

End-to-end Optimization of Constellation Shaping for Wiener Phase Noise Channels with a Differentiable Blind Phase Search

Andrej Rode, *Graduate Student Member, IEEE*, Benedikt Geiger, Shrinivas Chimmalgi, and Laurent Schmalen, *Fellow, IEEE*

Abstract—As the demand for higher data throughput in coherent optical communication systems increases, we need to find ways to increase capacity in existing and future optical communication links. To address the demand for higher spectral efficiencies, we apply end-to-end optimization for joint geometric and probabilistic constellation shaping in the presence of Wiener phase noise and carrier phase estimation. Our approach follows state-of-the-art bitwise auto-encoders, which require a differentiable implementation of all operations between transmitter and receiver, including the DSP algorithms. In this work, we show how to modify the ubiquitous blind phase search (BPS) algorithm, a popular carrier phase estimation algorithm, to make it differentiable and include it in the end-to-end constellation shaping. By leveraging joint geometric and probabilistic constellation shaping, we are able to obtain a robust and pilot-free modulation scheme improving the performance of 64-ary communication systems by at least 0.1 bit/symbol compared to square QAM constellations with neural demappers and by 0.05 bit/symbol compared to previously presented approaches applying only geometric constellation shaping.

Index Terms—Constellation shaping, end-to-end learning, optical fiber communication, phase noise

I. INTRODUCTION

IN recent years, more and more innovations, e.g., internet of things, 6G, and video streaming, continue to increase the demand for higher data rates. Optical fiber communication systems, in particular, have to bear the bulk of the traffic to interconnect geographical regions to provide connectivity to said innovations. Therefore, to keep up with the growing demand, increasing the network capacity is necessary. Ideally, the data rate of the physical layer employed in existing communication links shall be increased.

One way to increase the data rate is to increase the spectral efficiency of optical communication systems by applying constellation shaping. Both probabilistic and geometric shaping achieve a shaping gain over classical square quadrature amplitude modulation (QAM) constellations. Probabilistic shaping changes the probability of occurrence of constellation symbols arranged in the classical square QAM, while geometric shaping changes the position of constellation points. For the

additive white Gaussian noise (AWGN) channel, there exist good solutions for probabilistic [3], [4] and geometric [5]–[7] constellation shaping. For channels with memory or nonlinearities, closed-form analytical solutions for neither probabilistic nor geometric shaping are currently available and we need to resort to numerical optimization techniques. Especially for optical fiber communication systems, a popular approach to optimize constellation shaping is to apply machine learning (ML) [8]–[11]. An end-to-end (E2E) approach allows us to optimize constellations that maximize the achievable information rate (AIR) [12] in ubiquitous bit-interleaved coded modulation (BICM)-like systems with bit-metric decoders (BMDs). Classical approaches for constellation shaping usually lack efficient ways to optimize the bit labeling for such BICM system. In more recent works [13]–[15], joint geometric and probabilistic shaping optimized with the help of machine learning and in particular neural networks (NNs) was successfully investigated.

Another way to increase the effective net data rate is to reduce the number of transmitted pilot symbols. A portion of pilots are inserted to aid carrier phase estimation (CPE) [16]; the application of blind CPE reduces the need for pilot symbols. Reducing or fully eliminating the need for pilot symbols for CPE will be a benefit for future systems. A popular and widely used algorithm for blind CPE in optical communication systems is the blind phase search (BPS) algorithm. Its popularity stems from the fact that it can be implemented in a parallel and pipelined fashion [17], enabling CPE for high symbol rates. This is a big advantage over decision-directed CPE algorithms with feedback, which are popular in wireless communication systems because of their lower computational complexity. Due to the required feedback connection, a feed-forward, pipelined implementation is not straightforward.

Combining constellation shaping with auto-encoders and blind CPE is challenging, since the E2E optimization of constellation shaping with auto-encoders requires channel models and digital signal processing (DSP) algorithms that are differentiable. Differentiability is required to enable gradient descent based optimization using the back-propagation algorithm. Other approaches try to avoid the requirement of differentiable channels by introducing surrogate channel models [18] or by using techniques like reinforcement learning [19], genetic algorithms [20] or cubature Kalman filters [21]. These approaches come with a cost, as complexity of the training increases significantly and the rate of convergence is lower.

