

Blind Equalization and Channel Estimation in Coherent Optical Communications Using Variational Autoencoders

Vincent Lauinger, *Student Member, IEEE*, Fred Buchali and Laurent Schmalen, *Senior Member, IEEE*

Abstract—We investigate the potential of adaptive blind equalizers based on variational inference for carrier recovery in optical communications. These equalizers are based on a low-complexity approximation of maximum likelihood channel estimation. We generalize the concept of variational autoencoder (VAE) equalizers to higher order modulation formats encompassing probabilistic constellation shaping (PCS), ubiquitous in optical communications, oversampling at the receiver, and dual-polarization transmission. Besides black-box equalizers based on convolutional neural networks, we propose a model-based equalizer based on a linear butterfly filter and train the filter coefficients using the variational inference paradigm. As a byproduct, the VAE also provides a reliable channel estimation. We analyze the VAE in terms of performance and flexibility over a classical additive white Gaussian noise (AWGN) channel with inter-symbol interference (ISI) and over a dispersive linear optical dual-polarization channel. We show that it can extend the application range of blind adaptive equalizers by outperforming the state-of-the-art constant-modulus algorithm (CMA) for PCS for both fixed but also time-varying channels. The evaluation is accompanied with a hyperparameter analysis.

Index Terms—blind equalizers, channel estimation, variational inference, optical fiber communication

I. INTRODUCTION

THE digital transformation along with the modern lifestyle and the advent of video streaming platforms brought up a strong demand for high-speed and highly flexible communication systems. Precisely, the required data rates can only be provided by coherent optical communication systems along with high-order modulation formats and probabilistic constellation shaping (PCS) [2]. Due to its properties, e.g., easy rate adaption [3], a decreased gap to the additive white Gaussian noise (AWGN) channel capacity, increased energy efficiency [4] and a larger tolerance against fiber nonlinearities, PCS has become an essential ingredient of modern coherent optical communication systems [5], [6]. However, the use of PCS entails a more challenging carrier recovery than conventional square quadrature amplitude modulation (QAM)

formats. Often, data-aided or pilot-based algorithms are the only option nowadays, since an open issue in communications is the lack of optimum (but practical) blind adaptive channel equalizers. However, pilot symbols cannot transport information so they reduce data rate and limit the achievable net bit rate significantly. Hence, there is a strong need of blind channel equalizers, which can adapt to time-varying channels and transmission parameters. The saved data rate can be used to either increase the throughput or forward error correction (FEC) overhead.

In coherent optical communications, the standard algorithm for blind adaptive equalization of linear channels and symmetric complex modulation formats is the constant-modulus algorithm (CMA) [7]. It tries to reach a constant signal amplitude (radius) by adaptively equalizing the signal with trainable finite impulse response (FIR) filters. Thus, it is optimal for constant-amplitude formats such as M -ary phase shift keying (PSK), but it also converges for multi-amplitude formats such as M -ary QAM [8] where its criterion is sub-optimal. The multi-modulus algorithm (MMA) [9] is an extension for multi-amplitude formats, however, it suffers from its high implementation complexity and low convergence rate. Based on the same criterion, a non-linear blind neural network (NN) based equalization scheme was proposed in [10]. Since the CMA's criterion is independent of the signal phase, detection is only possible in combination with a carrier-phase estimation (CPE) block. The commonly used algorithms are the blind phase search [11] or the Viterbi-Viterbi algorithm [12], which both face performance degradation for PCS [13], [14]. Similarly, the CMA suffers from convergence issues for PCS as well [15], [16].

Optimally, we want to use the maximum likelihood (ML) criterion, which has been considered for blind equalization, e.g., in [17], [18], [19]. However, we are not aware of any blind ML based channel equalizer which has been seriously considered in real coherent optical communication systems. A promising approach is to approximate ML by *variational inference* via a variational autoencoder (VAE) [20], [21]. Variational inference is used for unsupervised and semi-supervised learning as well as generative models, however, there are not many applications in communications with notable exceptions being [22], [23], [24]. While [22] trains end-to-end transmission systems without inter-symbol interference (ISI) in a supervised manner, we, in contrast, focus on blind VAE-based equalization where we use unsupervised learning at the receiver to adjust the equalizer weights. Such an unsuper-

This work was carried out in the framework of the CELTIC-NEXT project AI-NET-ANTILLAS (C2019/3-3) and was funded by the German Federal Ministry of Education and Research (BMBF) under grant agreement 16KIS1316.

V. Lauinger and L. Schmalen are with Communications Engineering Lab (CEL), Karlsruhe Institute of Technology (KIT), 76131 Karlsruhe, Germany (email: vincent.lauinger@kit.edu, laurent.schmalen@kit.edu).

F. Buchali is with Nokia, 70469 Stuttgart, Germany (email: Fred.Buchali@nokia.com).

Parts of this paper have been presented at the Advanced Photonics Congress, Signal Processing in Photonic Communications (SPPCom), 2021 [1].

vised equalizer implementation has initially been presented in [23] and further extended in [24] towards unsupervised low-density parity-check (LDPC) decoding. However, only a simple quadrature PSK (QPSK) implementation was used. In this work, we generalize the approach of [23], [24] and propose essential extensions, including the application to oversampled dual-polarization (DP) signals and multi-amplitude PCS formats. We show that the generalized approach is independent of the equalizer architecture and train both a convolutional neural network (CNN) based equalizer as in [23], [24]—the VAE-NN—and a novel linear model-based equalizer with butterfly FIR filters—the VAE-LE. We evaluate the performance of the proposed equalizers on different linear channels and propose an extension for slowly time-varying channels.

This paper is structured as follows: in Sec. II, we introduce our system model, in Sec. III, we motivate variational inference for equalization, derive the VAE-based equalizer in a general form and explain how the proposed extensions can be incorporated. In Sec. IV, we discuss the implementation of the equalizers and propose an appropriate parameter update scheme, before we introduce our simulation environment and our results in Sec. V. We conclude the paper in Sec. VI.

II. SYSTEM MODEL

We start with the demonstration of the basic concept on a simple AWGN channel with ISI

$$\mathbf{y} = \mathbf{h}_{\text{sim}} * \mathbf{x} + \mathbf{n} ,$$

where the transmit vector \mathbf{x} is convolved with the simulated channel impulse response (IR) \mathbf{h}_{sim} and a white noise vector \mathbf{n} is added.

Then, we focus on a dispersive linear optical dual-polarization transmission to prove the concept in a practical environment (coherent optical transmission). We use the more natural description in the frequency domain by a linear channel matrix, i.e.,

$$\begin{bmatrix} \bar{\mathbf{y}}_{\text{TE}} \\ \bar{\mathbf{y}}_{\text{TM}} \end{bmatrix} = \mathbf{H}(f) \cdot \begin{bmatrix} \bar{\mathbf{x}}_{\text{TE}} \\ \bar{\mathbf{x}}_{\text{TM}} \end{bmatrix} + \begin{bmatrix} \bar{\mathbf{n}}_{\text{TE}} \\ \bar{\mathbf{n}}_{\text{TM}} \end{bmatrix} ,$$

with $\bar{\mathbf{z}} = \mathcal{F}\{\mathbf{z}\}$ being the Fourier transform of a time domain vector \mathbf{z} and {TE, TM} being indices describing the two polarizations of the light. Additionally, we emulate a slowly time-varying channel by changing the channel matrix' elements over time, so we can analyze the performance in a dynamical environment.

Per se, the optical channel is nonlinear, but the linear distortions are dominant in practical systems and have to be compensated, whereas the nonlinearity compensation is computationally demanding and usually provides a signal-to-noise ratio (SNR) gain of less than 1 dB [25, Sec. 6.9.3]. Hence, we focus in this paper on linear impairments and assumes potential nonlinearities to be either negligible or compensated by a separate digital signal processing (DSP) block, e.g., based on digital backpropagation, which can be switched on if required.

Further details on the simulation model including the specific parameters are provided later in Sec. V.

III. VARIATIONAL INFERENCE FOR EQUALIZATION

The goal of communications is to transmit data to a receiver, which has to fully recover the information without knowledge of the actually transmitted data. This can be interpreted as an inference problem, where the received samples $\mathbf{y} \in \mathbb{C}^N$ are observed variables while the transmitted samples $\mathbf{x} \in \mathbb{C}^N$ are unobservable latent variables. The optimum decision is based on the maximum of the a posteriori probability distribution [26, Ch. 4.1]

$$P(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x}) \cdot P(\mathbf{x})}{p(\mathbf{y})} ,$$

where $p(\mathbf{y}|\mathbf{x})$ is the likelihood of \mathbf{y} given \mathbf{x} , $P(\mathbf{x})$ is the prior probability and $p(\mathbf{y})$ is the observations' marginal density (also called the *evidence*). Throughout the paper, we denote probability mass functions (pmfs) by a capital $P(\cdot)$ and continuous densities by a lower case $p(\cdot)$. While the prior and the likelihood can usually be assumed as known or modeled well, the evidence is commonly intractable to compute, since the marginalization's complexity grows exponentially with the length and symbol order of \mathbf{x} , i.e. $p(\mathbf{y}) = \sum_{\mathbf{x}} P(\mathbf{x})p(\mathbf{y}|\mathbf{x})$.

In statistics, this is a common problem which can be solved by *variational inference*. It is also used in machine learning when the conditional has to be approximated efficiently and reliably [21]. In particular, this is the case in our problem where we usually require fast convergence to cope with time-dependent distortions. The main idea is to cast inference into an optimization problem, where the goal is to find an approximation $Q(\mathbf{x}|\mathbf{y}) \in \mathcal{Q}$ of the true a posteriori pmf $P(\mathbf{x}|\mathbf{y})$ from a family of approximate pmfs \mathcal{Q} , parameterized by free *variational parameters* [21], over the latent variables.

