanced error correction capability.

The remainder of this paper is organized as follows: In Section II we present the system model, briefly review factor graph-based symbol detection, and formulate the challenges and evaluation metrics of blind symbol detection. Section III proposes the EMBP algorithm as a blind detection algorithm as well as its data-driven version EMBP*. Section IV presents a comprehensive numerical evaluation of the proposed algorithms and compares the estimation and detection performance to other established methods. Section V concludes the paper.

Notation

Throughout the paper, we use upper case bold letters to denote matrices \mathbf{X} with entries $X_{m,n}$ at row m and column n . Lower case bold letters are used for column vectors \mathbf{x} . The i th element of \mathbf{x} is written as x_i . $\|\cdot\|$ denotes the Euclidean norm, $(\cdot)^T$ is the transpose of a matrix or vector and $(\cdot)^H$ is the conjugate transpose (Hermitian) operator. For a complex number $c \in \mathbb{C}$, $\text{Re}\{c\}$ ($\text{Im}\{c\}$) denotes its real (imaginary) part and c^* is its complex conjugate. The all-zeros and all-ones vectors are written as $\mathbf{0}$ and $\mathbf{1}$, respectively, and \mathbf{e}_i denotes a vector that has a 1 at position i and zeros everywhere else. The ceiling and flooring operation of a floating point number x are indicated by $\lceil x \rceil$ and $\lfloor x \rfloor$, respectively. $(i \bmod j)$ denotes the modulo operation, i.e., the remainder of Euclidean division of i by j with $i, j \in \mathbb{N}$. The probability density function of a continuous random variable y is denoted by $p_y(y)$ or $p(y)$ and the probability mass function of a discrete random variable x is $P_x(x)$ or $P(x)$. To keep the notation simple, we do not use a special notation for random variables since it is always clear from the context. The Gaussian distribution, characterized by its mean μ and variance σ^2 , is written as $\mathcal{N}(\mu, \sigma^2)$. The expected value of a random variable x is denoted by $\mathbb{E}_x\{x\}$ and the mutual information between x and y is $I(x; y)$. We use calligraphic letters to denote a set \mathcal{X} of cardinality $|\mathcal{X}|$.

II. SYSTEM MODEL AND PRELIMINARIES

A. System Model

We consider the transmission over a time-variant channel in the digital baseband which is impaired by linear ISI and AWGN [2]. The dynamics of the communication channel are assumed to be block-fading, i.e., the channel coefficients remain constant for a block of N consecutively transmitted symbols and change to an independent realization in the next block [29]. Each transmission block contains an information sequence $\mathbf{c} \in \mathcal{M}^N$, where each symbol c_n is sampled independently and uniformly from a constellation $\mathcal{M} = \{c_i \in \mathbb{C}, i = 1, \dots, M\}$. The corresponding bit pattern of length $m := \log_2(M)$ is denoted by $\mathbf{b}(c_n) = (b_1(c_n), \dots, b_m(c_n))$. The receiver observes the sequence

$$\mathbf{y} = \underbrace{\begin{pmatrix} h_0 & & & \\ \vdots & h_0 & \mathbf{0} & \\ h_L & \vdots & \ddots & \\ & h_L & & h_0 \\ \mathbf{0} & & \ddots & \vdots \\ & & & h_L \end{pmatrix}}_{=: \mathbf{H}} \underbrace{\begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_N \end{pmatrix}}_{=: \mathbf{c}} + \underbrace{\begin{pmatrix} w_1 \\ w_2 \\ \vdots \\ w_{N+L} \end{pmatrix}}_{=: \mathbf{w}},$$

where $\mathbf{h} = (h_0, \dots, h_L)^T \in \mathbb{C}^{L+1}$ describes the impulse response of the linear ISI channel of length $L+1$ and $w_n \sim \mathcal{CN}(0, \sigma^2)$ are independent noise samples from a complex circular Gaussian distribution. For each transmission block, $\tilde{\mathbf{h}} := \alpha \cdot \mathbf{h}$ is generated by $\tilde{\mathbf{h}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_{L+1})$ and $\alpha = 1/\|\tilde{\mathbf{h}}\|$. The channel is thus fully characterized by the parameter vector $\boldsymbol{\theta} := (h_0, \dots, h_L, \sigma^2)^T$ of length $L+2$. We define the signal-to-noise ratio (SNR) at the receiver as

$$\text{snr} := \frac{\|\mathbf{h}\|^2 \cdot \mathbb{E}_{\mathbf{h}}\{\|\mathbf{c}\|^2\}}{\mathbb{E}_{\mathbf{w}}\{\|\mathbf{w}\|^2\}} = \frac{\sum_{c \in \mathcal{M}} |c|^2}{M\sigma^2}.$$

B. Factor Graph-based Symbol Detection

1) *Symbol Detection*: For each block, our goal is to infer the information sequence \mathbf{c} from the observation \mathbf{y} at the receiver. Since the symbols and the channel parameters are independently sampled, this inference task can be approached separately for each block. In the context of Bayesian inference, we are interested in the symbol-wise a posteriori probabilities (APPs)

$$P(c_n = c|\mathbf{y}) = \sum_{\mathbf{c} \in \mathcal{M}^N, c_n = c} P(\mathbf{c}|\mathbf{y}), \quad n = 1, \dots, N. \quad (1)$$

2) *Factor Graph Representation*: Message passing on factor graphs [30] is a powerful tool to derive efficient algorithms for marginal inference such as the problem in (1). In order to apply the factor graph framework to the considered problem, we first need to find a suitable factorization of the joint APP distribution $P(\mathbf{c}|\mathbf{y})$. Applying Bayes' theorem, we obtain

$$P(\mathbf{c}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{c})P(\mathbf{c})}{p(\mathbf{y})} = \frac{p(\mathbf{y}|\mathbf{c})}{M^N p(\mathbf{y})} \propto p(\mathbf{y}|\mathbf{c}). \quad (2)$$

The proportionality \propto in (2) denotes that two terms are only differing in a factor independent of \mathbf{c} . The likelihood

$$p(\mathbf{y}|\mathbf{c}) = \frac{1}{(\pi\sigma^2)^N} \exp\left(-\frac{\|\mathbf{y} - \mathbf{H}\mathbf{c}\|^2}{\sigma^2}\right)$$

can be written as

$$p(\mathbf{y}|\mathbf{c}) \propto \exp\left(\frac{2\text{Re}\{\mathbf{c}^H \mathbf{H}^H \mathbf{y}\} - \mathbf{c}^H \mathbf{H}^H \mathbf{H} \mathbf{c}}{\sigma^2}\right). \quad (3)$$

Based on the observation model by Ungerboeck [22], we substitute

$$\mathbf{x} := \mathbf{H}^H \mathbf{y}, \quad \mathbf{G} := \mathbf{H}^H \mathbf{H},$$

which can be interpreted as the matched filtered versions of the observation and the channel matrix, respectively. By using

$$\begin{aligned} \mathbf{c}^H \mathbf{x} &= \sum_{n=1}^N x_n c_n^*, \\ \mathbf{c}^H \mathbf{G} \mathbf{c} &= \sum_{n=1}^N G_{n,n} |c_n|^2 - \sum_{n=1}^N \sum_{\substack{m=1 \\ m \neq n}}^N \text{Re}\{G_{n,m} c_m c_n^*\}, \end{aligned}$$

in (3), we obtain a factorized version of the likelihood

$$p(\mathbf{y}|\mathbf{c}) \propto \prod_{n=1}^N \left[e^{F_n(c_n)} \prod_{m=1}^{n-1} e^{I_{n,m}(c_n, c_m)} \right], \quad (4)$$




Fig. 1. Factor graph representation of (4) for a channel with memory $L = 2$.

where we introduced the local functions

$$F_n(c_n) := \frac{1}{\sigma^2} \text{Re}\{2x_n c_n^* - G_{n,n}|c_n|^2\} \quad (5)$$

$$I_{n,m}(c_n, c_m) := -\frac{2}{\sigma^2} \text{Re}\{G_{n,m} c_m c_n^*\}. \quad (6)$$

A factor graph visualizes a factorization, such as (4), using a bipartite graph [30]. Every variable c_n is represented by a circular node, a so-called variable node. Factor nodes represent the local functions of the factorization and are visualized by square nodes. A factor node is connected with a variable node via an undirected edge if and only if the corresponding factor is a function of this variable. For the factorization in (4), this results in the factor graph illustrated in Fig. 1. Due to the band structure of \mathbf{G} , many of the factor nodes $I_{n,m}(c_n, c_m)$ represent the identity function and can thus be omitted, leading to a sparse graph representation. Note that there exist different factorizations of $p(\mathbf{y}|\mathbf{c})$, which consequently lead to distinct graphs [31]. However, the specific structure of the factor graph in Fig. 1, originally proposed by Colavolpe et al. in [23], is particularly favorable in terms of computational complexity in the context of marginal inference based on the BP algorithm, which we discuss in the following.

3) *Belief Propagation*: The sum-product algorithm, also known as BP, is a message passing algorithm that operates in a graph and attempts to determine the marginalization towards each variable in the factor graph, respectively [30]. To this end, messages are propagated within the factor graph along its edges. In contrast to tree-structured graphs, where the messages can travel through the entire graph in a single forward and backward path, message passing in graphs with cycles yields an iterative algorithm, i.e., the messages keep traveling within the cyclic graph until some stopping criterion is fulfilled, e.g., a maximum number of message passing iterations is reached.

For the graph in Fig. 1, we denote the message¹ sent in iteration $t \in \mathbb{N}$ from a variable node c_n to a factor node $I_{n,m}(c_n, c_m)$ with $n > m$ as $\mu_{n,m}^{(t)}(c_n)$. The message traveling on the same edge but in the opposite direction is denoted as $\nu_{n,m}^{(t)}(c_m)$. All messages are initialized with an unbiased state, e.g., $\mu_{n,m}^{(0)}(c_n) = -\log M$. Then, T message passing iterations are performed. Each iteration t consists of the parallel update

of all messages from variable to factor nodes

$$\mu_{n,m}^{(t)}(c_n) = F_n(c_n) + \sum_{\substack{m'=1 \\ m' \neq m}}^{n-1} \nu_{n,m'}^{(t-1)}(c_n) + \sum_{n'=n+1}^N \nu_{n',n}^{(t-1)}(c_n), \quad (7a)$$

$$\mu_{n,m}^{(t)}(c_m) = F_m(c_m) + \sum_{m'=1}^{m-1} \nu_{m,m'}^{(t-1)}(c_m) + \sum_{\substack{n'=m+1 \\ n' \neq n}}^N \nu_{n',m}^{(t-1)}(c_m), \quad (7b)$$

followed by the parallel update of all messages from factor to variable nodes

$$\nu_{n,m}^{(t)}(c_n) = \log \sum_{c_m} \exp\left(I_{n,m}(c_n, c_m) + \mu_{n,m}^{(t)}(c_m)\right), \quad (8a)$$

$$\nu_{n,m}^{(t)}(c_m) = \log \sum_{c_n} \exp\left(I_{n,m}(c_n, c_m) + \mu_{n,m}^{(t)}(c_n)\right). \quad (8b)$$

The final BP result is obtained after the last message passing iteration $t = T$ by combining all incident messages at a variable node to the so-called beliefs via

$$b_n(c_n) = \gamma_n \exp\left(F_n(c_n) + \sum_{m'=1}^{n-1} \nu_{n,m'}^{(T)}(c_n) + \sum_{n'=n+1}^N \nu_{n',n}^{(T)}(c_n)\right), \quad (9)$$

where the $\gamma_n \in \mathbb{R}$ are chosen such that $b_n(c_n)$ are normalized probability distributions. The beliefs are an approximation of the symbol-wise APP distributions $P(c_n|\mathbf{y}) \approx b_n(c_n)$ and the quality of this approximation strongly depends on the number of iterations T as well as on the nature of the underlying factor graph. For the specific graph in Fig. 1, the approximation quality thus varies with the block length N , the channel parameters $\boldsymbol{\theta}$ and with the observation \mathbf{y} .