Parts of this work have been presented at the Optical Fiber Communications Conference (OFC) 2022 in paper [1] and at the European Conference on Optical Communication (ECOC) 2022 in paper [2].

A. Rode (rode@kit.edu), B. Geiger, S. Chimmalgi, and L. Schmalen are with the Communications Engineering Lab (CEL) at Karlsruhe Institute of Technology (KIT).

Direct E2E optimization through back-propagation and gradient descent is the most robust approach, since the effects of all elements between the sender and receiver, including DSP algorithms, are included during the training through the loss function. In order to have an E2E differentiable auto-encoder channel, all operations between encoder and decoder NNs have to be differentiable w.r.t the trainable parameters in the NNs. The regular BPS algorithm includes a non-differentiable argmin operation and hence cannot be used directly with gradient descent. Thus constellation shaping with surrogate channels or without requiring a differentiable channel have been studied for this particular use-case [18], [21], [22]. We however want to leverage the robustness and convergence speed of gradient descent and hence in [1] proposed a differentiable BPS algorithm. In this work, we propose to use the differentiable BPS algorithm from [1], to optimize constellations for the Wiener phase noise channel in an E2E manner. Besides geometric constellation shaping, we also investigate the potential of joint probabilistic and geometric constellation shaping.

The remainder of this work is organized as follows. In Sec. II, we explain how the non-differentiable argmin operation can be replaced with the differentiable *softmin with temperature* operation to implement the differentiable BPS algorithm. We introduce the system model for geometric constellation shaping (GCS) in Sec. III. The extension of the GCS system model to joint geometric and probabilistic constellation shaping (GeoPCS) is highlighted in Sec. IV. In Sec. V, we discuss the chosen simulation parameters for the E2E optimization and discuss results we obtain for constellation shaping with the differentiable BPS for the Wiener phase noise channel. We compare and discuss the results of our optimizations of GCS and GeoPCS in Sec. VI. In particular, we show a novel approach to jointly optimize geometric and probabilistic shaping for the Wiener phase noise channel including the BPS algorithm. We highlight differences to previous approaches optimizing the geometric shaping. We conclude the paper in Sec. VII.

II. DIFFERENTIABLE BLIND PHASE SEARCH

In a first step, we summarize the BPS algorithm as it was presented in [17] and introduce our modification to make the BPS differentiable in a second step. The BPS algorithm is described in Alg. 1 for a set of $M = 2^m$ constellation symbols $\mathcal{M} := \{c_1, \dots, c_M\}, c_i \in \mathbb{C}$ and a sequence of complex received symbols $\mathbf{z} = (z_1, \dots, z_k, \dots)$. The BPS is parametrized by the length of the averaging window $2N + 1$ and the number of test phases L , which define the granularity with which the BPS estimates and corrects phase errors.

The first step in the BPS algorithm when estimating the phase error $\hat{\phi}_k$ of received symbol z_k consists of finding the minimum distance $d_{k,\ell}$ between all known transmit constellation symbols c_i and the symbol z_k rotated by all test phases φ_ℓ with $\ell \in \{1, \dots, L\}$. In an equivalent, practical, implementation, the rotated constellations can be saved and the minimum distance for all test phases can be computed in parallel. After computing the distances $d_{k,\ell}$, an averaging step

Algorithm 1 Differentiable and Regular BPS

Constellation $\mathcal{M}, |\mathcal{M}| = 2^m$
 Received symbol $z_k \in \mathbb{C}$
 Test phases $\varphi \leftarrow (0, \frac{1}{L}2\pi, \dots, \frac{L-1}{L}2\pi)^\top$
for $\ell \leftarrow 1, \dots, L$ **do**
 $d_{k,\ell} \leftarrow \min_{c \in \mathcal{M}} |c - z_k \exp(-j\varphi_\ell)|^2$
end for
 $D_{k,\ell} \leftarrow \sum_{\tilde{k}=k-N}^{k+N} d_{\tilde{k},\ell} \quad \forall \ell \in \{1, \dots, L\}$
if differentiable **then**
 $\hat{\phi}_k \leftarrow \varphi^\top \text{softmin}_t(D_k)$
else
 $\hat{\ell}_k \leftarrow \arg \min_{\ell=1, \dots, L} D_{k,\ell}$
 $\hat{\phi}_k \leftarrow \varphi_{\hat{\ell}_k}$
end if
 Phase unwrapping $\tilde{\phi} \leftarrow \text{unwrap}(\hat{\phi})$
 $\hat{x}_k \leftarrow z_k \exp(-j\tilde{\phi}_k)$