A. The Evidence Lower Bound (ELBO)

A suitable objective function is the relative entropy $D_{\text{KL}}(P||Q)$, also called the Kullback-Leibler (KL) divergence, which is an information-theoretical measure of proximity. It is asymmetric, non-negative and convex with its minimum at $Q = P$ [27, Ch. 2]. Then, the optimization's goal is to find the best approximation to the true a posteriori probability for the observed variables \mathbf{y} by

$$\hat{Q}(\mathbf{x}|\mathbf{y}) = \arg \min_{Q \in \mathcal{Q}} D_{\text{KL}}(Q(\mathbf{x}|\mathbf{y})||P(\mathbf{x}|\mathbf{y})) .$$

With $\mathbb{E}_Q\{\cdot\} = \mathbb{E}_{Q(\mathbf{x}|\mathbf{y})}\{\cdot\}$ being the expectation regarding the variational approximation, the KL divergence can be expressed as¹

$$\begin{aligned} D_{\text{KL}}(Q(\mathbf{x}|\mathbf{y})||P(\mathbf{x}|\mathbf{y})) &= \mathbb{E}_Q\{\ln Q(\mathbf{x}|\mathbf{y})\} - \mathbb{E}_Q\{\ln P(\mathbf{x}|\mathbf{y})\} \\ &= \mathbb{E}_Q\{\ln Q(\mathbf{x}|\mathbf{y})\} - \mathbb{E}_Q\{\ln p(\mathbf{y}|\mathbf{x})\} \\ &\quad - \mathbb{E}_Q\{\ln P(\mathbf{x})\} + \ln p(\mathbf{y}) . \end{aligned} \quad (1)$$

Since the KL divergence depends on $\ln p(\mathbf{y})$, it is not easily computable and thus not suitable as objective function. However, the evidence is independent of $Q(\mathbf{x}|\mathbf{y})$, so the last term

¹Note that we use the natural logarithm with base e (\ln) instead of the logarithm with base 2 (\log_2), which is frequently used in communications and information theory.

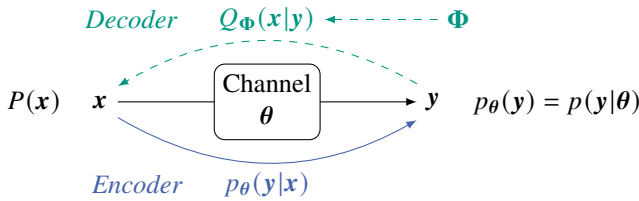


Fig. 1. Simple transmission system and its corresponding pdfs: \mathbf{x} and \mathbf{y} are realizations of random variable sequences $\{X_k\}$ and $\{Y_k\}$; θ is a discrete-time AWGN channel with impulse response \mathbf{h} and noise variance σ_w^2 ; $Q_{\Phi}(\mathbf{x}|\mathbf{y})$ is an approximation (in the parameter space Φ) of the real a posteriori pdf.

in (1) is only an additive constant regarding the optimization. Hence, maximizing the *evidence lower bound*

$$\begin{aligned} \text{ELBO}(Q) &= \mathbb{E}_Q\{\ln p(\mathbf{y}|\mathbf{x})\} + \mathbb{E}_Q\{\ln P(\mathbf{x})\} - \mathbb{E}_Q\{\ln Q(\mathbf{x}|\mathbf{y})\} \\ &= \underbrace{\mathbb{E}_Q\{\ln p(\mathbf{y}|\mathbf{x})\}}_{=:B} - \underbrace{D_{\text{KL}}(Q(\mathbf{x}|\mathbf{y})\|P(\mathbf{x}))}_{=:A} \end{aligned} \quad (2)$$

is equivalent to minimizing (1). This mirrors the usual balance between likelihood and prior, since the ELBO's first term B, the expected likelihood, favors densities which explain the observed data, while the second term A encourages the densities to be close to the prior [21]. The complexity of this optimization is defined by the complexity of the family Q .

By rearranging (1) and due to the KL divergence's non-negativity, we can show that the ELBO lower-bounds the (log) evidence, i.e.,

$$\begin{aligned} \ln p(\mathbf{y}) &= D_{\text{KL}}(Q(\mathbf{x}|\mathbf{y})\|P(\mathbf{x}|\mathbf{y})) + \text{ELBO}(Q) \\ &\geq \text{ELBO}(Q) . \end{aligned}$$

Since $\ln p(\mathbf{y})$ is a fixed upper bound, we sandwich $D_{\text{KL}}(Q(\mathbf{x}|\mathbf{y})\|P(\mathbf{x}|\mathbf{y}))$ by maximizing the ELBO, so we eventually minimize the KL divergence and find a good approximation $Q(\mathbf{x}|\mathbf{y})$ for the true a posteriori probability.

The concept can also be interpreted from a communication theory perspective with the likelihood $p(\mathbf{y}|\mathbf{x})$ as a probabilistic *encoder* and the a posteriori—or its variational approximation $Q(\mathbf{x}|\mathbf{y})$, respectively—as the corresponding *decoder*.² Precisely, during transmission, the data \mathbf{x} (latent variables) is encoded into the observable receive samples \mathbf{y} , while the receiver tries to decode the transmitted data again, estimating \mathbf{x} from \mathbf{y} . Assuming that the densities come from families of parametric distributions, the concept can be implemented as a *variational autoencoder* (VAE) using machine learning techniques, where typically both encoder and decoder are implemented as NNs [20], [24]. However, if we have a suitable model, e.g., of the encoder $p(\mathbf{y}|\mathbf{x})$, we do not have to apply an NN to learn it but we can use the model directly, as done in the following subsection.

B. The Variational Autoencoder (VAE)-based Equalizer

We assume a general transmission system through a (parameterized) channel θ , as depicted in Fig. 1. Then, the evidence

²In the deep learning literature, e.g., [20], an opposite definition is often found where $p(\mathbf{y}|\mathbf{x})$ is the decoder and $Q(\mathbf{x}|\mathbf{y})$ is the encoder. However, from a communications point-of-view we find our definition more intuitive.

$p_{\theta}(\mathbf{y})$ as well as the likelihood $p_{\theta}(\mathbf{y}|\mathbf{x})$ are parameterized by θ , while the variational approximation $Q_{\Phi}(\mathbf{x}|\mathbf{y})$ can be parameterized by a set of learnable parameters Φ , as denoted by the corresponding subscripts.

In other words, the channel (respectively the encoder) distorts the transmitted signal, while the decoder tries to find the mapping from the distorted received samples to the transmitted data. Thus, finding the optimum variational approximation corresponds to finding the optimized equalizer for this channel. Furthermore, $Q_{\Phi}(\mathbf{x}|\mathbf{y})$ gives a soft-decision on the received symbols, so the VAE-based equalizer also approximates an ML receiver [24].

Note that the exact values of θ are unknown, so the proposed equalization concept is blind and the model parameters are learned simultaneously with the decoder. In fact, the evidence $p_{\theta}(\mathbf{y}) = p(\mathbf{y}|\theta)$ can also be interpreted as the likelihood regarding the channel parameters, so variational inference also approximates maximum likelihood channel estimation. This byproduct can be used, e.g., for joint communication and sensing. See, e.g., [28] for an example of capturing acoustic signals based on the channel IR.

In the following, we derive the VAE-based equalizer in a generalized form compared to [24], where it is only derived for a toy model with QPSK. We try to keep repetitions as short as possible, but it is unavoidable at some points to highlight the generalizations we did. We start by assuming transmission over an AWGN channel parameterized by $\theta = (\mathbf{h}, \sigma_w^2)$ with finite IR \mathbf{h} and noise variance σ_w^2 . We can model the likelihood as

$$p_{\theta}(\mathbf{y}|\mathbf{x}) \sim \text{CN}(\mathbf{h} * \mathbf{x}, \sigma_w^2 \mathbf{I}_N) . \quad (3)$$

Further, we consider transmission of independently modulated square- M -QAM symbols x_i , so $\mathbf{x} = (x_1, \dots, x_N) = \mathbf{x}^I + \mathbf{j}\mathbf{x}^Q$ is a vector of N complex-valued symbols. Assuming further that I and Q have been modulated independently, than $x_i^I, x_i^Q \in \mathcal{A} = \{A_1, \dots, A_{\sqrt{M}}\}$ are conditionally independent given \mathbf{y} . Consequently, we can model

$$Q_{\Phi}(\mathbf{x}|\mathbf{y}) = \prod_{i=1}^N Q_{\Phi}(x_i^I|\mathbf{y}) \cdot Q_{\Phi}(x_i^Q|\mathbf{y}) \quad (4)$$

and define a vector $\mathbf{q}^{c, A_m}(\mathbf{y})$, $c \in \{I, Q\}$, containing the parametric functions evaluated for each symbol $\mathbf{q}^{c, A_m}(\mathbf{y}) := (Q_{\Phi}(x_1^c = A_m|\mathbf{y}), \dots, Q_{\Phi}(x_N^c = A_m|\mathbf{y}))$, which only depend on \mathbf{y} and Φ . Although we have a similar decoder model as [24], we consider multi-level signals and, thus, cannot simplify further by exploiting the normalization of probabilities.

Then, similarly to [24], we create a minimization problem by defining the loss function $\mathcal{L}(\theta, \Phi, \mathbf{y}) := -\text{ELBO}(Q) = A - B$ (see (2)), which only depends on both parameter spaces, θ and Φ , as well as the received samples \mathbf{y} .