A major benefit of BP on a suitable factor graph is its ability to approximate the APP with low computational complexity. The asymptotic complexity of the described algorithm is $\mathcal{O}(TNLM^2)$ [23]. In comparison, exact inference implemented by the BCJR algorithm [3] scales with $\mathcal{O}(NM^{L+1})$, i.e., the complexity grows exponentially with the memory of the channel L instead of linearly.

C. Problem Formulation

We consider symbol detection without CSI, i.e., the parameters $\boldsymbol{\theta}$ are unknown to both the transmitter and the receiver. An unsupervised estimate of the parameters can be found by maximizing the likelihood of the observation \mathbf{y} with respect to the channel parameters

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} p(\mathbf{y}|\boldsymbol{\theta}), \quad (10)$$

where we have included the unknown parameters $\boldsymbol{\theta}$ as latent variables in the model. Once $\hat{\boldsymbol{\theta}}$ is estimated, it can be used for CSI-based symbol detection. However, the exact ML solution in (10) is infeasible for reasonable block lengths $N \gg 1$. Nonetheless, leveraging intermediate results from symbol detection for the estimation task can help to reduce the complexity significantly. This motivates our goal to design

¹For reasons of numerical stability and practicability, we describe the messages and their updates in the logarithmic domain.

²Note the distinction to the message $\mu_{n,m}^{(t)}(c_m)$ which originates from variable node c_m , indicated by the function argument.

an efficient algorithm for *joint* channel estimation and symbol detection. In this context, we are facing two major challenges:

- C1 To achieve a good detection performance that is comparable to a coherent detection with available CSI.
- C2 To keep the computational complexity low such that it can be implemented on a receiver with sufficiently low power consumption and/or delay.

D. Evaluation Metrics

Suitable objective functions for *joint* estimation and detection algorithms can either evaluate the detection or the estimation performance. The quality of an estimated channel impulse response $\hat{\mathbf{h}}$ can be quantified by the squared error

$$\|\hat{\mathbf{h}} - \mathbf{h}\|^2 = \sum_{\ell=0}^L |\hat{h}_\ell - h_\ell|^2$$

to the ground truth \mathbf{h} . For statistics over multiple channels and their respective estimations, we will consider the median and the mean squared error (MSE).

As a figure of merit for the detection performance, we consider the maximum achievable information rate between the channel input and the detector output. Many practical transmission systems use bit-interleaved coded modulation (BICM), which decouples the symbol detection from a binary soft-decision FEC [32]. In a BICM system, the symbol detector soft output $P(c_n|\mathbf{y}, \boldsymbol{\theta})$ is converted by a bit-metric decoder (BMD) to binary soft information

$$P(b_i(c_n) = b|\mathbf{y}, \boldsymbol{\theta}) = \sum_{c \in \mathcal{M}_i^{(b)}} P(c_n = c|\mathbf{y}, \boldsymbol{\theta}), \quad b \in \{0, 1\}$$

with $\mathcal{M}_i^{(b)} := \{c \in \mathcal{M} : b_i(c) = b\}$. The resulting bit-wise APPs are typically expressed in log-likelihood ratios (LLRs)

$$L_{n,i}(\mathbf{y}) := \log \left(\frac{P(b_i(c_n) = 0|\mathbf{y}, \boldsymbol{\theta})}{P(b_i(c_n) = 1|\mathbf{y}, \boldsymbol{\theta})} \right).$$

After interleaving, the LLRs are fed to a bit-wise soft-decision FEC. By interpreting the BMD as a mismatched detector, the bitwise mutual information (BMI) is an achievable information rate for BICM [32]. The calculation of the BMI, detailed in [33], considers the BMD by assuming m parallel sub-channels transmitting on a binary basis instead of one symbol-based channel. Assuming independent transmit bits, the BMI is defined as the sum of mutual informations³ $I(b_i; y)$ of m unconditional bit-wise channel transmissions:

$$\text{BMI} := \sum_{i=1}^m I(b_i(c); y).$$

By a sample mean estimate over D labeled data tuples, a feasible approximation is given by [33]

$$\text{BMI} \approx \log_2(M) - \frac{1}{DN} \sum_{i=1}^m \sum_{n=1}^N \sum_{d=1}^D \log_2 \left(\exp \left(-(-1)^{b_{n,i}(c_n^{(d)})} L_{n,i}(\mathbf{y}^{(d)}) \right) + 1 \right),$$

where $c_n^{(d)}$ and $\mathbf{y}^{(d)}$ are from a labeled data set \mathcal{D} , which will be formally defined later in (16) in Sec. III-D.

³The mutual information is a measure between two random variables. We avoid a distinct notation for random variables as it is clear from the context.

III. FACTOR GRAPH-BASED JOINT ESTIMATION AND DETECTION USING EM AND BP

We propose a joint estimation and detection scheme for the problem formulated in Sec. II-C. First, we derive closed-form update expressions for the EM-based parameter estimation in Sec. III-A and show how the BP algorithm can be applied to effectively reduce the computational complexity. As this formulation on its own may not successfully cope with C1-C2 in short coherence durations, we empower our method in two aspects: (i) we incorporate dedicated initialization methods, as detailed in Sec. III-C; and (ii) we optimize the algorithm by introducing a small set of key parameters within the algorithm that are tuned in a data-driven fashion as a model-based machine learning model in offline training [34], as detailed in Sec. III-D.

A. Expectation Maximization

The EM algorithm is an established technique for finding ML parameter estimates in probabilistic models that contain latent variables [35]. The basic idea is to maximize a simple lower bound of the likelihood in (10) in a 2-stage iterative algorithm. To this end, the log-likelihood function is decomposed into [10]

$$\begin{aligned} \log p(\mathbf{y}|\boldsymbol{\theta}) &= \log p(\mathbf{c}, \mathbf{y}|\boldsymbol{\theta}) - \log P(\mathbf{c}|\mathbf{y}, \boldsymbol{\theta}) \\ &= \underbrace{\sum_{\mathbf{c}} Q(\mathbf{c}) \log \left(\frac{p(\mathbf{c}, \mathbf{y}|\boldsymbol{\theta})}{Q(\mathbf{c})} \right)}_{=: \mathcal{L}(Q, \boldsymbol{\theta})} - \underbrace{\sum_{\mathbf{c}} Q(\mathbf{c}) \log \left(\frac{P(\mathbf{c}|\mathbf{y}, \boldsymbol{\theta})}{Q(\mathbf{c})} \right)}_{=: D_{\text{KL}}(Q\|P)}, \end{aligned} \quad (11)$$

where $Q(\mathbf{c})$ can be any normalized trial distribution of \mathbf{c} . Jointly optimizing the log-likelihood in (11) with respect to $\boldsymbol{\theta}$ and *all* possible $Q(\mathbf{c})$ is typically infeasible. Due to the non-negativity of the Kullback-Leibler (KL) divergence $D_{\text{KL}}(Q\|P) \geq 0$, the functional $\mathcal{L}(Q, \boldsymbol{\theta})$ is a so-called evidence lower bound (ELBO), satisfying

$$\log p(\mathbf{y}|\boldsymbol{\theta}) \geq \mathcal{L}(Q, \boldsymbol{\theta}).$$

The EM algorithm maximizes the ELBO in multiple consecutive steps. Starting with an initial guess for the parameters $\hat{\boldsymbol{\theta}}^{(0)}$, the EM algorithm iteratively performs

- 1) **E-step:** Fix $\hat{\boldsymbol{\theta}}^{(t-1)}$ and maximize $\mathcal{L}(Q, \boldsymbol{\theta})$ w.r.t. Q :

$$\begin{aligned} Q &= \arg \max_Q \mathcal{L}(Q, \boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(t-1)}} \\ &= P(\mathbf{c}|\mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}). \end{aligned} \quad (12)$$

- 2) **M-step:** Fix Q and maximize $\mathcal{L}(Q, \boldsymbol{\theta})$ w.r.t. $\boldsymbol{\theta}$:

$$\begin{aligned} \hat{\boldsymbol{\theta}}^{(t)} &= \arg \max_{\boldsymbol{\theta}} \mathcal{L}(Q, \boldsymbol{\theta}) \Big|_{Q=P(\mathbf{c}|\mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})} \\ &= \arg \max_{\boldsymbol{\theta}} \mathbb{E}_{P(\mathbf{c}|\mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})} \{ \log p(\mathbf{c}, \mathbf{y}|\boldsymbol{\theta}) \} \\ &=: \arg \max_{\boldsymbol{\theta}} \tilde{Q}(\boldsymbol{\theta} | \hat{\boldsymbol{\theta}}^{(t-1)}). \end{aligned} \quad (13)$$

Despite the fact that both the E-step and the M-step only maximize a lower bound of the likelihood, respectively, the

$$\hat{\sigma}^{2(t)} = \frac{1}{N} \left[B - \sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \left(C(c_n, \mathbf{h}^{(t-1)}) - \sum_{m < n} \sum_{c_m} P(c_m | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) D(c_n, c_m, \mathbf{h}^{(t-1)}) \right) \right], \quad (14)$$

$$\begin{aligned} \hat{h}_\ell^{(t)} = & \left(\sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) y_{n+\ell} c_n^* - \sum_{\substack{k=0 \\ k \neq \ell}}^L \sum_{c_{n-|\ell-k|}} P(c_{n-|\ell-k|} | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) h_k^{(t-1)} \right. \\ & \left. \left(\text{Re}\{c_{n-|\ell-k|} c_n^*\} - j \cdot \text{Im}\{c_{n-|\ell-k|} c_n^*\} \text{sign}\{\ell - k\} \right) \right) \cdot \left(\sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) |c_n|^2 \right)^{-1} \end{aligned} \quad (15)$$

EM algorithm has the appealing property of monotonically increasing the log-likelihood in every iteration [25]:

$$\log p(\mathbf{y} | \hat{\boldsymbol{\theta}}^{(t)}) \geq \log p(\mathbf{y} | \hat{\boldsymbol{\theta}}^{(t-1)}).$$

For a more detailed derivation of (12) and (13) and an in-depth treatment of the EM algorithm, we refer the reader to [10].

For the application of the EM algorithm to our problem in (10), we need to solve the optimization problems of the E-step and M-step, respectively. The optimal solution for the E-step is given in (12) by the APP distribution $Q(c) = P(c | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) = \prod_{n=1}^N P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})$. Finding the symbol-wise posteriors $P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})$ based on the channel observation \mathbf{y} and some model assumptions $\hat{\boldsymbol{\theta}}^{(t-1)}$ is a probabilistic inference task similar to the problem discussed in Sec. II-B. In other words, the E-step of the EM algorithm resolves to soft-output symbol detection using the CSI $\hat{\boldsymbol{\theta}}^{(t-1)}$.

Driven by the requirement of finding a *low-complexity* solution, we propose to use the factor graph-based BP algorithm discussed in Sec. II-B to efficiently perform the E-step by approximating the symbol-wise posteriors $P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})$ with the message passing beliefs $b_n(c_n)$.