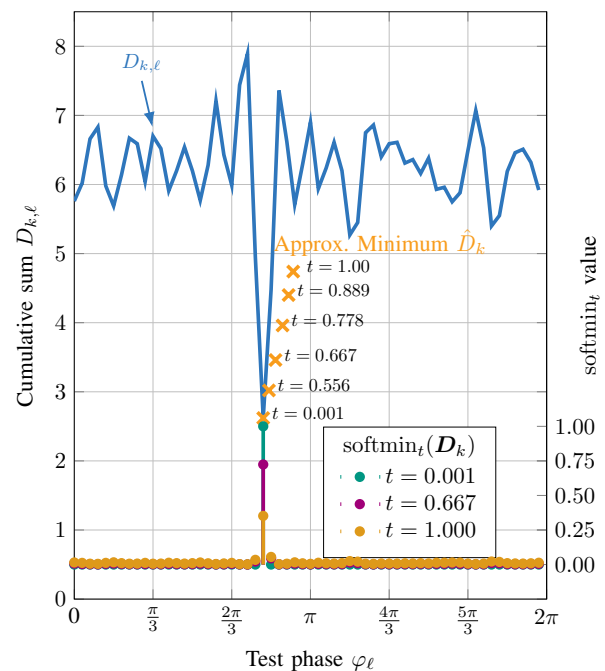


Fig. 1. $\text{softmin}_t(D_k)$ applied on the cumulative sum $D_{k,\ell}$ to approximate arg min. The approximated minimum values \hat{D}_k marked in the plot by “x” are given by $\hat{D}_k = D_k^\top \text{softmin}_t(D_k)$.

is performed by calculating the cumulative sum $D_{k,\ell}$ per test phase φ_ℓ using N previous and N following distances. This step increases the robustness against AWGN while decreasing the resilience to higher laser linewidths. The cumulative sum introduces fringe effects at the start and the end of a BPS block, which we will need to remove during optimization to resemble the behavior of a system operating in a steady state. After computing the cumulative sum $D_{k,\ell}$ the following operations between differentiable BPS and regular BPS differ. With regular BPS, the test phase index $\hat{\ell}_k$ that minimizes the cumulative sum $D_{k,\ell}$ is found using the arg min operation and the corresponding test phase $\hat{\phi}_k = \varphi_{\hat{\ell}_k}$ is used as a carrier

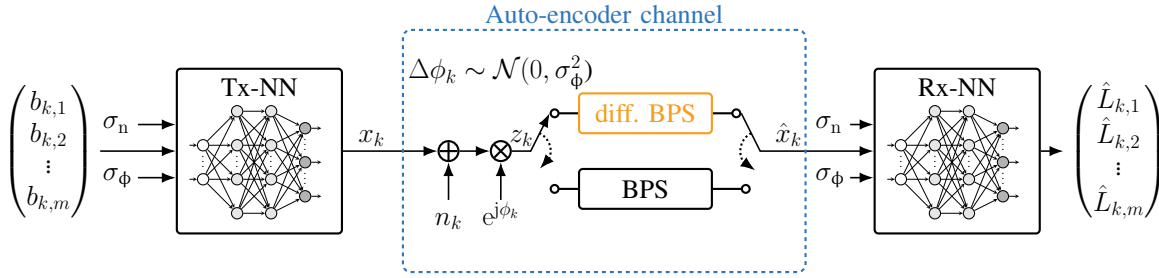


Fig. 2. Bitwise auto-encoder system model for GS with differentiable BPS

phase estimate. We apply phase unwrapping to $\hat{\phi}$ and obtain $\hat{\phi}$, which we use to correct the received symbols z . In the differentiable BPS, we replace the $\arg\min$ operation with a differentiable approximation softmin_t , which we call *softmin with temperature*. In the following section, we introduce softmin_t in more detail. To obtain a phase estimate $\hat{\phi}_k$, the dot product between the result of softmin_t and the vector of test phases φ has to be computed. The degree to which softmin_t approximates the $\arg\min$ operation is controlled with the temperature parameter t .