The first term A can be easily computed by the standard formula of the KL divergence, i.e.,

$$\begin{aligned} A(\Phi) &= D_{\text{KL}}(Q_{\Phi}(\mathbf{x}|\mathbf{y})\|P(\mathbf{x})) \\ &= \sum_{i=1}^N \sum_{m=1}^{\sqrt{M}} \sum_{c \in \{I, Q\}} \left[q_i^{c, A_m}(\mathbf{y}) \cdot \ln \frac{q_i^{c, A_m}(\mathbf{y})}{P(x_i^c = A_m)} \right] . \end{aligned} \quad (5)$$

Since the term A does not break down to the entropy as in [24] (due to the assumption of a uniform prior pmf), an important feature is the inclusion of the prior density $P(\mathbf{x})$ into the loss function, which implies the adaption to, e.g., PCS [2], [3]. Since state-of-the-art blind equalizers struggle with non-uniform priors [15], [16], [29], simple inclusion is one of the major benefits of this concept.

The second term B can be re-written, similarly to [24], as

$$\begin{aligned} \mathbf{B}(\boldsymbol{\theta}, \boldsymbol{\Phi}, \mathbf{y}) &= \mathbb{E}_Q \{ \ln p_{\boldsymbol{\theta}}(\mathbf{y}|\mathbf{x}) \} \\ &= -N \ln(\pi \sigma_w^2) - \frac{1}{\sigma_w^2} \cdot \underbrace{\mathbb{E}_Q \{ \|\mathbf{y} - \mathbf{h} * \mathbf{x}\|^2 \}}_C, \quad (6) \end{aligned}$$

so B tries to find the best channel estimate \mathbf{h} regarding the least-squares of the observation \mathbf{y} to the prediction ($\mathbf{h} * \mathbf{x}$), also referred to as *autoencoder distortion* [24].

Further, with the operator $\text{Re}\{\cdot\}$ returning the real part of a complex number, $(\cdot)^H$ denoting a vector's conjugate-complex and transpose, $(\cdot)^T$ the transpose, and assuming that \mathbf{x} , \mathbf{h} and \mathbf{y} are column vectors, C becomes

$$\begin{aligned} \mathbf{C}(\boldsymbol{\theta}, \boldsymbol{\Phi}, \mathbf{y}) &= \mathbb{E}_Q \{ (\mathbf{y} - \mathbf{h} * \mathbf{x})^H (\mathbf{y} - \mathbf{h} * \mathbf{x}) \} \\ &= \|\mathbf{y}\|^2 - 2 \text{Re} \{ \mathbf{y}^H \cdot \mathbb{E}_Q \{ \mathbf{h} * \mathbf{x} \} \} + \mathbb{E}_Q \{ \|\mathbf{h} * \mathbf{x}\|^2 \} \\ &= \|\mathbf{y}\|^2 - 2 (\mathbf{y}^I)^T \mathbf{D}^I - 2 (\mathbf{y}^Q)^T \mathbf{D}^Q + \mathbf{E}, \quad (7) \end{aligned}$$

with

$$\begin{aligned} \mathbf{D}^I &= \mathbf{h}^I * \mathbb{E}_Q \{ \mathbf{x}^I \} - \mathbf{h}^Q * \mathbb{E}_Q \{ \mathbf{x}^Q \}, \\ \mathbf{D}^Q &= \mathbf{h}^I * \mathbb{E}_Q \{ \mathbf{x}^Q \} + \mathbf{h}^Q * \mathbb{E}_Q \{ \mathbf{x}^I \}, \\ \mathbf{E} &= \|\mathbf{D}^I\|^2 + \|\mathbf{D}^Q\|^2 + \\ &|\mathbf{h}|^2 * \left(\mathbb{E}_Q \left\{ (\mathbf{x}^I)^2 \right\} + \mathbb{E}_Q \left\{ (\mathbf{x}^Q)^2 \right\} - \mathbb{E}_Q \{ \mathbf{x}^I \}^2 - \mathbb{E}_Q \{ \mathbf{x}^Q \}^2 \right). \end{aligned}$$

In contrast to [24], we keep the vector notation, which helps to identify an efficient implementation, and we also have to compute the expectations for any M -QAM symbol from the VAE-based equalizer's output by ($c \in \{I, Q\}$)

$$\begin{aligned} \mathbb{E}_Q \{ \mathbf{x}^c \} &= \sum_{m=1}^{\sqrt{M}} \mathbf{q}^{c, A_m}(\mathbf{y}) A_m, \\ \mathbb{E}_Q \{ (\mathbf{x}^c)^2 \} &= \sum_{m=1}^{\sqrt{M}} \mathbf{q}^{c, A_m}(\mathbf{y}) A_m^2. \end{aligned}$$

Similarly to [24], we find an analytical solution for σ_w^2 by partially differentiating the loss function and equating it to zero. In fact, A (see (5)) does not depend on σ_w^2 and

$$\frac{\partial \mathbf{B}}{\partial \sigma_w^2} = -N \frac{1}{\sigma_w^2} + \frac{\mathbf{C}}{\sigma_w^4} \stackrel{!}{=} 0 \quad \Rightarrow \quad \sigma_w^2 = \frac{\mathbf{C}}{N}. \quad (8)$$

Hence, inserting (8) into the loss function and omitting all additive constants yields

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}, \boldsymbol{\Phi}, \mathbf{y}) := \text{D}_{\text{KL}}(Q_{\boldsymbol{\Phi}}(\mathbf{x}|\mathbf{y}) \| P(\mathbf{x})) + N \ln \mathbf{C}. \quad (9)$$

Here, we emphasize that the equalizer can also be designed for any integer oversampling factor N_{os} . If the equalizer incorporates downsampling, e.g., by convolution with stride N_{os} , all vectors can be defined accordingly. However, the

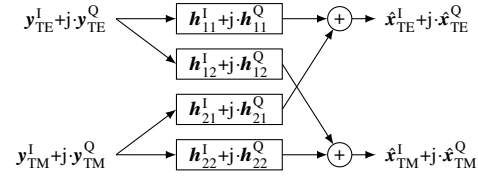


Fig. 2. Complex-valued 2×2 multiple-input multiple-output (MIMO)-system

size of the expectation vectors does not match the size of the observations \mathbf{y} anymore, so the term C would not be computable. Since the loss is summed over all samples, we can simply match the vectors by inserting $(N_{\text{os}} - 1)$ zeros between consecutive samples of $\mathbf{q}^{c, A_m}(\mathbf{y})$.

C. Extension Towards Coherent Optical Communication Systems

Light is always traveling as combination of two orthogonal polarizations, which can be independently modulated but rotate during propagation through a standard fiber. In combination with other effects like polarization mode dispersion (PMD) and chromatic dispersion (CD), the receiver observes a superposition of both polarizations similar to a classical multiple-input multiple-output (MIMO) channel with cross-talk. Although we focus on DP systems in this work, the proposed VAE-based equalization scheme can be extended towards any kind of MIMO system accordingly.

In the DP case, the transmitted and received sequences have to be represented by the matrices $\mathbf{X} = [\mathbf{x}_{\text{TE}} \ \mathbf{x}_{\text{TM}}]$ and $\mathbf{Y} = [\mathbf{y}_{\text{TE}} \ \mathbf{y}_{\text{TM}}]$, respectively, where the indices denote the transversal electric (TE) and transversal magnetic (TM) polarizations of the laser beam.

Considering the orthogonality, we can model

$$\begin{aligned} Q_{\boldsymbol{\Phi}}(\mathbf{X} | \mathbf{Y}) &= Q_{\boldsymbol{\Phi}}(\mathbf{x}_{\text{TE}} | \mathbf{Y}) \cdot Q_{\boldsymbol{\Phi}}(\mathbf{x}_{\text{TM}} | \mathbf{Y}), \\ p_{\boldsymbol{\theta}}(\mathbf{Y} | \mathbf{X}) &= p_{\boldsymbol{\theta}}(\mathbf{y}_{\text{TE}} | \mathbf{X}) \cdot p_{\boldsymbol{\theta}}(\mathbf{y}_{\text{TM}} | \mathbf{X}), \\ P(\mathbf{X}) &= P(\mathbf{x}_{\text{TE}}) \cdot P(\mathbf{x}_{\text{TM}}), \end{aligned}$$

and derive the loss function similarly to the AWGN channel case, which yields

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}, \boldsymbol{\Phi}, \mathbf{Y}) = \tilde{\mathcal{L}}_{\text{TE}}(\boldsymbol{\theta}, \boldsymbol{\Phi}, \mathbf{Y}) + \tilde{\mathcal{L}}_{\text{TM}}(\boldsymbol{\theta}, \boldsymbol{\Phi}, \mathbf{Y}). \quad (10)$$

In principal, the losses per polarization are calculated similarly to the proposed case for a single channel. However, in order to incorporate cross-talk between the polarizations, the mean of the likelihood's circular-symmetric normal distribution is not anymore ($\mathbf{h} * \mathbf{x}$) (see (3)), but depends on the superposition of both polarizations. In this work, we implement it as a complex-valued 2×2 MIMO-system as depicted in Fig. 2, which is based on a physical model [30], [31]. Alternatively, a real-valued 4×4 system can be implemented, which has additional degrees of freedom to compensate transceiver impairments.

IV. REALIZATION OF THE EQUALIZER

Commonly in the machine learning community, the VAE's *encoder* $p_{\boldsymbol{\theta}}(\mathbf{y}|\mathbf{x})$ and the *decoder* $Q_{\boldsymbol{\Phi}}(\mathbf{x}|\mathbf{y})$ are implemented as

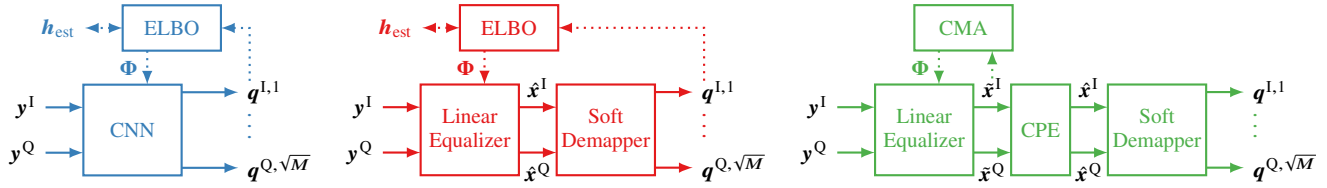


Fig. 3. Receivers considered in this paper for rectangular M -QAM formats (with PCS): VAE-NN (left), VAE-LE (middle); and reference CMA (right).