Based on the APP distribution $P(c | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})$ or an approximation thereof, the M-step updates the parameter estimates $\hat{\boldsymbol{\theta}}^{(t-1)}$. The optimization problem in (13) cannot be solved jointly for the entire parameter vector $\boldsymbol{\theta}$ in closed form. However, we can find a closed-form solution for the maximization of each scalar parameter θ_ℓ by solving $\frac{\partial}{\partial \theta_\ell} \tilde{Q}(\boldsymbol{\theta} | \hat{\boldsymbol{\theta}}^{(t-1)}) \Big|_{\theta_\ell = \hat{\theta}_\ell^{(t-1)}} = 0$, as stated in the following theorem:

Theorem 1. *The setting of $\hat{\boldsymbol{\theta}}^{(t)} = (\hat{h}_0, \dots, \hat{h}_L, \hat{\sigma}^2)^T$ which solves the optimization problem in (13) for the system model detailed in Sec. II-A along each dimension of $\boldsymbol{\theta}$ independently is given by (14) and (15), where B , $C(c_n, \mathbf{h})$ and $D(c_n, c_m, \mathbf{h})$ are given in (A.3).*

Proof: The proof is given in Appendix A. ■

Theorem 1 implies that the update of a single parameter $\hat{\theta}_\ell^{(t)}$ does not directly depend on its respective predecessor $\hat{\theta}_\ell^{(t-1)}$. This dependency is merely given indirectly via the posterior distribution $P(c | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})$ in the E-step. On the other hand, the update equations directly depend on the estimates of other parameters, e.g., the update of $\hat{h}_\ell^{(t)}$ in (15) is a function of $\hat{h}_k^{(t-1)}$ with $k \neq \ell$. This suggests that for optimal parameter updates, the M-step should only update

Algorithm 1: EMBP

Data: $\mathbf{y}, \hat{\boldsymbol{\theta}}^{(0)}$
1 $\nu_{n,m}^{(0)}(c_k) \leftarrow -\log M, k = m, n$ // BP init.
2 **for** $t = 1, \dots, T$ **do**
3 **E-step:** BP-based symbol detection
4 $\mu_{n,m}^{(t)}(c_k), k = m, n \leftarrow$ BP msg. updates (7)
5 $\nu_{n,m}^{(t)}(c_k), k = m, n \leftarrow$ BP msg. updates (8)
6 $b_n(c_n)^{(t)} \leftarrow$ Compute beliefs (9)
7 **M-step:** Select $K^{(t)}$ parameters $\hat{\boldsymbol{\theta}}_\subset^{(t-1)}$ to update
8 $\hat{\boldsymbol{\theta}}_\subset^{(t)} \leftarrow$ update $\hat{\boldsymbol{\theta}}_\subset^{(t-1)}$ using (14), (15) based on
 the approximation $b_n^{(t)}(c_n) \approx P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)})$
Result: $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}^{(T)}, b_n(c_n) = b_n^{(T)}(c_n)$

one element of $\hat{\boldsymbol{\theta}}$ at a time and perform another E-step before updating the next element. In this case, Theorem 1 guarantees the optimality of the respective parameter update in the M-step. However, the element-wise parameter updates come with the price of potentially more EM steps and thus a higher computational complexity. In general, we select the parameters that are updated in parallel during one particular M-step t by defining a sub-vector $\hat{\boldsymbol{\theta}}_\subset^{(t)} := (\hat{\theta}_{J_1}, \dots, \hat{\theta}_{J_{K^{(t)}}})^T$ with $\{J_1, \dots, J_{K^{(t)}}\} \subseteq \{1, \dots, L+2\}$ and only updating the $K^{(t)}$ elements $\hat{\boldsymbol{\theta}}_\subset^{(t)} \leftarrow \hat{\boldsymbol{\theta}}_\subset^{(t-1)}$ while leaving the remaining parameter estimates unchanged.

We summarize the proposed EM procedure for joint estimation and detection in Alg. 1. The algorithm accepts the channel observation \mathbf{y} as well as an initialization for the parameters $\hat{\boldsymbol{\theta}}^{(0)}$ as inputs. Then, it performs T EM steps, each consisting of one BP message passing iteration in the E-step followed by the update of a selection $\hat{\boldsymbol{\theta}}_\subset^{(t-1)}$ of parameter estimates in the M-step. Note that multiple BP iterations per E-step can be performed by skipping the parameter updates in one or multiple consecutive EM steps, i.e., by setting the update vectors $\hat{\boldsymbol{\theta}}_\subset^{(t-1)}$ to length $K^{(t)} = 0$ in the respective EM steps t . The algorithm outputs a final CSI estimate $\hat{\boldsymbol{\theta}}$ and the beliefs $b_n(c_n), n = 1, \dots, N$ as the result of symbol detection. Because the proposed joint estimation and detection algorithm performs BP in a factor graph for the E-step of the EM procedure, we denote it *EMBP algorithm*.

B. Complexity

We consider the complexity of the proposed EMBP algorithm based on Table I, which provides the number of

TABLE I

NUMBER OF ADDITIONS (ADD), MULTIPLICATIONS (MULT), AND $\log \Sigma \exp$ -COMPUTATIONS IN THE STEPS OF THE EMBP ALGORITHM COMPARED TO OTHER ALGORITHMS

Operation	ADD (real-valued)	MULT (real-valued)	$\log \Sigma \exp$	asymptotic complexity
E-step(BP)	$NM(2LM + 4L + 4) + N(4L + 4) + L(M^2 + 2L + 3) + 1$	$NM(3) + N(4L + 4) + L(3M^2 + 2L + 4) + M + 2$	$NM(2LM + 1)$	$\mathcal{O}(NLM^2)$
M-step($\hat{\sigma}^2$)	$NM(2LM + 4) + N(4L + 6) + L(2L + 5) + 1$	$NM(3LM + 3) + N(4L + 6) + L(2L + 6) + M + 2$	-	$\mathcal{O}(NLM^2)$
M-step(\hat{h}_ℓ)	$NM(4L + 6)$	$NM(6L + 9)$	-	$\mathcal{O}(NLM)$
VAE-LE-step	$12N(L + 1) + 7NM + 18N$	$13N(L + 1) + 9NM + 14N$	-	$\mathcal{O}(NL + NM)$
MAP (BCJR)	$NM^{L+1}(2L + 8)$	$NM^{L+1}(4L + 9)$	$M^{L+1}(3N + 2L)$	$\mathcal{O}(NLM^L)$

required operations for the different steps of the EMBP algorithm. For detailed remarks concerning Table I, we refer the interested reader to Appendix B. To obtain the overall complexity of the EMBP algorithm for a given transmission scenario, we can accumulate the executed operations of the required E/M-steps in each iteration. If every parameter is updated a constant number of times, the asymptotic complexity of the EMBP algorithm to perform blind estimation and joint detection becomes $\mathcal{O}(TNLM^2 + NL^2M)$, where the number of iterations T scales either constantly or linearly with L , depending on the concurrency of the parameter updates in the M-step (line 7 in Alg. 1).

C. EMBP Initialization

An important aspect of the EMBP algorithm, which we did not discuss yet, is the choice of a suitable parameter initialization $\hat{\theta}^{(0)}$. Initialization has been shown to substantially influence the quality of EM-based estimation since a poor initialization might cause the EM algorithm to converge to a locally (instead of globally) optimal solution [10], [28]. For example, in the context of Gaussian mixture models, one of the most widely used applications of the EM algorithm, it has been shown in [36] that the EM algorithm converges with high probability to a bad local optimum which can be arbitrarily worse than that of any global optimum if the number of mixture components is 3 or larger.

For our application, the importance of a proper initialization of the EMBP algorithm can immediately be seen if we consider the initialization of the parameter estimates \hat{h}_ℓ , $\ell = 0, \dots, L$ with their expected values $\mathbb{E}\{h_\ell\} = 0$. In this case, the EM algorithm deadlocks in the local optimum of the prior information⁴ $P(c|\mathbf{y}, \hat{\theta}^{(t-1)}) = P(c) = 1/M$ and $\hat{h}_\ell^{(t)} = 0$, $\forall t$. A more practical option is the initialization with an impulse $\hat{\mathbf{h}}^{(0)} = \mathbf{h}_\delta := \mathbf{e}_{\lceil L/2 \rceil}$, which corresponds to the assumption of a memoryless channel.

We propose to initialize the estimate $\hat{\theta}^{(0)}$ of the EMBP algorithm by employing the variational autoencoder-based linear equalizer (VAE-LE), a lightweight blind channel equalizer [7]. To briefly introduce the VAE-LE, we revisit the derivation of the EM algorithm in Sec. III-A. Tractable ML estimation is achieved by simplifying the objective from the log-likelihood function $\log p(\mathbf{y}|\theta)$ to the ELBO $\mathcal{L}(Q, \theta)$, which

is jointly maximized with respect to θ and $Q(c)$ [6]. The EM algorithm simplifies the problem of joint maximization by alternately fixing θ or Q , which leads to the easier subproblems (12) and (13) in the E-step and M-step, respectively. In contrast, the idea of the VAE-LE is to perform a truly *joint* maximization of $\mathcal{L}(Q, \theta)$ and instead reduce the set of possible distributions $Q(c)$ over which the optimization is performed. Following the idea of variational inference, the trial distribution $Q(c)$ is chosen from a family of probability distributions \mathcal{Q}_Φ which are parametrized by the so-called variational parameters Φ [37].

In particular, the VAE-LE restricts the search space \mathcal{Q}_Φ to distributions $Q_\Phi(c)$ which are the result of linear channel equalization using a finite impulse response (FIR) filter with impulse response $\Phi \in \mathbb{C}^{L_{LE}+1}$. To be more precise, a distribution $Q_\Phi(c) = \prod_{n=1}^N Q_\Phi(c_n)$ is obtained from the output $\hat{c}_{LE} \in \mathbb{C}^N$ of the linear equalizer by applying a soft-demapping based on a Gaussian assumption:

$$Q_\Phi(c_n = c) = \frac{\exp\left(-\frac{|\hat{c}_{n,LE} - c|^2}{\sigma_{VAE}^2}\right)}{\sum_{c' \in \mathcal{M}} \exp\left(-\frac{|\hat{c}_{n,LE} - c'|^2}{\sigma_{VAE}^2}\right)}, \quad c \in \mathcal{M}.$$

To perform blind channel equalization, the VAE-LE jointly maximizes $\mathcal{L}(Q, \theta)$ with respect to the channel parameters θ and the taps of the linear equalization filter Φ . Typically, this maximization is based on gradient methods, e.g., applying S_{VAE} iterations of the Adam algorithm [38]. It is important to note that this optimization is completely unsupervised and is performed online as an integral part of the algorithm. For more details and practical considerations we refer the reader to [7]. Due to its restriction to linear equalization, we presume the VAE-LE to be in general less performant compared to the EMBP algorithm that employs a non-linear detector using BP on factor graphs. On the other hand, we expect the VAE-LE to be more robust against poor initializations, since the simple structure of the FIR-based detector is less prone to get stuck in local optima. In addition, the restriction to a linear filter in the equalizer leads to a particularly low complexity that scales linearly with the block length N , the channel memory L and the size of the constellation M per VAE iteration. This is why we propose to use a few iterations S_{VAE} of the VAE-LE to initialize the parameter estimates $\hat{\theta}^{(0)}$ of the EMBP algorithm.

⁴We assume that the constellation \mathcal{M} is axisymmetric.

D. Factor Graph-based Deep Learning

The proposed EMBP algorithm is a joint estimation and detection scheme that is strongly based on its underlying probabilistic model: the EM algorithm aims at maximizing the likelihood $p(\mathbf{y}|\boldsymbol{\theta})$ and the BP-based symbol detection in the E-step is performed in a factor graph that represents a factorization, i.e., conditional independencies among the random variables in $p(\mathbf{y}|\mathbf{c})$. The reliance on such modeling brings forth several advantages, like the interpretability of intermediate results, e.g., every BP message in the factor graph represents a marginal probability distribution. Further advantages are theoretical guarantees, for instance, the EM algorithm provably increases the likelihood in each step.