NNs with the parameters θ and Φ , but this is no requirement. In the application of the VAE concept to communications, the transmission channel is forming the encoder, while the decoder can be either an NN, as in in [24], or an FIR filter system with a soft-demapping block, as proposed in this work.

In comparison to the FIR filter system, the NN

- carries out a classification task and, hence, combines equalizer and demapper,
- is potentially capable of compensating non-linearities,
- requires more learnable parameters since its dimensionality depends on the modulation order,
- comprises more hyperparameters which have to be tuned,
- provides no access to the output constellation (since it only outputs the approximations $q^{c,A_m}(y)$), which may prevent the inclusion into state-of-the-art DSP chains.

In Fig. 3, we show the block diagrams of the investigated adaptive equalizers. The VAE-NN employs a CNN with two one-dimensional convolutional layers as in [24]. Adaptions are necessary to the final layer, namely exchanging the sigmoid with a softmax, to transform the multilevel output into probabilities, and we apply an exponential linear unit (ELU) instead of a softsign activation to the first layer, which provides better results. If applicable, a stride in the final layer downsamples the output. We found that the second layer's kernel size can be fixed to a small value (3 to 5), while the first layer's kernel size remains a hyperparameter. Precisely, the first layer's kernel size and the length of the estimated channel impulse response can be any odd integer (to ensure symmetry around a major central tap), which we optimized during our simulations. Typical values for both have been around 29 (for $N_{os} = 2$) and 11 (for $N_{os} = 1$). All equalizers and the simulation environment are implemented in Python with the PyTorch library.

A. The Proposed VAE-LE Scheme

We further propose the VAE-LE scheme, which is based on a classical 2×2 butterfly equalizer system with complex-valued FIR filters as depicted in Fig. 2. It uses the same filter system with F taps per filter as the reference CMA, which allows the integration into state-of-the-art DSP chains [31]. Note, however, that the computation of the cost function may not be as simple as for the CMA and needs to be adapted as it requires a soft-demapper output and backpropagation through the latter as well as possibly some further DSP blocks. Future work may take into account backpropagation through

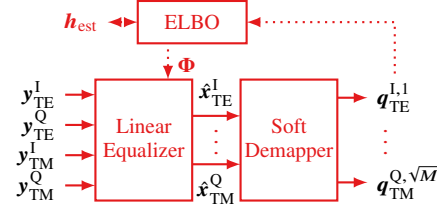


Fig. 4. Structure of the VAE-LE for DP systems—the other equalizers are implemented accordingly.

consecutive, possibly non-differentiable DSP blocks (see, e.g., [32] for an example) or the complexity reduction of the update algorithm. Precisely, the VAE-LE uses one filter system for the equalization, where its weights are Φ in the derivation above, and a second similar filter system as channel model for the estimation. Both parameter sets are independent, but they are trained simultaneously. We also initialize both similarly, i.e., we only initialize the real part of the filters h_{ii} with a 1 at the center tap, while all other taps (including the imaginary parts) are zero (Dirac initialization). Similarly to the CMA, if used along with soft-decision FEC, the VAE-LE also requires a soft-demapping block to transform the constellation output into the variational approximations. The implemented structure of the VAE-LE for a DP system is displayed in Fig. 4, which is representative for the implementation of the other equalizers.

We assume the general case of squared- M -QAM transmission, where the modulation symbols' prior pmf follows a Maxwell-Boltzmann distribution with normalization constant C_ν and shaping parameter $\nu \geq 0$, i.e.,

$$p_{MB}(x_i^c) = C_\nu e^{-\nu(x_i^c)^2}, \quad c \in \{I, Q\},$$

with $p_{MB}(x_i) = p_{MB}(x_i^I)p_{MB}(x_i^Q)$. For $\nu = 0$, $p_{MB}(x_i^c)$ becomes a uniform distribution. The Maxwell-Boltzmann distribution is the preferred distribution for PCS [2], [3], so we can cover both the case of uniform and PCS-QAM transmission by considering this prior pmf. We characterize the different formats using the constellation entropy

$$\mathcal{H} = -\mathbb{E}_{p_{MB}}\{\log_2 p_{MB}(x_i)\} \stackrel{(i)}{\leq} \log_2 M,$$

with equality in (i) if $\nu = 0$ (uniform distribution). Since optimum demapping is based on the maximum a posteriori (MAP) criterion [26, Ch. 4.1], we want to find

$$\begin{aligned} x_{dec,i}^c &= \arg \max_{x_i^c \in \mathcal{X}^c} p(x_i^c | y_i^c) \\ &= \arg \max_{x_i^c \in \mathcal{X}^c} p(y_i^c | x_i^c) p(x_i^c) \end{aligned}$$

$$= \arg \max_{x_i^c \in \mathcal{X}^c} \left(-\frac{(y_i^c - x_i^c)^2}{2\sigma_w^2} - \nu (x_i^c)^2 \right)$$

for a Gaussian likelihood [4]. Translating this to *soft*-demapping and adapting it to the M -QAM case with $x_i^c \in \mathcal{A} = \{A_1, \dots, A_{\sqrt{M}}\}$ yields

$$\mathbf{q}_i^c(\mathbf{y}) = \text{softmax} \left(-\frac{(\hat{x}_i^c \cdot \mathbf{1}_{\sqrt{M}} - \mathbf{A})^2}{2\sigma_w^2} - \nu \mathbf{A}_{\text{sq}} \right),$$

where $\mathbf{q}_i^c(\mathbf{y}) = (q_i^{c,A_1}, \dots, q_i^{c,A_{\sqrt{M}}})$, \hat{x}_i^c is the equalizer's output constellation (real or imaginary part) at time instant i , $\mathbf{A} = (A_1, \dots, A_{\sqrt{M}})$, $\mathbf{A}_{\text{sq}} = (A_1^2, \dots, A_{\sqrt{M}}^2)$, $\mathbf{1}_{\sqrt{M}}$ is a vector of \sqrt{M} ones, and $\text{softmax}(\mathbf{z}) = \exp(\mathbf{z}) / (\sum_{i=1}^L \exp(z_i))$ with $\mathbf{z} \in \mathbb{R}^L$.

B. Parameter Update Schemes

Next, we show how to adapt the VAE's training procedure to resemble the online-update of classical gradient-descent-based equalizers (such as the CMA) and to enable tracking of time-varying channels. Instead of separating the dataset into a training, test, and validation set as in classical supervised machine learning systems, we can directly train on the same data which we evaluate. Therefore, we continuously buffer the received data stream, slice it into consecutive mini-batches (of length $N_B \cdot N_{\text{os}}$ samples) and feed them to the equalizer. Then, an Adam optimizer [33] constantly updates the weights after each mini-batch. For the VAE-LE, the equalizer outputs all N_B equalized symbols and the soft-demapper translates them into the corresponding variational approximations, while the VAE-NN's CNN directly outputs the N_B approximated probabilities.

For both, we start the slicing of the next batch at the end of the former, so there are consecutive slices without any gap or overlapping in between. However, there is no requirement for having no overlapping, so we can also start the next slice only $N_{\text{flex}} \cdot N_{\text{os}}$ (instead of $N_B \cdot N_{\text{os}}$) samples after the start of the former slice and, thus, reduce the equalizer output from N_B to N_{flex} symbols. This results in an overlap of $(N_B - N_{\text{flex}}) \cdot N_{\text{os}}$ samples. Hence, each sample is considered for $n = \lfloor N_B / N_{\text{flex}} \rfloor$ consecutive update steps and, if (N_B / N_{flex}) is no integer, some samples are also considered for a further update step. This boosts convergence speed by sacrificing computational complexity due to more frequent updates. We call this generalized implementation *VAEflex* and introduce it for the evaluation of the time-varying channel. Throughout this work, we focus on adaptive channel equalization assuming an infinitely long random data sequence. Hence, we do not have to worry about overfitting during training for a specific channel and data sequence, although time-varying effects require a continuous re-adaptation of the filters.

The batch-wise update incorporates an implicit averaging of the loss over $N_B \cdot N_{\text{os}}$ samples, while the CMA updates the filter taps after each processed symbol by gradient descent as in its standard implementation [7], [8]. We observe that the VAE-LE with batch-wise training, which only updates every N_B symbols, had a significantly shorter computation time on

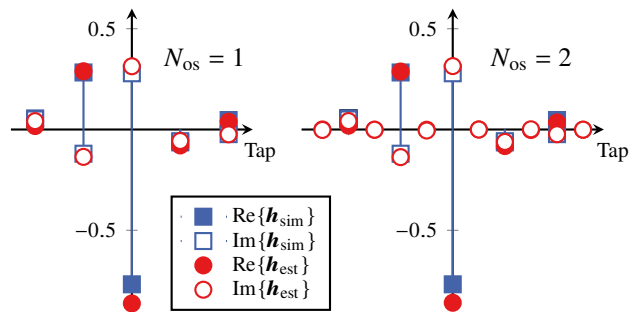


Fig. 5. Simulated channel IR \mathbf{h}_{sim} (similar to [23], [24]) and estimations \mathbf{h}_{est} by the VAE-LE at 20 dB for $N_{\text{os}} = 1$ and $N_{\text{os}} = 2$ sps (without pulse shaping).

a standard laptop's CPU as the CMA with its symbol-wise update.