Nonetheless, we have introduced several approximations in the derivation of the EMBP algorithm to keep it feasible in terms of complexity. First, we have approximated the true marginals $P(c_n|\mathbf{y}, \boldsymbol{\theta})$ with the beliefs $b_n(c_n)$ as the result of BP on a *cyclic* factor graph. Second, we have simplified the ML parameter estimation by using the EM algorithm as a 2-step iterative approach that maximizes the ELBO instead of the log-likelihood. To overcome the performance loss that comes with these approximations and at the same time maintain a low complexity, we leverage methods of model-based deep-learning [34] to develop a refined data-driven version of the EMBP algorithm in the following.

1) *Enhancing BP accuracy*: Several approaches have been proposed in the literature to improve the performance of BP in cyclic factor graphs, including the generalized BP algorithm by Yedidia et al. [39], and the concave-convex procedure (CCCP) which offers convergence guarantees [40]. Importantly, these generalized message-passing algorithms can be integrated into the proposed EMBP framework, as it only requires an iterative algorithm that refines approximations of the posterior distribution. However, the aforementioned algorithms significantly increase the computational complexity which is not in line with the focus of this work. A common method to improve the performance of BP while maintaining its low complexity is *neural BP*, which unrolls the BP iterations in the factor graph to a feed-forward network and assigns trainable weights to every edge in the unrolled network [41]. Neural BP can significantly improve the detection performance for cyclic factor graphs based on the Ungerboeck observation model if the algorithm is optimized for one specific channel [42]. However, our investigations show that it does not generalize well for a broad range of channel characteristics, as it is required for the application of blind symbol detection on block-fading channels in this work.

A much simpler generalization of the BP algorithm, which we expect to better generalize on fading channels, is the introduction of “momentum” in the message updates [43], i.e., replacing the BP message updates in (7), (8) with a convex combination of the new and the old messages, respectively. For example, the message $\mu_{n,m}^{(t)}(c_n)$ in iteration t is replaced with $\beta_{BP} \cdot \mu_{n,m}^{(t)}(c_n) + (1 - \beta_{BP}) \cdot \mu_{n,m}^{(t-1)}(c_n)$. By choosing $0 < \beta_{BP} < 1$, the idea is to improve the convergence behavior of the message passing scheme compared to the original BP ($\beta_{BP} = 1$) while retaining the same fixed points of

the iterative message passing. We generalize this idea, where a constant $\beta_{BP} \in (0, 1]$ is used among all message updates [43], by using an individual weight $\beta_{BP}^{(t)} \in \mathbb{R}$ for each BP iteration $t = 1, \dots, T$, thus unfolding these iterations into a trainable machine learning architecture. We expect that this additional degree of freedom allows the message passing to improve its convergence behavior during the early BP iterations (global optimization) and at the same time to reduce the overall number of iterations by speeding up the local optimization during the final BP iterations. Note that this momentum-based BP is a generalization of BP and by setting all weights $\beta_{BP} = 1$, we recover the original BP algorithm. Hence, by an optimal choice of the weights β_{BP} with respect to a given figure of merit, the overall performance of the EMBP algorithm will only be improved.

To find the optimal weights $\beta_{BP}^{(t)}$, we use a training data set

$$\mathcal{D} = \{(\mathbf{c}^{(d)}, \boldsymbol{\theta}^{(d)}, \mathbf{y}^{(d)}), d = 1, \dots, D\}, \quad (16)$$

consisting of randomly and independently sampled information sequences $\mathbf{c}^{(d)}$, channel realizations $\boldsymbol{\theta}^{(d)}$, and a corresponding noisy channel observation $\mathbf{y}^{(d)}$. We iteratively apply the EMBP algorithm to elements of \mathcal{D} , evaluate the results towards one of the objective functions introduced in Sec. II-D (MSE or BMI), and backpropagate its gradient with respect to the trainable parameters which are then optimized using Adam.

2) *Learning the EM update schedule*: Another part of EMBP which we optimize is the selection of parameters $\hat{\boldsymbol{\theta}}_{\mathcal{C}}^{(t)}$ that are updated in each M-step of the EM algorithm (see line 7 in Alg. 1). Finding an optimal EM update schedule addresses challenges related to: optimizing the detection performance (C1), e.g., “Is it favorable to perform multiple BP iterations before updating the next parameter?”, and reducing complexity (C2), e.g., “Given a limited number of EM iterations T and of total parameter updates $K_{EM} := \sum_{t=1}^T K^{(t)}$, how does an optimal update schedule look like?”. Learning a schedule is generally non-trivial as it corresponds to a discrete and combinatorial optimization.

For the purpose of integrating this problem into the discussed gradient-based end-to-end optimization framework, we convert the schedule learning task into the optimization of continuous weights, by leveraging the concept of momentum in the EM parameter updates. Similarly to the discussed modification of the BP message updates, we replace the original EM parameter estimates $\hat{\boldsymbol{\theta}}_k^{(t)}$ with momentum-based updates $\beta_{EM,k}^{(t)} \cdot \hat{\boldsymbol{\theta}}_k^{(t)} + (1 - \beta_{EM,k}^{(t)}) \cdot \hat{\boldsymbol{\theta}}_k^{(t-1)}$ where we have introduced the weights $(\beta_{EM,1}^{(t)}, \dots, \beta_{EM,L+2}^{(t)})^T =: \boldsymbol{\beta}_{EM}^{(t)} \in \mathbb{R}^{L+2}$, $t = 1, \dots, T$. In order to allow every possible schedule to be learned during the training procedure, we activate the computation of all parameter updates in the M-step by setting $\hat{\boldsymbol{\theta}}_{\mathcal{C}}^{(t)} = \hat{\boldsymbol{\theta}}^{(t)}$ for all EM iterations t . The momentum weights are initialized with a serial schedule $\boldsymbol{\beta}_{EM}^{(t)} = \mathbf{e}_{((t-1) \bmod L+2)+1}$ at the beginning of the training phase, i.e., exactly one parameter is updated alternately per EM iteration. Then, the weights $\boldsymbol{\beta}_{EM}^{(t)}$ are tuned as part of the gradient-based optimization, converting the algorithm into a trainable discriminative architecture [44].

To learn an update schedule and reduce the overall number of parameter updates $K_{EM} < T \cdot (L + 2)$, we add the L_1 regularizer that is known to encourage coefficients' sparsity [10, Chap. 3.1]

$$\mathcal{L}_{EM}(K') := \sum_{\beta_k \in \beta_{\min}(K')} |\beta_k|$$

to the original loss function (MSE or BMI), where $\beta_{\min}(K')$ is the set of K' scalar elements in the vectors $\beta_{EM}^{(t)}$, $t = 1, \dots, T$ with the smallest absolute value. Thereby, the regularization penalizes more than $T \cdot (L + 2) - K'$ parameter updates by enforcing K' scalar elements of the vectors $\beta_{EM}^{(t)}$, $\forall t$ to be close to zero, i.e., they are effectively not being updated. During training, we initialize $K' = 0$ and gradually increase K' until the target number of parameter updates $K_{EM} = T \cdot (L + 2) - K'$ is reached. After the training procedure, we implement the learned schedule by removing the K' elements from $\hat{\theta}_C^{(t)}$ that correspond to the respective elements in $\beta_{\min}(K')$. At this point, EMBP is reduced in complexity by reducing the number of required parameter update computations by K' . For the remaining parameter updates, we keep and apply the learned momentum weights $\beta_{EM}^{(t)}$ as a simple and low-complexity generalization of the full EM updates ($\beta_{EM}^{(t)} = 1$).

IV. EXPERIMENTAL STUDY

For the numerical experiments that are considered in this section, we will assume the transmission blocks to contain $N = 100$ symbols from a binary phase-shift keying (BPSK) constellation. Additional simulation results for a quadrature phase-shift keying (QPSK) transmission are provided in Appendix D. Unless explicitly mentioned otherwise, we apply $S_{VAE} = 10$ iterations of the VAE-LE with $L_{LE} = 2L$ and a learning rate $l_{VAE} = 0.1$ of the Adam optimizer. For the EMBP algorithm, we apply $T = 3 \cdot (L + 2)$ steps of the EMBP algorithm with a serial parameter update schedule, i.e., every parameter θ_ℓ is updated 3 times.

A. Channel Estimation Analysis

We start the numerical evaluation by analyzing the channel estimation performance of the proposed EMBP algorithm and the VAE-LE, for different initializations of the channel parameters $\hat{\mathbf{h}}_{\text{init}}$. To this end, we sample 10^5 random channels with memory $L = 5$ and corresponding channel observations \mathbf{y} at $\text{snr} = 10$ dB. Fig. 2 evaluates the squared error $\|\hat{\mathbf{h}} - \mathbf{h}\|^2$ of the VAE-LE estimate $\hat{\mathbf{h}}_{VAE}$ and of the EMBP algorithm ($\hat{\mathbf{h}}_{EMBP}$) for three different initialization methods: (a) $\hat{\mathbf{h}}_{\text{init}} = \mathbf{h} + \sqrt{\gamma}\mathbf{h}_w$ where $\mathbf{h}_w \sim \mathcal{CN}(0, 1)$ is standard complex normally distributed, i.e., the initial estimate $\hat{\mathbf{h}}_{\text{init}}$ is a noisy version of the true channel impulse response \mathbf{h} and $\gamma \in \mathbb{R}^+$ indicates the variance of the noise. Note that this method is genie-aided due to its dependence on \mathbf{h} ; (b) blind initialization with an impulse $\hat{\mathbf{h}}_{\text{init}} = \mathbf{h}_\delta$; and (c) initialization of the EMBP algorithm by using the estimation result of the VAE-LE which is, in turn, initialized according to method (b).

The results of method (a) are given by the line plots in Fig. 2, where the horizontal axis indicates the squared




Fig. 2. Squared error of the channel estimation $\hat{\mathbf{h}}$ for the EMBP and VAE-LE algorithm at $\text{snr} = 10$ dB for 10^5 random channels with $L = 5$, based on different initializations $\hat{\mathbf{h}}_{\text{init}}$: (a) $\hat{\mathbf{h}}_{\text{init}} = \mathbf{h} + \sqrt{\gamma}\mathbf{h}_w$ with $\mathbf{h}_w \sim \mathcal{CN}(0, \mathbf{I}_{L+1})$, $\gamma \in \mathbb{R}^+$ (solid line: mean, dashed line: median, dotted lines: 25/75 percentiles), (b) $\hat{\mathbf{h}}_{\text{init}} = \mathbf{h}_\delta$ (circles: mean), (c) initialization of the EMBP algorithm using the result \mathbf{h}_{VAE} of the VAE-LE with $\hat{\mathbf{h}}_{\text{init}} = \mathbf{h}_\delta$ (star: mean). The horizontal axis indicates the MSE of the initialization $\hat{\mathbf{h}}_{\text{init}}$.

initialization error $\|\hat{\mathbf{h}}_{\text{init}} - \mathbf{h}\|^2 = \gamma \cdot (L + 1)$ averaged over all 10^5 channel realizations. The vertical axis plots the quality of the channel estimation by the VAE-LE and EMBP algorithm. Both algorithms have a monotonic improvement of the estimation quality with a decreasing initialization error. The MSE of the VAE-LE saturates to $\|\hat{\mathbf{h}}_{VAE} - \mathbf{h}\|^2 = 0.24$ for small initialization errors. We conjecture that this error floor is caused by the limited capabilities of the linear equalizer in the joint estimation and detection process of the VAE-LE. In comparison, the MSE of the EMBP algorithm saturates to a much smaller error floor. We can observe that the mean and median substantially differ from each other which indicates the existence of outliers for the EMBP algorithm.