The constant modulus criterion is not phase sensitive, so the CMA can only equalize the amplitude and requires a consecutive CPE stage, which we implement using the Viterbi-Viterbi algorithm [12] with averaging over 501 symbols. This CPE stage has, to the best of our knowledge, not been considered in [23], [24] (or was insufficient), which could explain the severe observed symbol error rate (SER) penalties in their QPSK simulations. In fact, we would expect significantly better results for a PSK-transmission over an AWGN channel, for which the CMA is well suited, especially if sufficiently long data sequences are used.

V. RESULTS

We evaluated the proposed equalizers both for simulations of a simple AWGN channel with ISI and an optical DP transmission system as well as a time-varying channel as introduced in Sec. II. The source code is available online [34].

A. Simulation Environment

Unless stated differently, our transmitter model consists of a source, which outputs a random sequence of M -QAM symbols, and a root-raised cosine (RRC) pulse-shaping filter with roll-off $\alpha = 0.1$. We use an oversampling factor of $N_{\text{os}} = 2$ sps throughout our simulations and incorporate downsampling to the equalizer. Since we omit matched filtering (see [26, Ch. 9]) but expect the equalizers to learn it, the receiver faces ISI from the RRC pulse-shaping additionally to the simulated channel.

In the AWGN channel, we convolve the transmitted sequence with the complex-valued channel IR

$$\mathbf{h}_{\text{sim}} = [0.055 + j0.05, 0.283 - j0.120, -0.768 + j0.279, -0.064 - j0.058, 0.047 - j0.023]$$

already used in [23], [24] (also shown in Fig 5), which we oversample by inserting $(N_{\text{os}} - 1)$ zeros between consecutive samples and interpolate by convolving it with the RRC pulse. We add real-valued AWGN with a variance of $\sigma_w^2/2$ on both real and imaginary part and take the oversampling into account for the SNR calculation.

TABLE I
SIMULATION PARAMETERS

Parameter	Value	Parameter	Value
γ_{hv}	0.1π	φ_{IQ}	0.01π
D_{pmd}	$0.1 \text{ ps}/\sqrt{\text{km}}$ [35]	L_{pmd}	1000 km
β_{cd}^1	$-26 \text{ ps}^2/\text{km}$	L_{cd}	1 km

¹ equals $D_{cd} = 20 \text{ ps nm}^{-1} \text{ km}^{-1}$ at $\lambda = 1550 \text{ nm}$ (see [37, Eq. (2.3.5)])

Monte-Carlo simulations of an AWGN channel without ISI provide a baseline SER per SNR and modulation format, which we denote by “No ISI” in the result plots.

Our optical DP transmission model follows [35] and [36] and includes a static (input and output) IQ-phase-shift φ_{IQ} and a static rotation of the reference polarization to the fiber's principal state of polarization (PSP), called HV-phase-shift γ_{hv} . Additionally, we simulate both first-order PMD caused by the differential group delay $\tau_{pmd} = D_{pmd}\sqrt{L_{pmd}}$ between the PSPs over a fiber length L_{pmd} , and residual CD, which is defined by the fiber's group velocity dispersion (GVD) parameter β_{cd} [37, Ch. 2.3] times the uncompensated fiber length L_{cd} . We display all parameter values in Table I.

By assuming—similarly to [35]—that all non-linear effects and phase noise are either negligible or compensated, and that we transmit a complex-valued DP signal $\mathbf{x} = [x_{TE}(t) \ x_{TM}(t)]$, we can model the fiber by the linear time-invariant frequency domain channel matrix

$$\mathbf{H}(f) = \mathbf{R}^T \begin{pmatrix} e^{j\pi\tau_{pmd}f} & 0 \\ 0 & e^{-j\pi\tau_{pmd}f} \end{pmatrix} \mathbf{R} \cdot e^{-j2\pi^2\beta_{cd}L_{fiber}f^2}$$

where $\mathbf{R} = e^{-j\varphi_{IQ}} \begin{pmatrix} \cos \gamma_{hv} & \sin \gamma_{hv} \\ -\sin \gamma_{hv} & \cos \gamma_{hv} \end{pmatrix}$. Again, we add AWGN on both polarizations and I/Q-components with a variance of $\sigma_w^2/2$ each. To simulate a time-varying channel, we change the HV-phase-shift after each frame of $N_{frame} = 10,000$ symbols. Precisely, we extend $\tilde{\gamma}_{hv} = \gamma_{hv} + \Delta\gamma_{hv} \cdot kT_{frame}$ with the deviation of the HV-shift $\Delta\gamma_{hv}$, the frame duration $T_{frame} = N_{frame}/R_S$, the symbol rate R_S and the frame-index k as an integer indexing the discrete time (in multiples of the frame duration). It should be highlighted that we apply deviations frame-wise, i.e., we neglect deviations within each frame of N_{frame} symbols.

For evaluation, we chose the averaging scheme sketched in Fig. 6, which might not seem to be straight-forward at first glance but fits to our needs, i.e., it returns a reliable SER estimate to compare different hyperparameter settings but also provides insights into the convergence behavior. In particular, the analysis of the average convergence behavior requires a certain temporal averaging, but the final analysis requires also an averaging over multiple execution runs, since the behavior could vary between different runs.

To get the desired insights, we first have to slice the equalizer output z_j^{pol} (per simulation run j and per polarization $pol \in \{TE, TM\}$) into $N_{ind} = 170$ consecutive frames $\mathbf{z}_{j,k}^{pol} \in \mathbb{C}^{N_{frame}}$ to allow the evaluation at different time steps, respectively frame-indices $k = 1 \dots N_{ind}$. Then, we estimate each frame's scalar $\widehat{SER}_{j,k}^{pol}$ after hard decision taking into account the

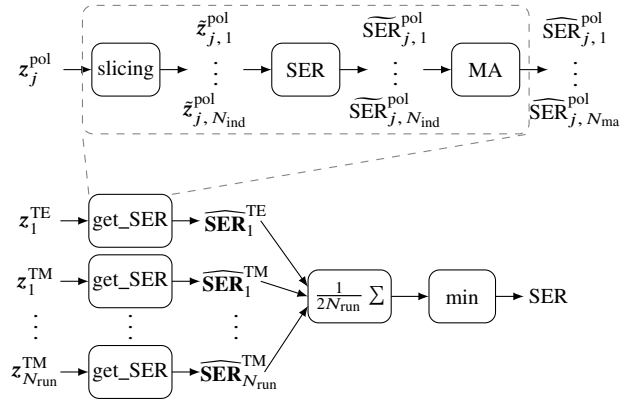


Fig. 6. Block diagram sketching the result averaging— $\mathbf{z}_{j,k}^{pol} \in \mathbb{C}^{N_{frame}}$ are the slices (frames) of the vector containing the corresponding equalizer output with $pol \in \{TE, TM\}$, MA is a moving average filter with length $F_{ma} = 10$, so $N_{ma} = (N_{ind} - F_{ma} + 1)$, and min is an operator which returns a vector's element with the minimum value.

prior distribution of the constellation symbols [4]. Since we perform blind equalization, we also need to compensate for possible time-shifts, I/Q-flips and phase-rotations (in multiples of $\frac{\pi}{4}$). Furthermore, we discard symbols that may be incorrect due to boundary effects (but having at least 8,000 symbols per frame remaining for evaluation—hence, we can evaluate the potential performance of each algorithm). Eventually, we perform a moving average with filter length $F_{ma} = 10$ and get a sequence of $N_{ma} = N_{ind} - F_{ma} + 1$ estimates $\widehat{SER}_{j,k}^{pol}$, each being evaluated on approximately 80,000 ... 100,000 symbols per frame-index k (per polarization and per run). This allows the analysis of the equalizers' convergence for each polarization and run independently.

For further evaluations, we need one reliable scalar SER estimate for each algorithm and hyperparameter setting. Thus, we carry out $N_{run} = 10$ (unless stated otherwise) independent simulation runs and average the results over all successful runs, i.e., runs where the SER drops below a pre-defined threshold of 0.3, by taking the element-wise mean of the vectors $\widehat{SER}_j^{pol} \in \mathbb{R}^{N_{ma}}$, which contain the estimates per run and polarization after moving average. If unsuccessful runs occurred, we display the amount by a small number next to the corresponding data point in the result figures. Finally, we get a reliable scalar SER estimate by taking the element with the minimum value from the vector containing the averaged estimates for all remaining N_{ma} frame-indices. We chose the minimum over the mean to display the best (averaged) performance of the equalizers, since we already have a sufficient averaging (the mean of $N_{pol} \cdot N_{run}$ estimates where each is averaged over at least 80,000 symbols), but can prevent distortions from potential outliers. In fact, we did not observe any outliers in all our simulations, so we conjecture that the estimates' variance is small and the difference between taking the minimum and the mean (after convergence) is insignificant.

Furthermore, we use a learning rate scheduler for the optical DP channel which halves the learning rate ϵ_{lr} of each equalizer

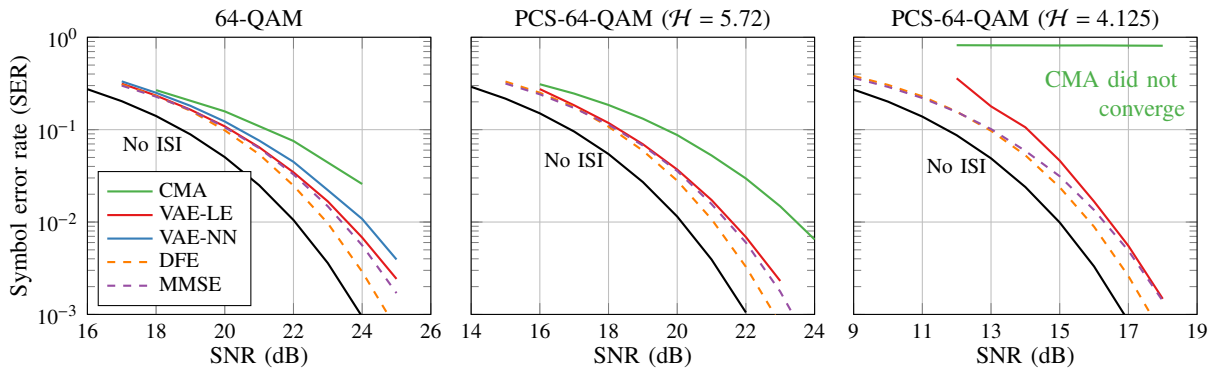


Fig. 7. Results for uniform and PCS-64-QAM through an AWGN channel with ISI and without learning rate reduction—see also [1]; the simulations for the *non-blind* DFE and MMSE are without pulse-shaping at $N_{os} = 1$ sps.