Another noteworthy effect is the waterfall behavior of the median squared error for the EMBP algorithm. Nonlinear estimators are well-known to exhibit such a waterfall effect where the squared error increases abruptly below a certain threshold [45]. This behavior is typically caused by very noisy observations which lead to a transition from local to global estimation errors in the ML estimates [46].

Comparing the results to initialization method (b), it is worth mentioning that the latter can surpass the MSE results of method (a) with the same average initialization error for the EMBP algorithm. Finally, we analyze the proposed method (c) where the VAE-LE is leveraged to produce an initialization for the EMBP algorithm. The resulting estimates of the EMBP algorithm outperform the estimation performance of method (a) with the same mean squared initialization error $\|\hat{\mathbf{h}}_{\text{init}} - \mathbf{h}\|^2 = 0.29$. This confirms the VAE-LE as a very well-suited method to initialize the EMBP algorithm.

We furthermore evaluate the squared error $\|\hat{\mathbf{h}} - \mathbf{h}\|^2$ for the considered methods as a function of the SNR in Fig. 3. In general, the estimation quality improves for higher snr. For the EMBP algorithm initialized according to method (b)




Fig. 3. MSE versus snr for 10^7 random channels with $L = 5$ for various algorithms and initialization methods (solid line: mean, dashed line: median).

with \mathbf{h}_δ , we can again observe the deviation of the mean from the median especially for high snr. This deviation can be significantly reduced by using the VAE-LE for initialization: the MSE of the EMBP algorithm with method (c) is much closer to its respective median.

We compare the results of the blind estimation methods to a pilot-based channel estimation. To this end, we fix a portion of the N information symbols in the transmission block to pseudo-random pilot symbols that are known at the receiver. Based on these pilot symbols, the receiver performs an ML estimation of \mathbf{h} .⁵ Fig. 3 shows the performance of the consequent ML estimation based on $N \cdot 10\%$ and $N \cdot 20\%$ pilot symbols (gray lines). Further, we consider a decision-directed (DD) scheme which uses not only the pilot symbols for channel estimation but also incorporates the remaining received symbols in \mathbf{y} to acquire more precise CSI. To this end, we perform MAP detection using the already obtained pilot-based channel estimates. A refined ML channel estimation is then performed by incorporating the entire data block based on the pilot symbols together with a hard decision of the MAP detection for the remaining symbols. The results are shown in Fig. 3 (red lines) and confirm the good estimation performance of the blind EMBP scheme with initialization method (c), which consistently outperforms the DD-MAP baseline with 10% pilots and competes with the variant that uses 20% pilots.

B. Symbol Detection Evaluation

Based on the same data that was sampled for the results in Fig. 3, we evaluate the detection performance of the considered joint estimation and detection schemes in terms of the BER in Fig. 4. For low snr, the BER performance of the VAE-LE is relatively close to the optimal coherent MAP performance. This performance gap increases significantly for higher snr where the BER of the VAE-LE only improves marginally. The sampled channels can introduce severe ISI where linear equal-




Fig. 4. BER over snr for various detection schemes, averaged over 10^7 random channels with $L = 5$.

izers like the FIR filter of the VAE-LE⁶ perform poorly [2]. The EMBP algorithm performs significantly better. For low snr, it can be compared to MAP detection performance based on the DD-MAP estimation baseline with 20% pilots. In the high SNR regime, the EMBP algorithm runs into an error floor. This effect is mainly caused by the suboptimality of the BP detector due to many short cycles in the underlying factor graph, which can cause a non-convergent behavior of the BP algorithm. Especially for high snr, the dynamic range of the factor nodes (5) and (6) is further increased [42] and the suboptimality of BP becomes the dominant error source. To reduce this error floor, we can leverage the momentum-based BP message updates as introduced in Sec. III-D1. We optimize the weights $\beta_{BP}^{(t)}$, $t = 1, \dots, L + 2$ for generic channels, i.e., each sample in the training data set \mathcal{D} consists of an independent channel realization and the snr is uniformly sampled in $[0, 12]$ dB. This generic training can be performed offline and prior to evaluation. We apply batch gradient descent optimization with 200 batches each containing 1000 transmission blocks. As shown in Fig. 4, the refined EMBP* algorithm with optimized momentum weights β_{BP}^* can effectively reduce the error floor⁷ and thereby reduces the gap to coherent MAP detection significantly for a target BER = 10^{-2} .

A surprising observation in Fig. 4 is that the blind EMBP algorithm clearly outperforms the coherent BP detector for snr > 3 dB. Although the EMBP algorithm was demonstrated to produce good channel estimates in Fig. 3, it seems counterintuitive that a blind detector is able to outperform its coherent counterpart which has full access to the CSI. While this intuition is valid for optimal MAP detection, it is not given for suboptimal detectors such as BP on cyclic factor graphs, whose suboptimality directly and heavily depends on the characteristics of the underlying channel assumption. To exemplify this phenomenon, we take a closer look at the transmission of 10^4 random information blocks over one

⁵For the pilot-aided baseline methods, we only estimate the channel impulse response \mathbf{h} and assume perfect knowledge of σ^2 , as opposed to the blind schemes.

⁶For channels with less severe ISI, the VAE-LE typically performs significantly better, as demonstrated in [7].

⁷The lowered but remaining error floor of the EMBP* algorithm in Fig. 4 vanishes by the application of FEC, as demonstrated in Appendix C.




Fig. 5. BER and ELBO $\mathcal{L}(Q, \theta)$ of the exact APP distribution Q_{APP} and the approximate distribution Q_{BP} for different channel assumptions $\mathbf{h} = \alpha \mathbf{h}_0$.

specific channel with impulse response $\mathbf{h}_0 = (0.3 - 0.3j, 0.6 - 0.1j, 0.6 - 0.3j)^T$ and $\text{snr} = 10$ dB. The coherent BP detector ($\text{BER} = 0.21$) is significantly outperformed by the EMBP algorithm with a $\text{BER} = 0.051$. The reason for this improved detection performance becomes obvious if we evaluate the mean of the channel estimates $\hat{\mathbf{h}}_{0,\text{EMBP}} = (0.22 - 0.26j, 0.51 - 0.063j, 0.51 - 0.26j)^T$: the channel taps of the EMBP estimate are consistently reduced in their amplitude compared to the true channel \mathbf{h}_0 . This directly influences the dynamic range of the factor nodes (5) and (6), known to notably affect the convergence of BP on factor graphs with cycles [47].

But why does the EMBP algorithm converge to such a “surrogate” channel $\hat{\mathbf{h}}_{0,\text{EMBP}}$, despite the fact that it is based on ML estimation? To answer this question, we recall that the EM algorithm simplifies ML estimation to the maximization of the ELBO $\mathcal{L}(Q, \theta)$ which is only equivalent to maximizing the likelihood function if the E-step is using *exact* APP distributions $P(c|\mathbf{y}, \hat{\theta}^{(t-1)})$. Fig. 5 visualizes $\mathcal{L}(Q, \theta)$ both for the true posterior distribution $Q_{\text{APP}} = P(c|\mathbf{y}, \theta)$ and for the approximative BP solution $Q_{\text{BP}} = \text{BP}(\mathbf{y}, \theta)$ over a one-dimensional subspace $\{\alpha \mathbf{h}_0 : \alpha \in \mathbb{R}^+\}$ of the complete estimation space of \mathbf{h} at $\text{snr} = 10$ dB. For the APP distribution, $D_{\text{KL}}(Q_{\text{APP}}||P) = 0$ holds and thus $\mathcal{L}(Q_{\text{APP}}, \theta) = \log p(\mathbf{y}|\theta)$. The maximum of the ELBO $\mathcal{L}(Q_{\text{APP}}, \theta)$ is therefore at $\alpha = 1$, i.e., the EM algorithm converges to the true channel $\mathbf{h} = \mathbf{h}_0$. In contrast, the KL divergence between Q_{BP} and P is non-vanishing for $\alpha > 0.6$ and rapidly grows with increasing α . This degrading accuracy of the BP algorithm for large α can also be observed when evaluating the BER of the BP-based detection which has its minimum at $\alpha = 0.61$. The maximum of the ELBO for BP consequently lies at $\alpha = 0.66$ between the ML solution ($\alpha = 1.0$) and a small KL divergence ($\alpha < 0.6$). This spot corresponds to the “surrogate” channel to which the EMBP algorithm converges by maximizing the ELBO $\mathcal{L}(Q_{\text{BP}}, \theta)$ ⁸. Thereby, EMBP finds a channel representation that is better suited for BP-based symbol detection compared to the ground truth channel.

⁸Note that the damping of the channel taps with a *scalar* factor $\alpha \in \mathbb{R}^+$ in Fig. 5 is only for illustrative purposes and does not represent the true maximum within the complete estimation space of θ .

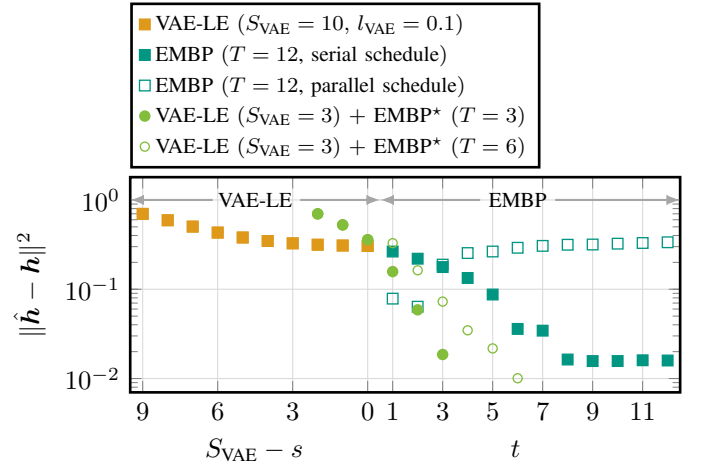


Fig. 6. Mean squared estimation error $\|\hat{\mathbf{h}} - \mathbf{h}\|^2$ after each iteration of the VAE-LE and the EMBP algorithm with different EM parameter update schedules. The randomly sampled channels have $\text{snr} = 10$ dB and memory $L = 5$, i.e., 7 parameters to estimate (including σ^2).

The channels considered so far follow a uniformly distributed power-delay profile (PDP) as defined in Sec. II-A. However, many realistic communication channels follow a PDP that is non-uniformly distributed, e.g., exponentially decaying. The behavior of the proposed EMBP algorithm on this alternative channel characteristic behaves qualitatively similar to the previously discussed results. Detailed evaluations for this case are provided in Appendix D, which also contains results for a QPSK transmission.

C. Complexity Reduction using Model-based Deep Learning

We study the capability of the proposed model-based deep learning methods to reduce the computational complexity of the EMBP algorithm. To this end, we consider a transmission over randomly sampled channels with memory $L = 5$ and optimize the EMBP algorithm towards minimum MSE of the channel estimates $\hat{\mathbf{h}}$ while restraining the algorithm’s complexity. More specifically, we reduce the number of VAE-LE iterations for the initialization from $S_{\text{VAE}} = 10$ to $S_{\text{VAE}} = 3$ and limit the EM iterations to $T = 6$ or $T = 3$, respectively. Based on this restriction $T < L + 2 = 7$, the EMBP algorithm would not be able to update each parameter (including σ^2) at least once if it followed a serial update schedule where only one parameter is updated per EM iteration. We apply the gradient-based end-to-end training method, discussed in Sec. III-D, to find a more effective EM update schedule for this setting. Besides the parameters $\beta_{\text{EM}}^{(t)}$, $t = 1, \dots, T$ that define the EM update schedule, we also include the momentum BP weights $\beta_{\text{BP}}^{(t)}$ in the end-to-end optimization. We use 2500 batches for the offline training and each batch contains 1000 random channel realizations. For the schedule learning, we fix the maximum number of parameter updates to $K_{\text{EM}} = 24$.

To mitigate the effect of the reduced number of VAE-LE iterations S_{VAE} , we additionally apply a hyperparameter search to obtain the optimized learning rates $(l_{\text{VAE}}^{(1)}, l_{\text{VAE}}^{(2)}, l_{\text{VAE}}^{(3)}) = (0.1, 0.16, 0.3)$ for each VAE-LE iteration $s = 1, 2, 3$. The left side in Fig. 6 shows the squared

EMBP* iteration t	$\beta_{EM}^{(t)}$						
	\hat{h}_1	\hat{h}_2	\hat{h}_3	\hat{h}_4	\hat{h}_5	\hat{h}_6	$\hat{\sigma}^2$
1	0.0	0.0	0.0	0.63	0.0	0.0	0.0
2	0.76	0.84	0.0	0.0	0.82	0.78	0.0
3	0.70	0.64	1.60	0.0	0.80	0.83	0.0
4	0.77	0.96	0.0	1.26	0.0	0.97	0.93
5	0.64	0.0	0.84	0.0	0.84	0.0	0.0
6	1.15	1.18	1.39	1.34	1.19	1.16	0.0

Fig. 7. Optimized EM schedule of the EMBP* algorithm based on the momentum weights $\beta_{EM}^{(t)}$ for the complexity constraints $T = 6$ and $K_{EM} = 24$.

estimation error after each VAE-LE iteration compared to a baseline with $S_{VAE} = 10$ VAE-LE steps using a constant learning rate $l_{VAE} = 0.1$. The results are averaged over 10^5 random channels with $\text{snr} = 10$ dB. We can observe that the VAE-LE with optimized learning rates reduces the estimation error faster and achieves the same $\text{MSE} = 0.36$ after only 3 iterations for which the baseline VAE-LE requires 6 steps.

Based on the VAE-LE initialization, the right side in Fig. 6 shows the estimation performance of the EMBP algorithm over the EM iterations t . Using a serial schedule, the MSE monotonically decreases. The stagnation at $t = 7$ is due to the parameter estimate $\hat{\sigma}^2$ not being *directly* included in the evaluation of the squared error $\|\hat{\mathbf{h}} - \mathbf{h}\|^2$. The second baseline in Fig. 6 uses a parallel schedule for EMBP, i.e., all parameters are simultaneously updated in every EM step t . As a result, the MSE decreases significantly after the first EM iteration but worsens after $t = 2$ due to an unstable behavior of the fully parallel updates. The EMBP* algorithm with optimized update schedules shows a fast yet stable convergence and the estimation results are comparable to those of the serial EM schedule but with less iterations, i.e., with a reduced computational complexity.

The momentum weights $\beta_{EM}^{(t)}$ which define the corresponding learned schedule for $T = 6$ are visualized in Fig. 7. In the first EMBP iteration $t = 1$ there is only one parameter \hat{h}_4 updated with small weight, i.e., the algorithm basically performs two BP iterations before starting to update the channel estimate. In the following EMBP iterations $t = 2, \dots, 5$, multiple parameters are updated simultaneously. Only in the last iteration $t = 6$, all channel taps are updated in parallel.

D. Comparison to Data-aided Receivers

At last, we compare the blind EMBP algorithm to state-of-the-art data-aided deep receivers. As a model-agnostic approach, we consider a recurrent neural network (RNN) with window size $L_h + 1$ followed by a linear layer and the softmax function. Further, we use the ViterbiNet detector [17] as a model-based deep learning method. We consider the transmission over a time-invariant channel with a real-valued and exponentially decaying impulse response $\mathbf{h}_e = (0.802, 0.487, 0.295, 0.179)^T$. Both data-aided detectors are provided with one transmission block, i.e., $N = 100$ BPSK symbols, of labeled training data based on




Fig. 8. Comparison of the blind EMBP algorithm to data-aided deep receivers (RNN and ViterbiNet) with respect to the BER over snr for the channel \mathbf{h}_e , averaged over 10^5 random transmission blocks.

which they perform 500 epochs of gradient descent optimization to adapt to the channel \mathbf{h}_e . For a more involved elaboration on both considered algorithms, their respective training, and an open-source implementation, we refer to [21].

Fig. 8 reports the BER versus snr, averaged over 10^5 transmission blocks. Both the RNN-based detector and the ViterbiNet algorithm were trained specifically for each snr value and the evaluation results are averaged over 50 independent training runs with varying initialization. In contrast, the EMBP detector is adaptive to the channel by design without the requirement of channel-specific training. This also holds for the EMBP* algorithm which was optimized a priori based on randomly sampled channels with memory $L = 3$ and uniformly distributed $\text{snr} \sim \mathcal{U}[0, 12]$ dB, but *not* specifically on the channel \mathbf{h}_e , i.e., it does not require any pilots during the actual transmission. In the low snr regime, the EMBP algorithm is close to coherent MAP performance with an snr gap of 1.5 dB for a target $\text{BER} = 10^{-1}$. It thereby significantly outperforms the ViterbiNet detector and the RNN-based method which have an snr gap of 4.5 dB and 8.5 dB, respectively. For high snr, the severe ISI of the channel leads to a very poor performance of the coherent BP detector. The EMBP algorithm performs considerably better but also runs into an error floor. For $\text{snr} > 7$ dB, it is thereby outperformed by the ViterbiNet detector whose suboptimality mainly originates from an imprecise CSI due to the limited amount of training data. The EMBP* algorithm has a significantly reduced error floor and outperforms the ViterbiNet detector in the considered high snr regime with a 3 dB gain for a target $\text{BER} = 10^{-2}$.

V. CONCLUSION

We studied the problem of blind joint channel estimation and symbol detection for linear block-fading channels with memory and introduced the EMBP algorithm which interweaves the iterations of EM and BP on a suitable factor graph. The VAE-LE was found as a suitable method to initialize the EMBP algorithm whose convergence is sensitive to the initialization. We further leveraged techniques of model-based deep learning to develop the refined algorithm EMBP* which

uses momentum in the updates of BP as well as in the EM updates to significantly improve the detection performance and reduce the algorithm's complexity with only a few learnable parameters. A complexity analysis and various numerical simulations demonstrate that the proposed EMBP* algorithm fulfills the ambition of a fully blind low-complexity symbol detection algorithm that performs comparable to pilot-based or coherent symbol detectors.

APPENDIX

A. Derivation of the EM Updates

We consider the optimization problem (13). A necessary condition for the optimal solution $\hat{\theta}^{(t)}$ is

$$\mathbf{0} = \nabla_{\theta} \tilde{Q}(\theta | \hat{\theta}^{(t-1)}) \Big|_{\theta=\hat{\theta}^{(t)}}. \quad (\text{A.1})$$

To rewrite the term

$$\begin{aligned} \tilde{Q}(\theta | \hat{\theta}^{(t-1)}) &= \sum_{\mathbf{c}} P(\mathbf{c} | \mathbf{y}, \hat{\theta}^{(t-1)}) \log p(\mathbf{y}, \mathbf{c} | \theta) \\ &= \sum_{\mathbf{c}} P(\mathbf{c} | \mathbf{y}, \hat{\theta}^{(t-1)}) \log \left(\frac{p(\mathbf{y} | \mathbf{c}, \theta)}{M^N} \right), \end{aligned}$$

we use $P(\mathbf{c} | \theta) = P(\mathbf{c}) = M^{-N}$. The likelihood function $p(\mathbf{y} | \mathbf{c}, \theta)$ can be decomposed in the same manner as in (3). However, note that we must not omit the terms that are independent of \mathbf{c} in this case, since we are maximizing over θ instead of \mathbf{c} . This leads to

$$\begin{aligned} \log \left(\frac{p(\mathbf{y} | \mathbf{c}, \theta)}{M^N} \right) &= -N \log(M\pi\sigma^2) \\ &\quad + \left(\frac{-\mathbf{y}^H \mathbf{y} + 2\text{Re}\{\mathbf{c}^H \mathbf{x}\} - \mathbf{c}^H \mathbf{G} \mathbf{c}}{\sigma^2} \right) \\ &= -A(\sigma^2) + \frac{-B + C(\mathbf{c}, \mathbf{h}) - D(\mathbf{c}, \mathbf{h})}{\sigma^2}, \end{aligned} \quad (\text{A.2})$$

where we introduced the terms

$$\begin{aligned} A(\sigma^2) &:= N \log(M\pi\sigma^2), \quad B := \sum_{n=1}^{N+L} |y_n|^2, \\ C(\mathbf{c}, \mathbf{h}) &:= \sum_{n=1}^N \underbrace{2\text{Re}\{x_n c_n^*\} - G_{n,n} |c_n|^2}_{=: C_n(\mathbf{c}_n, \mathbf{h})}, \quad \text{and} \\ D(\mathbf{c}, \mathbf{h}) &:= \sum_{n=1}^N \sum_{m < n} \underbrace{2\text{Re}\{G_{m,n} c_m c_n^*\}}_{=: D_{n,m}(\mathbf{c}_n, \mathbf{c}_m, \mathbf{h})}. \end{aligned} \quad (\text{A.3})$$

To keep the notation uncluttered, we use the shorter notations $C(\mathbf{c}_n, \mathbf{h})$ and $D(\mathbf{c}_n, \mathbf{c}_m, \mathbf{h})$ for the terms $C_n(\mathbf{c}_n, \mathbf{h})$ and $D_{n,m}(\mathbf{c}_n, \mathbf{c}_m, \mathbf{h})$, respectively. Note that the terms \mathbf{x} and $G_{n,m}$ are a function of the channel impulse response \mathbf{h} which is part of θ . This means that (A.1) is a nonlinear system of equations that can generally not be solved in closed form. However, we can find a closed-form solution for the maximization along each dimension of θ , respectively, as suggested by Theorem 1. Therefore, we solve each row $i = \ell$ in (A.1) for one element θ_ℓ , respectively, leveraging the formulation of the

log-likelihood in (A.2). The partial derivative with respect to $\theta_{L+2} = \sigma^2$ yields

$$\begin{aligned} 0 &= \frac{\partial}{\partial \sigma^2} \tilde{Q}(\theta | \hat{\theta}^{(t-1)}) \Big|_{\sigma^2=\hat{\sigma}^{2(t)}} \\ \Leftrightarrow 0 &= \sum_{\mathbf{c}} P(\mathbf{c} | \mathbf{y}, \hat{\theta}^{(t-1)}) \left(\frac{-N}{\hat{\sigma}^{2(t)}} + \frac{B - C(\mathbf{c}, \mathbf{h}) + D(\mathbf{c}, \mathbf{h})}{(\hat{\sigma}^{2(t)})^2} \right) \\ \Leftrightarrow \hat{\sigma}^{2(t)} &= \frac{1}{N} \left[B - \sum_{n=1}^N \sum_{\mathbf{c}_n} P(\mathbf{c}_n | \mathbf{y}, \hat{\theta}^{(t-1)}) \left(C(\mathbf{c}_n, \mathbf{h}) \right. \right. \\ &\quad \left. \left. - \sum_{m < n} \sum_{\mathbf{c}_m} P(\mathbf{c}_m | \mathbf{y}, \hat{\theta}^{(t-1)}) D(\mathbf{c}_n, \mathbf{c}_m, \mathbf{h}) \right) \right], \end{aligned} \quad (\text{A.4})$$

i.e., we found the closed-form solution (A.4) for the EM update of the parameter σ^2 .