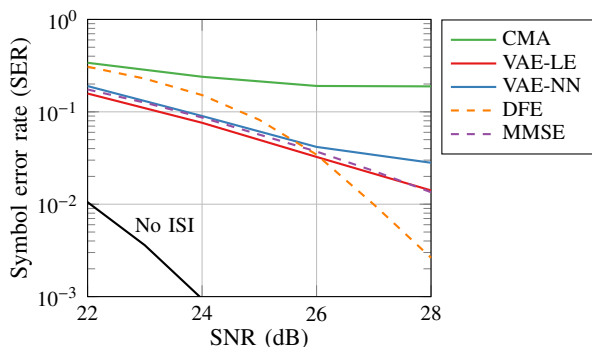


Fig. 8. Results for uniform 64-QAM through an AWGN channel with ISI and without learning rate reduction—see also [1]; $\mathbf{h}_{sim,2} = [0.055 + j0.017, -1.345 - j0.452, 1.007 + j1.152, 0.348 + j0.315]$; the simulations for the *non-blind* DFE and MMSE are without pulse-shaping at $N_{os} = 1$ sps.

after every frame-index k being an integer-multiple of 20, so the given ϵ_r represents the initial value.

B. Simple AWGN Channel with ISI

Figure 7 depicts the proposed equalizers’ performance for both uniform and PCS-64-QAM at $N_{os} = 2$ sps. We tuned the lengths of the filters F , the batches N_B , the kernels, and the learning rate. Later, for the DP simulations, we will show how the hyperparameters influence the equalizers’ performance. Besides the baseline “No ISI”-curve and the CMA as reference, we also evaluated two *non-blind* equalizers at $N_{os} = 1$ sps without pulse-shaping, namely the non-linear decision feedback equalizer (DFE) [26, Ch. 9.5] (with 10 taps each for both feed-forward and feedback filters) and the linear minimum mean squared error (MMSE) equalizer [26, Ch. 9.4] (with 20 taps).

For uniform 64-QAM, the CMA converges with a significant penalty to the MMSE, while both VAE-based equalizers operate close to the MMSE. For a linear channel, the VAE-LE outperforms the VAE-NN, which might originate from a closer bounding of the family \mathcal{Q} around the optimum $\hat{\mathcal{Q}}(x|y) \in \mathcal{Q}$. We found qualitatively similar results for other ISI channels as well, e.g., in Fig. 8 for channel \mathbf{h}_2 from [24]. For the

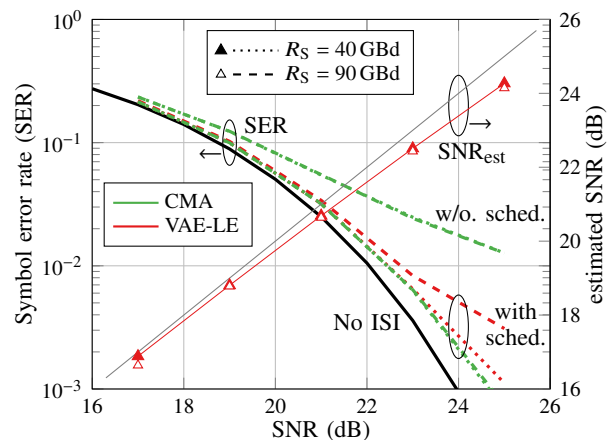


Fig. 9. Performance of the VAE-LE and the CMA for uniform DP-64-QAM with $\epsilon_r = 0.5 \cdot 10^{-3}$, $F = 25$ and $N_B = 350$: SER, left axis, and estimated SNR, right axis, vs SNR for $R_S = 40$ GBd and $R_S = 90$ GBd. The CMA is evaluated with ($\epsilon_r = 0.5 \cdot 10^{-3}$) and without (w/o. - $\epsilon_r = 5 \cdot 10^{-5}$) learning rate scheduler.

sake of computational complexity, we restrict ourselves in the following to the VAE-LE.

As already forecasted in the early 1990s [15], [16], the CMA fails to converge for PCS formats which approximate a Gaussian prior; however, the VAE-LE even reaches the *non-blind* MMSE’s performance for higher SNRs, although we used the same soft-demapper as for the uniform QAM in these AWGN channel simulations (instead of the optimal PCS-adapted version as introduced in Sec. IV). Furthermore, Fig. 5 demonstrates the VAE-LE’s capability of estimating the channel IR. We depict the estimate without averaging while processing uniform 64-QAM at SNR = 20 dB for 1 and 2 sps without pulse-shaping.

C. Optical Dual-polarization Transmission

The results for the application of the VAE-LE are depicted in Fig. 9. The right axis shows the estimated SNR (averaged similarly to the SER) while evaluating DP-64-QAM at various SNRs. The VAE-LE always underestimates the SNR, but only within a fraction of a decibel (dB). Along with the results

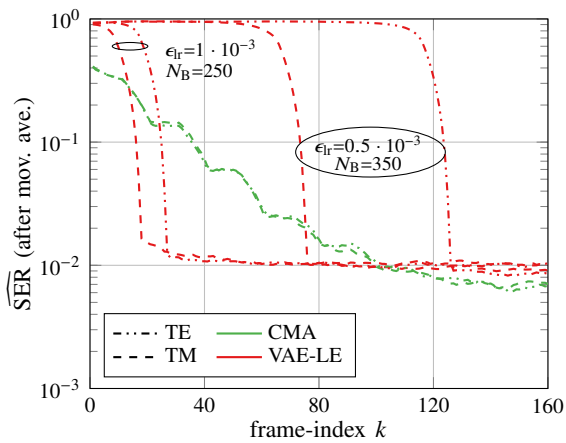


Fig. 10. Convergence behavior of the VAE-LE and the CMA ($\epsilon_{lr} = 0.5 \cdot 10^{-3}$) for uniform DP-64-QAM with $R_S = 90$ GBd, 23 dB and $F = 25$.

of Fig. 5, this demonstrates the potential for joint communications and sensing, since the estimation of the channel parameters is a valuable byproduct for tracking and evaluating the channel IR and SNR without interfering communications.

The left axis in Fig. 9 depicts the corresponding SERs. Similar to the results in Fig. 7, the CMA has a significant penalty without learning rate scheduler. Since the scheduler halves the learning rate after every frame-index k being an integer-multiple of 20, the CMA’s performance-convergence trade-off is avoided and it reaches marginally lower SERs as the VAE-LE. For high symbol rates and high SNRs, the VAE-LE deviates from the ideal “No ISI”-curve. In all other cases, both equalizers stay within 1 dB of the “No ISI”-curve and converge to it for low SNRs.

The equalizers’ convergence behavior differs significantly, as depicted in Fig. 10, where we display one simulation run’s N_{ma} SER estimates after the moving average filter. It should be noted that we do not consider a fixed amount of training data but conduct online-learning on the (in theory, infinitely long) received data sequence and that the frame-index corresponds to discrete time steps during training, i.e., one frame has a duration of T_{frame} and corresponds to $N_{os} \cdot N_{frame}$ samples processed by the equalizer. The CMA converges gradually with a decreasing slope. The distinct “plateaus” are caused by the learning rate scheduler. The VAE-LE shows a “waterfall-like” curve. Also, the VAE-LE starts at a significantly higher SER. Although the initial learning rate and batch size strongly influence the VAE-LE’s time convergence, the curve’s waterfall-like shape remains. Interestingly, the CMA seems to optimize both polarizations equally, while the VAE-LE first focuses on one polarization before trailing the second one.

The influence of the main hyperparameters is analyzed in Fig. 11. Interestingly, the symbol rate R_S and the filter length F influence the CMA significantly less than the VAE-LE, which suffers from performance penalties for low F . The reason might be that the VAE-LE has to compensate for both amplitude and phase offset while an extra CPE compensates the latter for the CMA. Since a high symbol rate mainly increases the ISI, it is reasonable that it only influences the

VAE-LE for high SNR.

The VAE-LE suffers from convergence issues for high F and high N_B (or ϵ_{lr}), but is very stable for small F and small N_B even under high symbol rates. While changes in the initial learning rate affect the CMA strongly, the VAE-LE has a negligible penalty over a relatively large range.

We show the results for PCS in a DP optical channel in Fig. 12. With the adapted soft-demapping and the learning rate scheduler, the VAE-LE is able to follow the “No ISI”-curve within 1 dB penalty even for high symbol rates and low SNR. Lower filter lengths F and batch-sizes N_B as well as higher ϵ_{lr} are necessary to ensure a high probability for convergence, especially for high symbol rates. The VAE-LE is potentially capable of converging for $R_S > 100$ GBd and strong shaping in our simulation model, but the probability for non-convergence is relatively high for typical working points. Still, the VAE-LE significantly outperforms the CMA for PCS, where the latter does not converge at all.