To simplify the derivation of the parameter updates for \mathbf{h} , but without loss of optimality, we split the complex-valued parameters $h_\ell =: h_{R,\ell} + j h_{I,\ell}$ into their real and imaginary parts and solve for $h_{R,\ell}$ and $h_{I,\ell}$, respectively. The partial derivative with respect to $h_{R,\ell}$ yields the condition

$$\begin{aligned} 0 &= \frac{\partial}{\partial h_{R,\ell}} \tilde{Q}(\theta | \hat{\theta}^{(t-1)}) \Big|_{h_{R,\ell}=\hat{h}_{R,\ell}^{(t)}} \\ &= \sum_{\mathbf{c}} P(\mathbf{c} | \mathbf{y}, \hat{\theta}^{(t-1)}) \frac{\partial}{\partial h_{R,\ell}} (C(\mathbf{c}, \mathbf{h}) - D(\mathbf{c}, \mathbf{h})) \Big|_{h_{R,\ell}=\hat{h}_{R,\ell}^{(t)}}. \end{aligned} \quad (\text{A.5})$$

Using the relations $\frac{\partial}{\partial h_{R,\ell}} \text{Re}\{x_n c_n^*\} = \text{Re}\{y_{n+\ell} c_n^*\}$ and

$$\begin{aligned} \frac{\partial}{\partial h_{R,\ell}} G_{n,n} &= \frac{\partial}{\partial h_{R,\ell}} \left(\sum_{k=0}^L |h_k|^2 \right) = 2h_{R,\ell}, \quad \text{we find that} \\ \frac{\partial}{\partial h_{R,\ell}} C(\mathbf{c}, \mathbf{h}) &= \sum_{n=1}^N \left(2 \frac{\partial}{\partial h_{R,\ell}} \text{Re}\{x_n c_n^*\} - |c_n|^2 \frac{\partial}{\partial h_{R,\ell}} G_{n,n} \right) \\ &= 2 \sum_{n=1}^N (\text{Re}\{y_{n+\ell} c_n^*\} - |c_n|^2 h_{R,\ell}). \end{aligned} \quad (\text{A.6})$$

Substituting (A.6) into (A.5) yields

$$\begin{aligned} &\sum_{\mathbf{c}} P(\mathbf{c} | \mathbf{y}, \hat{\theta}^{(t-1)}) \frac{\partial}{\partial h_{R,\ell}} C(\mathbf{c}, \mathbf{h}) \\ &= 2 \sum_{n=1}^N \sum_{\mathbf{c}_n} P(\mathbf{c}_n | \mathbf{y}, \hat{\theta}^{(t-1)}) (\text{Re}\{y_{n+\ell} c_n^*\} - |c_n|^2 h_{R,\ell}). \end{aligned} \quad (\text{A.7})$$

For the partial derivative of $D(\mathbf{c}_n, \mathbf{c}_m, \mathbf{h})$ we obtain

$$\begin{aligned} \frac{\partial}{\partial h_{R,\ell}} D(\mathbf{c}_n, \mathbf{c}_m, \mathbf{h}) &= 2 \frac{\partial}{\partial h_{R,\ell}} \text{Re}\{G_{m,n} c_m c_n^*\} \\ &= 2 \frac{\partial}{\partial h_{R,\ell}} \text{Re} \left\{ \sum_{k=0}^{L-(n-m)} h_{k+(n-m)}^* h_k c_m c_n^* \right\} \\ &= 2 \text{Re}\{c_m c_n^*\} (h_{R,\ell+(n-m)} + h_{R,\ell-(n-m)}) \\ &\quad + 2 \text{Im}\{c_m c_n^*\} (h_{I,\ell+(n-m)} - h_{I,\ell-(n-m)}). \end{aligned} \quad (\text{A.8})$$

To rewrite the sums in

$$\frac{\partial}{\partial h_{R,\ell}} D(\mathbf{c}, \mathbf{h}) = \sum_{n=1}^N \sum_{m < n} \frac{\partial}{\partial h_{R,\ell}} D(\mathbf{c}_n, \mathbf{c}_m, \mathbf{h}),$$

$$\hat{h}_{R,\ell}^{(t)} = \left(\sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \text{Re}\{y_{n+\ell} c_n^*\} - \sum_{\substack{k=0 \\ k \neq \ell}}^L \sum_{c_{n-|\ell-k|}} P(c_{n-|\ell-k|} | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \right. \\ \left. \left(\text{Re}\{c_{n-|\ell-k|} c_n^*\} h_{R,k} + \text{Im}\{c_{n-|\ell-k|} c_n^*\} \text{sign}\{\ell - k\} h_{I,k} \right) \right) \cdot \left(\sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) |c_n|^2 \right)^{-1} \quad (\text{A.12})$$

$$\hat{h}_{I,\ell}^{(t)} = \left(\sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \text{Im}\{y_{n+\ell} c_n^*\} - \sum_{\substack{k=0 \\ k \neq \ell}}^L \sum_{c_{n-|\ell-k|}} P(c_{n-|\ell-k|} | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \right. \\ \left. \left(\text{Re}\{c_{n-|\ell-k|} c_n^*\} h_{I,k} - \text{Im}\{c_{n-|\ell-k|} c_n^*\} \text{sign}\{\ell - k\} h_{R,k} \right) \right) \cdot \left(\sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) |c_n|^2 \right)^{-1} \quad (\text{A.13})$$

we can use the equivalences

$$\sum_{n=1}^N \sum_{m < n} 2\text{Re}\{c_m c_n^*\} (h_{R,\ell+(n-m)} + h_{R,\ell-(n-m)}) \\ = \sum_{n=1}^N \sum_{\substack{k=0 \\ k \neq \ell}}^L 2\text{Re}\{c_{n-|\ell-k|} c_n^*\} h_{R,k}, \quad (\text{A.9})$$

and

$$\sum_{n=1}^N \sum_{m < n} 2\text{Im}\{c_m c_n^*\} (h_{I,\ell+(n-m)} - h_{I,\ell-(n-m)}) \\ = \sum_{n=1}^N \sum_{\substack{k=0 \\ k \neq \ell}}^L 2\text{Im}\{c_{n-|\ell-k|} c_n^*\} \text{sign}\{\ell - k\} h_{I,k}. \quad (\text{A.10})$$

Based on (A.8) and by using the expressions (A.9) and (A.10), we find for the second term in (A.5):

$$\sum_{\mathbf{c}} P(\mathbf{c} | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \frac{\partial}{\partial h_{R,\ell}} D(\mathbf{c}, \mathbf{h}) \\ = 2 \sum_{n=1}^N \sum_{c_n} P(c_n | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \sum_{\substack{k=0 \\ k \neq \ell}}^L \sum_{c_{n-|\ell-k|}} P(c_{n-|\ell-k|} | \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t-1)}) \\ \left(\text{Re}\{c_{n-|\ell-k|} c_n^*\} h_{R,k} + \text{Im}\{c_{n-|\ell-k|} c_n^*\} \text{sign}\{\ell - k\} h_{I,k} \right). \quad (\text{A.11})$$

Applying (A.7) and (A.11) in (A.5) and solving for $h_{R,\ell}$ finally yields the update equation for $\hat{h}_{R,\ell}^{(t)}$ given in (A.12). In a similar fashion, we can solve $0 = \frac{\partial}{\partial h_{I,\ell}} \tilde{Q}(\boldsymbol{\theta} | \hat{\boldsymbol{\theta}}^{(t-1)})$ for $h_{I,\ell}$ which leads to the update equation (A.13) for $\hat{h}_{I,\ell}^{(t)}$. By studying the two equations (A.12) and (A.13), we can observe that the update of $\hat{h}_{R,\ell}^{(t)}$ is independent of $\hat{h}_{I,\ell}^{(t-1)}$ and that, vice versa, the update of $\hat{h}_{I,\ell}^{(t)}$ does not depend on $\hat{h}_{R,\ell}^{(t-1)}$. Therefore, we can update $\hat{h}_{R,\ell}^{(t)}$ and $\hat{h}_{I,\ell}^{(t)}$ in parallel without loss of optimality, i.e., $\hat{h}_\ell = \hat{h}_{R,\ell}^{(t)} + j\hat{h}_{I,\ell}^{(t)}$, thus proving the theorem. \square

B. Complexity

We analyze the complexity of the proposed EMBP algorithm by counting the number of real-valued additions (ADD),

multiplications (MULT) and $\log \Sigma \exp$ -operations that are required in the E-step (1 BP iteration) and in the M-step to update the parameter estimates $\hat{\sigma}^2$ and \hat{h}_ℓ . Complex additions are counted as 2 ADD operations and complex multiplications are counted as 4 MULT and 2 ADD operations. We neglect interim results that are data-independent, e.g., $|c_n|^2$, as they can be precomputed and stored. Also note that, depending on the parameter update schedule of the EMBP algorithm, some data-dependent interim results can be reused. For simplicity, however, we consider them to be recomputed each time. The results are provided in Table I.

The EMBP* algorithm requires some extra operations for the convex combination in the momentum updates. For the momentum updates of the BP messages, $N(2L+1)$ real-valued messages of dimension M are scaled and added to a scaled version of the old messages, which results in $2NM(2L+1)$ MULT and $NM(2L+1)$ ADD operations. For the momentum-based parameter update of one complex-valued element in $\hat{\boldsymbol{\theta}}$, it requires 4 MULT and 2 ADD operations.

We further provide some remarks about the complexity of the pilot-based and decision-directed MAP detection schemes that are considered in Sec. IV for comparison. For the pilot-based MAP detector with N_p pilots, the complexity of the MAP detector (see Table I) is supplemented by the complexity of the pilot-based ML estimator, which can be implemented by an $(L+1) \times N_p$ complex-valued matrix-vector multiplication. The DD-MAP scheme builds upon the previously mentioned pilot-based MAP detector. Additionally, it requires the ML estimation based on the hard symbol decision of the first MAP detection run. This can be implemented by an $(L+1) \times N$ complex-valued matrix-vector multiplication and an algorithm to solve an $(L+1) \times (L+1)$ Toeplitz system, e.g., the Levinson-Durbin recursion. Furthermore, a second run of the MAP detector is required.

C. FEC-BER Performance

We consider an FEC scheme applied to the BPSK transmission scenario with block length $N = 100$ and memory $L = 5$, as described in Sec. IV with the results shown in Fig. 3 and Fig. 4. For the non-blind baseline algorithms with 20% pilots we apply a low-density parity-check (LDPC) code of

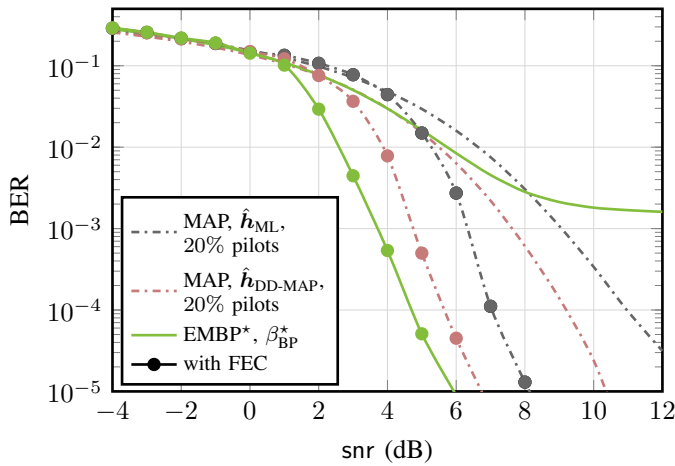


Fig. 9. BER and FEC-BER over snr for a BPSK transmission over 10^7 random channels with $L = 5$. The FEC-BER results are based on an LDPC code of rate 0.8 for the pilot-based schemes and a code rate of 0.64 for the EMBP* algorithm, such that all compared schemes effectively transmit 64 information bits.

rate 0.8, and for the blind EMBP* algorithm, we use an LDPC code of rate 0.64. Thereby, all compared schemes effectively transmit 64 information bits. For channel decoding, we use a normalized min-sum decoder with a maximum number of 20 layered BP iterations. Fig. 9 compares the BER and FEC-BER performance of the EMBP* algorithm with the pilot-based schemes. With FEC, the EMBP* algorithm consistently outperforms the pilot-based MAP detector with a 2 dB SNR gain and the pilot-based DD-MAP baseline with a gain of 1 dB.

D. Channels with Non-uniform PDP

We evaluate the estimation and detection performance of the EMBP algorithm on 10^7 random channels with a non-uniformly distributed PDP. In contrast to the channel model defined in Sec. II-A, the channel taps h_ℓ , $\ell = 0, \dots, L$ are independently sampled from a complex-circular Gaussian distribution with exponentially decaying variance $\mathcal{N}(0, \exp(-\ell))$. For better comparison, the power of the sampled channel impulse responses is normalized as before. Fig. 10 plots the squared estimation error over the snr and Fig. 11 shows the respective BER performance.

We additionally simulate the transmission of $N = 100$ QPSK symbols per transmission block on the same batch of 10^7 channels. The evaluation results for the squared estimation error and the BER over snr are shown in Fig. 12 and Fig. 13, respectively.

REFERENCES

- [1] L. Schmid, T. Raviv, N. Shlezinger, and L. Schmalen, "Optimization of iterative blind detection based on expectation maximization and belief propagation," *preprint arXiv:2408.02312*, 2024.
- [2] J. Proakis and M. Salehi, *Digital Communications*, 5th ed. McGraw Hill, Nov. 2007.
- [3] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error rate (corresp.)," *IEEE Trans. Inf. Theory*, vol. 20, no. 2, pp. 284–287, Mar. 1974.
- [4] L. Tong and S. Perreau, "Multichannel blind identification: From subspace to maximum likelihood methods," *Proc. IEEE*, vol. 86, no. 10, pp. 1951–1968, Oct. 1998.




Fig. 10. MSE over snr for various algorithms (solid line: mean, dashed line: median). Evaluation for a BPSK transmission over 10^7 random channels with $L = 5$ and exponentially decaying PDP.




Fig. 11. BER over snr for a BPSK transmission over 10^7 random channels with $L = 5$ and exponentially decaying PDP.

- [5] D. Godard, "Self-recovering equalization and carrier tracking in two-dimensional data communication systems," *IEEE Trans. Commun.*, vol. 28, no. 11, pp. 1867–1875, 1980.
- [6] A. Caciularu and D. Burshtein, "Unsupervised linear and nonlinear channel equalization and decoding using variational autoencoders," *IEEE Trans. on Cogn. Commun. Netw.*, vol. 6, no. 3, pp. 1003–1018, Sep. 2020.
- [7] V. Lauinger, F. Buchali, and L. Schmalen, "Blind equalization and channel estimation in coherent optical communications using variational autoencoders," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 9, pp. 2529–2539, Sep. 2022.
- [8] M. Ghosh and C. L. Weber, "Maximum-likelihood blind equalization," *Optical Engineering*, vol. 31, no. 6, pp. 1224–1228, 1992.
- [9] G. Kaleh and R. Vallet, "Joint parameter estimation and symbol detection for linear or nonlinear unknown channels," *IEEE Trans. Commun.*, vol. 42, no. 7, pp. 2406–2413, 1994.
- [10] C. M. Bishop, *Pattern recognition and machine learning*. New York: Springer, 2006.
- [11] L. Dai, R. Jiao, F. Adachi, H. V. Poor, and L. Hanzo, "Deep learning for wireless communications: An emerging interdisciplinary paradigm," *IEEE Wireless Commun.*, vol. 27, no. 4, pp. 133–139, 2020.
- [12] S. Dörner, J. Clausius, S. Cammerer, and S. ten Brink, "Learning joint detection, equalization and decoding for short-packet communications," *IEEE Trans. Commun.*, vol. 71, no. 2, pp. 837–850, Feb. 2023.
- [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, Nov. 2018.




Fig. 12. MSE over snr for various algorithms (solid line: mean, dashed line: median). Evaluation for a QPSK transmission over 10^7 random channels with $L = 5$ and exponentially decaying PDP.




Fig. 13. BER over snr for a QPSK transmission over 10^7 random channels with $L = 5$ and exponentially decaying PDP.

[14] N. Shlezinger, J. Whang, Y. C. Eldar, and A. G. Dimakis, "Model-based deep learning," *Proc. IEEE*, vol. 111, no. 5, pp. 465–499, May 2023.

[15] H. He, S. Jin, C.-K. Wen, F. Gao, G. Y. Li, and Z. Xu, "Model-driven deep learning for physical layer communications," *IEEE Wireless Commun.*, vol. 26, no. 5, pp. 77–83, Oct. 2019.

[16] N. Shlezinger and Y. C. Eldar, "Model-based deep learning," *Foundations and Trends® in Signal Processing*, vol. 17, no. 4, pp. 291–416, 2023.

[17] 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.*, vol. 19, no. 5, pp. 3319–3331, May 2020.

[18] —, "Data-driven factor graphs for deep symbol detection," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Los Angeles, CA, USA, Jun. 2020, pp. 2682–2687.

[19] T. Raviv, S. Park, O. Simeone, Y. C. Eldar, and N. Shlezinger, "Online meta-learning for hybrid model-based deep receivers," *IEEE Trans. Wireless Commun.*, vol. 22, no. 10, pp. 6415–6431, 2023.

[20] T. Raviv and N. Shlezinger, "Data augmentation for deep receivers," *IEEE Trans. Wireless Commun.*, vol. 22, no. 11, pp. 8259–8274, 2023.

[21] T. Raviv, S. Park, O. Simeone, Y. C. Eldar, and N. Shlezinger, "Adaptive and flexible model-based AI for deep receivers in dynamic channels," *IEEE Wireless Commun.*, vol. 31, no. 4, pp. 163–169, 2024.

[22] G. Ungerboeck, "Adaptive maximum-likelihood receiver for carrier-modulated data-transmission systems," *IEEE Trans. Commun.*, no. 5, pp. 624–636, May 1974.

[23] G. Colavolpe, D. Fertonani, and A. Piemontese, "SISO detection over linear channels with linear complexity in the number of interferers," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 8, pp. 1475–1485, 2011.

[24] Y. Liu, L. Brunel, and J. J. Boutros, "Joint channel estimation and decoding using Gaussian approximation in a factor graph over multipath

channel," in *Proc. IEEE Int. Symp. Pers. Indoor Mobile Radio Commun. (PIMRC)*, Tokyo, Japan, Sep. 2009, pp. 3164–3168.

[25] A. W. Eckford, "Channel Estimation in Block Fading Channels Using the Factor Graph EM Algorithm," in *Proc. Biennial Symposium on Communications*, Kingston, Ontario, Canada, May 2004.

[26] J. Dauwels, S. Korl, and H.-A. Loeliger, "Expectation maximization as message passing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Adelaide, Australia, Sep. 2005, pp. 583–586.

[27] T. Wang and S. C. Liew, "Joint channel estimation and channel decoding in physical-layer network coding systems: an EM-BP factor graph framework," *IEEE Trans. Commun.*, vol. 13, no. 4, pp. 2229–2245, Apr. 2014.

[28] E. Shireman, D. Steinley, and M. J. Brusco, "Examining the effect of initialization strategies on the performance of Gaussian mixture modeling," *Behav. Res. Methods*, vol. 49, no. 1, pp. 282–293, Feb. 2017.

[29] W. Yang, G. Durisi, T. Koch, and Y. Polyanskiy, "Block-fading channels at finite blocklength," in *Proc. IEEE Symp. Wireless Commun. Syst. (ISWCS)*, Aug. 2013.

[30] F. R. Kschischang and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.

[31] L. Schmid and L. Schmalen, "Neural enhancement of factor graph-based symbol detection," in *Proc. IEEE Int. Workshop on Signal Process. Adv. Wireless Commun. (SPAWC)*, Oulu, Finland, Jul. 2022.

[32] A. Guillén i Fàbregas, A. Martinez, and G. Caire, "Bit-interleaved coded modulation," *Foundations and Trends® in Signal Processing*, vol. 5, no. 1–2, pp. 1–153, 2008.

[33] A. Alvarado, T. Fehenberger, B. Chen, and F. M. J. Willems, "Achievable information rates for fiber optics: Applications and computations," *J. Lightw. Technol.*, vol. 36, no. 2, pp. 424–439, Jan. 2018.

[34] N. Shlezinger, Y. C. Eldar, and S. P. Boyd, "Model-based deep learning: On the intersection of deep learning and optimization," *IEEE Access*, vol. 10, pp. 115 384–115 398, 2022.

[35] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Stat. Soc.: Ser. B*, vol. 39, no. 1, pp. 1–38, Sep. 1977.

[36] C. Jin, Y. Zhang, S. Balakrishnan, M. J. Wainwright, and M. I. Jordan, "Local maxima in the likelihood of Gaussian mixture models: Structural results and algorithmic consequences," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 29, 2016.

[37] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *J. Amer. Stat. Assoc.*, vol. 112, no. 518, pp. 859–877, Apr. 2017.

[38] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, San Diego, CA, USA, May 2015.

[39] J. S. Yedidia, W. T. Freeman, and Y. Weiss, "Generalized belief propagation," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 13, pp. 689–695, 2001.

[40] A. L. Yuille, "CCCP algorithms to minimize the Bethe and Kikuchi free energies: Convergent alternatives to belief propagation," *Neural Computation*, vol. 14, no. 7, pp. 1691–1722, 2002.

[41] E. Nachmani, E. Marciano, L. Lugosch, W. J. Gross, D. Burshtein, and Y. Beery, "Deep learning methods for improved decoding of linear codes," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 1, pp. 119–131, Feb. 2018.

[42] L. Schmid and L. Schmalen, "Low-complexity near-optimum symbol detection based on neural enhancement of factor graphs," *IEEE Trans. Commun.*, pp. 7562–7575, Nov. 2022.

[43] K. Murphy, Y. Weiss, and M. I. Jordan, "Loopy belief propagation for approximate inference: An empirical study," in *Proc. Int. Conf. Uncertainty Artif. Intell. (UAI)*, Stockholm, Sweden, 1999, pp. 467–475.

[44] N. Shlezinger and T. Routenberg, "Discriminative and generative learning for linear estimation of random signals [lecture notes]," *IEEE Signal Process. Mag.*, vol. 40, no. 6, pp. 75–82, 2023.

[45] H. Van Trees, *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory*. Hoboken, NJ, USA: Wiley, Apr. 2004.

[46] F. Athley, "Threshold region performance of maximum likelihood direction of arrival estimators," *IEEE Trans. Signal Process.*, vol. 53, no. 4, pp. 1359–1373, 2005.

[47] A. T. Ihler, J. W. Fisher III, and A. S. Willsky, "Loopy belief propagation: Convergence and effects of message errors," *J. Mach. Learn. Res.*, vol. 6, no. 31, pp. 905–936, 2005.