D. Time-varying Channel

We analyze the influence of time-dependent channel deviations to the equalizers by a frame-wise increasing HV-shift $\tilde{\gamma}_{hv} = \gamma_{hv} + \Delta\gamma_{hv} \cdot kT_{frame}$ within the optical DP model. Figure 13 depicts the SER for different slopes $\Delta\gamma_{hv}$. Since we did not employ the learning rate scheduler in this evaluation, the hyperparameter values differ from the ones used in the time-invariant case. Due to its rather slow convergence speed, the CMA has to operate at rather high learning rates and entails a severe penalty. The wide range of possible learning rates allows the VAE-LE to optimize its working point towards low SERs for low $\Delta\gamma_{hv}$ or a high tolerance towards deviations by moderate penalties. The VAEflex with $N_B = 100$ and $N_{flex} = 10$ accelerates training significantly, which makes it tolerant towards deviations by still reaching very low SERs. Although we did not optimize the VAEflex as thoroughly as the other algorithms, it converges until a low (respectively, early) frame-index k at the cost of a higher computational complexity. Hence, an option would be to switch between the VAEflex and the batch-wise VAE during operation by changing N_{flex} between 1 and N_B .

For comparison, we also implemented a CMA with a batch-wise updating scheme as proposed in [38], which we denote *CMAbatch*. Additionally, we also extended this scheme with a flexible update rule akin to the VAEflex. We denote the resulting equalizer *CMAflex*. Although Fig. 13 shows that both the *CMAbatch* and the *CMAflex* perform better than the symbol-wise CMA for this time-varying channel and without the learning rate scheduler, the gain is relatively small and both are outperformed by the VAE-based equalizers. Especially the *CMAflex* performs very similar to the *CMAbatch* and, in contrast to the VAEflex, the flexible update rule is not able to accelerate training.

VI. CONCLUSION

In this paper, we proposed the new VAE-LE, a model-based approach with linear butterfly FIR filters, which is trained by the VAE-based learning paradigm with extensions

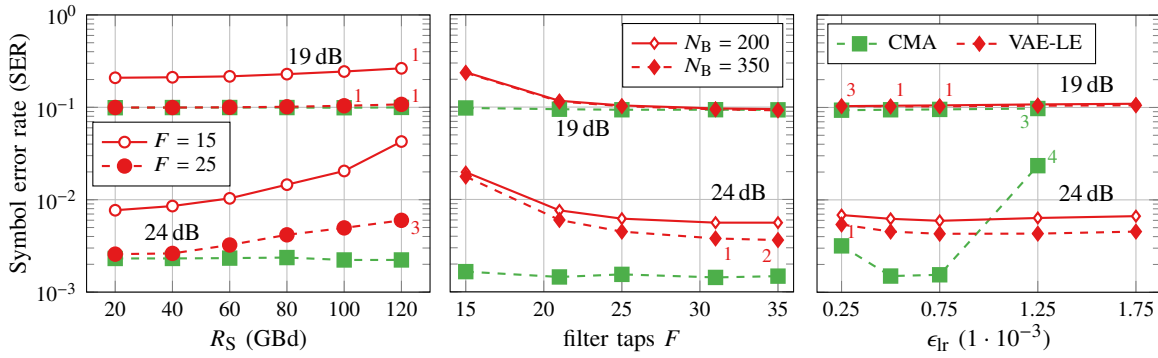


Fig. 11. Hyperparameter analysis of the VAE-LE (red, —) and the CMA (green, —) for uniform DP-64-QAM with $R_S = 90$ GBd and the default values of $F = 25$, $\epsilon_r = 0.5 \cdot 10^{-3}$ and $N_B = 350$.

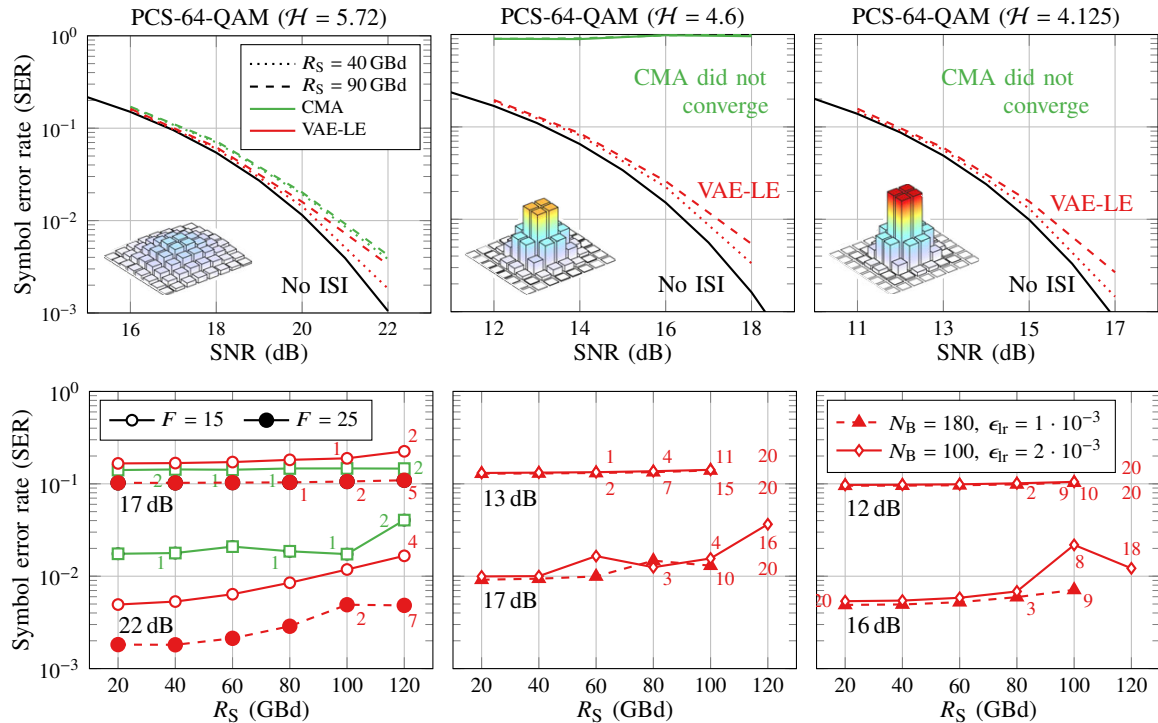


Fig. 12. Results for DP-PCS-64-QAM with VAE-LE (red, —) and CMA (green, —); for $\mathcal{H} = 5.72$: $\epsilon_r = 0.5 \cdot 10^{-3}$ and $N_B = 350$; for $\mathcal{H} = 4.6$ and $\mathcal{H} = 4.125$: $\epsilon_r = 1 \cdot 10^{-3}$ (VAE-LE), $N_B = 180$, $F = 15$, and $N_{run} = 20$ for better statistics.

towards practically relevant optical communication systems with oversampling and DP-PCS-QAM transmission.

For AWGN channels with ISI, the *blind* VAE-based equalizers can approach the performance of the *non-blind* MMSE equalizer for both uniform and PCS formats. The proposed VAE-LE outperforms the previously introduced VAE-NN for this linear channel. The VAE-LE equalizer also converges in a dispersive optical DP system but shows a negligible penalty to the CMA for uniform formats. However, for PCS formats where the CMA fails to converge without modifications, the VAE-LE still approaches the ideal reference within 1 dB.

The VAE-LE's rapid convergence behavior is advantageous for time-varying channels, where the gradually converging CMA performs significantly worse. Our proposed VAEflex

update scheme with flexible step length is a powerful alternative if convergence speed is a key factor. Additionally, we have shown that the VAE-LE is able to estimate both the communication channel taps and the noise variance very well, which can be an enabler for joint communications and sensing.

While we focused on linear channels in this work, the extension towards nonlinear impairments might be a possible direction of future research.

ACKNOWLEDGMENT

The authors would like to thank Luca Schmid for providing the code for the DFE and MMSE equalizers.

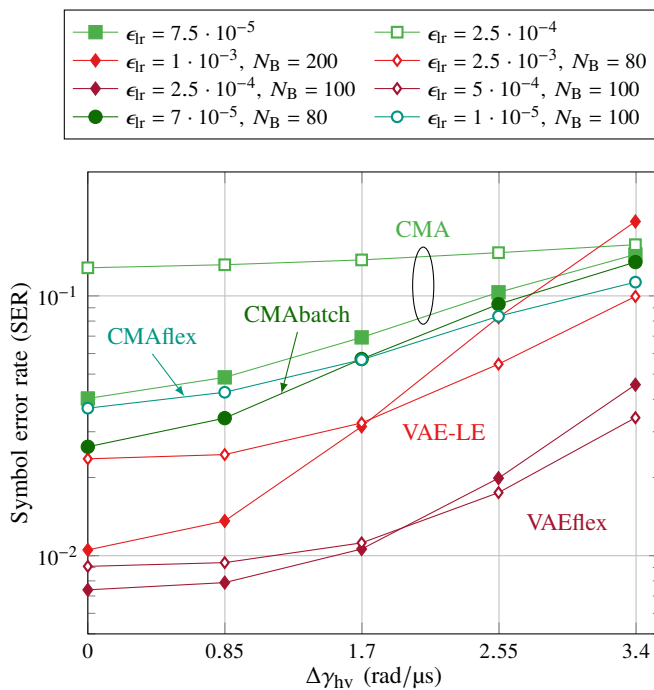


Fig. 13. Results for uniform DP-64-QAM transmitted over a time-varying optical DP channel with SNR = 23 dB, $R_S = 90$ GBd and $F = 25$ (here, no learning rate scheduler is used): CMA, CMAbatch, VAE-LE as well as CMAflex and VAEflex with $N_{flex} = 10$. Note that all deviations are simulated frame-wise with $T_{frame} = 0.111 \mu s$.

REFERENCES

- [1] V. Lauinger, F. Buchali, and L. Schmalen, "Blind equalization for coherent optical communications based on variational inference," in *Proc. Advanc. Photon. Cong.: Sign. Process. Photon. Commun. (SPPCom)*. Opt. Soc. America (OSA), 2021, pp. SpTh1D-6.
- [2] G. Böcherer, F. Steiner, and P. Schulte, "Bandwidth efficient and rate-matched low-density parity-check coded modulation," *IEEE Trans. Commun.*, vol. 63, no. 12, pp. 4651–4665, 2015.
- [3] F. Buchali, F. Steiner, G. Böcherer, L. Schmalen, P. Schulte, and W. Idler, "Rate adaptation and reach increase by probabilistically shaped 64-QAM: An experimental demonstration," *J. Lightw. Technol.*, vol. 34, no. 7, pp. 1599–1609, 2016.
- [4] J. Cho and P. J. Winzer, "Probabilistic constellation shaping for optical fiber communications," *J. Lightw. Technol.*, vol. 37, no. 6, pp. 1590–1607, 2019.
- [5] H. Sun *et al.*, "800G DSP ASIC design using probabilistic shaping and digital sub-carrier multiplexing," *J. Lightw. Technol.*, vol. 38, no. 17, pp. 4744–4756, Sep 2020.
- [6] F. Buchali, V. Lauinger, M. Chagnon, K. Schuh, and V. Aref, "CMOS DAC supported 1.1 Tb/s/λ DWDM transmission at 9.8 bit/s/Hz over DCI distances," *J. Lightw. Technol.*, vol. 39, no. 4, pp. 1171–1178, Oct. 2020.
- [7] 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.
- [8] R. Johnson, P. Schniter, T. J. Endres, J. D. Behm, D. R. Brown, and R. A. Casas, "Blind equalization using the constant modulus criterion: A review," *Proc. IEEE*, vol. 86, no. 10, pp. 1927–1950, 1998.
- [9] J. Yang, J.-J. Werner, and G. A. Dumont, "The multimodulus blind equalization and its generalized algorithms," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 5, pp. 997–1015, 2002.
- [10] C. You and D. Hong, "Nonlinear blind equalization schemes using complex-valued multilayer feedforward neural networks," *IEEE Trans. Neural Netw.*, vol. 9, no. 6, pp. 1442–1455, 1998.
- [11] T. Pfau, S. Hoffmann, and R. Noé, "Hardware-efficient coherent digital receiver concept with feedforward carrier recovery for M -QAM constellations," *J. Lightw. Technol.*, vol. 27, no. 8, pp. 989–999, 2009.

- [12] A. J. Viterbi and A. M. Viterbi, "Nonlinear estimation of PSK-modulated carrier phase with application to burst digital transmission," *IEEE Trans. Inf. Theory*, vol. 29, no. 4, pp. 543–551, 1983.
- [13] D. A. Mello, F. A. Barbosa, and J. D. Reis, "Interplay of probabilistic shaping and the blind phase search algorithm," *J. Lightw. Technol.*, vol. 36, no. 22, pp. 5096–5105, 2018.
- [14] Q. Zhang and C. Shu, "Viterbi and Viterbi algorithm based phase recovery for probabilistically shaped signals," *J. Lightw. Technol.*, vol. 39, no. 5, pp. 1364–1370, 2020.
- [15] C. K. Chan and J. J. Shynk, "Stationary points of the constant modulus algorithm for real Gaussian signals," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 12, pp. 2176–2181, 1990.
- [16] E. Zervas, J. G. Proakis, and V. Eyuboglu, "Effects of constellation shaping on blind equalization," in *Proc. Adapt. Sign. Process.*, vol. 1565. Internat. Soc. for Opt. and Photon. (SPIE), 1991, pp. 178 – 187.
- [17] M. Ghosh and C. L. Weber, "Maximum-likelihood blind equalization," *Opt. Engineer.*, vol. 31, no. 6, pp. 1224–1228, 1992.
- [18] L. Tong and S. Perreau, "Multichannel blind identification: From subspace to maximum likelihood methods," *Proc. IEEE*, vol. 86, no. 10, pp. 1951–1968, 1998.
- [19] G. K. Kaleb 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.
- [20] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. Internat. Conf. Learn. Represent. (ICLR)*, 2015.
- [21] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *J. Americ. Statist. Assoc.*, vol. 112, no. 518, pp. 859–877, 2017.
- [22] V. Raj and S. Kalyani, "Design of communication systems using deep learning: A variational inference perspective," *IEEE Trans. on Cogn. Commun. Netw.*, vol. 6, no. 4, pp. 1320–1334, 2020.
- [23] A. Caciularu and D. Burshtein, "Blind channel equalization using variational autoencoders," in *Proc. Internat. Conf. Commun. (ICC)*. IEEE, 2018, pp. 1–6.
- [24] —, "Unsupervised linear and nonlinear channel equalization and decoding using variational autoencoders," *IEEE Trans. on Cogn. Commun. Netw.*, vol. 6, no. 3, pp. 1003–1018, 2020.
- [25] B. Mukherjee, I. Tomkos, M. Tornatore, P. Winzer, and Y. Zhao, *Springer Handbook of Optical Networks*. Springer Nature, 2020.
- [26] J. Proakis and M. Salehi, *Digital Communications, 5th edition*. McGraw-Hill Higher Education, 2008.
- [27] T. M. Cover and J. A. Thomas, *Elements of Information Theory, 2nd edition*. John Wiley & Sons, 2006.
- [28] C. Dorize, S. Guerrier, E. Awwad, and J. Renaudier, "Capturing acoustic speech signals with coherent MIMO phase-OTDR," in *Proc. Eur. Conf. Opt. Commun. (ECOC)*, Dec 2020, pp. 1–4.
- [29] M. Arikawa, M. Sato, and K. Hayashi, "Wide range rate adaptation of QAM-based probabilistic constellation shaping using a fixed FEC with blind adaptive equalization," *Opt. Express*, vol. 28, no. 2, pp. 1300–1315, 2020.
- [30] R. Raheli and G. Picchi, "Synchronous and fractionally-spaced blind equalization in dually-polarized digital radio links," in *Proc. Internat. Conf. Commun. (ICC)*. IEEE, 1991, pp. 156–161.
- [31] S. J. Savory, "Digital coherent optical receivers: Algorithms and sub-systems," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 5, pp. 1164–1179, 2010.
- [32] A. Rode, B. Geiger, and L. Schmalen, "Geometric constellation shaping for phase-noise channels using a differentiable blind phase search," in *Proc. Opt. Fiber Commun. Conf. (OFC)*, 2022, p. Th2A.32.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *preprint arXiv:1412.6980*, 2014.
- [34] V. Lauinger, F. Buchali, and L. Schmalen, *Source code*, Available online: <https://github.com/kit-cel/vae-equalizer>, 2022, [Online].
- [35] E. Ip and J. M. Kahn, "Digital equalization of chromatic dispersion and polarization mode dispersion," *J. Lightw. Technol.*, vol. 25, no. 8, pp. 2033–2043, 2007.
- [36] S. J. Savory, "Digital filters for coherent optical receivers," *Opt. Express*, vol. 16, no. 2, pp. 804–817, 2008.
- [37] G. P. Agrawal, *Fiber-Optic Communication Systems, 4th edition*. John Wiley & Sons, 2010.
- [38] D. E. Crivelli *et al.*, "Architecture of a single-chip 50 Gb/s DP-QPSK/BPSK transceiver with electronic dispersion compensation for coherent optical channels," *IEEE Trans. Circuits Syst. I*, vol. 61, no. 4, pp. 1012–1025, Apr. 2014.



Vincent Lauinger (Student Member, IEEE) received both the B. Sc. ('17) and M. Sc. ('20) degree in electrical engineering and information technology from the Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany. During his Master studies, he put an emphasis on signal processing and optical communications. In 2019 and 2020, he was with Nokia Bell Labs, Stuttgart, working on digital pre-distortion for high-speed optical communication systems and on novel equalization techniques. Since 2020, he is research and teaching assistant at the

Communications Engineering Lab (CEL) of KIT, where he is pursuing his Ph.D. degree. His research interests include digital signal processing, machine learning in communications, and information theory.



Fred Buchali received the diploma degree in electrical engineering in 1988 and the Ph.D. degree in 1991 from Humboldt-University Berlin. In 1999 he joined the optical transmission systems group of Alcatel SEL Research and Innovation, where he contributed to optical and electrical impairment mitigation techniques at 10 and 40 Gb/s by analog electrical and optical techniques. During emerging coherent transmission technology he substantially contributed to the discovery of OFDM in optics. His current research interests cover coherent single carrier

systems towards 1 Tb/s using higher order modulation formats including probabilistic shaping. He published more than 100 journal and conference papers and holds more than 30 patent applications.



Laurent Schmalen (Senior Member, IEEE) received the Dipl.-Ing. degree in electrical engineering and information technology and the Dr.-Ing. Degree from the RWTH Aachen University of Technology, Aachen, Germany. From 2011 to 2019, he was with Bell Labs as a member of Technical Staff and also the Department Head of the Coding in Optical Communications Department from 2016 to 2019. From 2014 to 2019, he was a Guest Lecturer with the University of Stuttgart, Germany. He is currently a Full Professor with Karlsruhe Institute

of Technology (KIT), where he heads the Communications Engineering Lab (CEL). His research interests include forward error correction, modulation formats, and information theory for future optical networks. He was a recipient and a co-recipient of several awards, including the E-Plus Award for his Ph.D. thesis, the Best Paper Award of the 2010 ITG Speech Communication Conference, the 2013 Best Student Paper Award at the IEEE Signal Processing Systems Workshop, and both the 2016 and 2018 Journal of Lightwave Technology Best Paper Awards. He serves as an associate editor for the IEEE Transactions on Communications.