A. Softmin with Temperature

We derive the softmin_t from the softmin operation. For a vector $\mathbf{x} = (x_1, \dots, x_n)^T$, the i -th element of $\text{softmin}_t(\mathbf{x})$ is given by

$$\text{softmin}(x_i) := (\text{softmin}(\mathbf{x}))_i = \frac{\exp(-x_i)}{\sum_{j=1}^n \exp(-x_j)}. \quad (1)$$

To improve the approximation of the true minimum, we scale the input vector to the softmin operation with $1/t$:

$$\text{softmin}_t(x_i) := (\text{softmin}_t(\mathbf{x}))_i = \frac{\exp(-\frac{x_i}{t})}{\sum_{j=1}^n \exp(-\frac{x_j}{t})}. \quad (2)$$

Now, for $0 < t < 1$, the minimum and maximum are further separated in value than for the unscaled input. Therefore the following softmin operation returns a higher weight for the minimum value. In Fig. 1, we show the change in the approximation of $\arg\min$ with softmin_t for a varying parameter t . Decreasing t improves the approximation: For small values of t , most elements of the softmin_t output vector are zero; only a single non-zero element at a single test phase which leads to the smallest distance in the cumulative sum vector \mathbf{D}_k persists. This corresponds closely to the $\arg\min$ operation. The approximated minimum value $\hat{D}_k = \mathbf{D}_k^T \text{softmin}_t(\mathbf{D}_k)$ is plotted in Fig. 1 only for illustrative purposes, as we are only interested in the phase estimate $\hat{\phi}_k$ which is given by $\hat{\phi}_k = \varphi^T \text{softmin}_t(\mathbf{D}_k)$. For $t = 1$, softmin_t is equivalent to the softmin operation and we can observe a significant offset from the true $\arg\min$. In [23], the authors show how to smoothly approximate $\arg\max$ with a similar approach applied to the softmax function. In our use case, the softmin operation does not approximate the $\arg\min$ sufficiently and consequently the returned phase estimate has a significant offset from the true phase unless the extra parameter t is used.

III. GEOMETRIC CONSTELLATION SHAPING SYSTEM MODEL

Auto-encoders have been successfully used for unsupervised learning of efficient latent representations in the wider machine learning community and are quite naturally applicable to the design of communication systems. Namely, in an auto-encoder, information is first processed by an encoder NN to get a representation in a latent space, which usually has a smaller dimension than the input. A decoder NN is then used to reconstruct the original information from the information in the latent space. No labeling in the latent space is required, instead, representations of the information can be extracted from the latent space. We can cast this auto-encoder concept to a communication problem: the encoder is a mapper which maps bit vectors to symbols on the complex plane and the decoder is a demapper trying to recover the embedded information from complex, noisy received symbols. In order to learn representations that are efficient in the presence of channel impairments, the symbols—generated by the encoder NN (denoted in the following by “Tx-NN”)—are impaired by a communication channel, e.g., by AWGN, dispersion, or multi-path propagation. These impairments are partially removed with classical DSP algorithms and the recovered complex symbols are processed by a decoder NN (denoted in the following by “Rx-NN”) to return a log likelihood ratio (LLR) vector. By optimizing encoder and decoder NNs to minimize the binary cross entropy (BCE) loss [12], an efficient mapping can be found by back-propagation and gradient descent with, e.g., the Adam algorithm [24]. The learned encoder and decoder do not necessarily have to be used in a practical implementation to reap the benefits of this optimization. The encoder can be replaced by a lookup table and the bitwise decision regions can be approximated by simpler structure than a deep neural network. In [25], we present a way to learn a decoder with comparable complexity to a Gaussian demapper.

Our proposed system model in Fig. 2 is built according to a bitwise auto-encoder. At time step k , a bit vector $\mathbf{b}_k = (b_{1,k}, \dots, b_{m,k})^T$ is first mapped to a corresponding one-hot vector $\hat{\mathbf{b}}_k$ of length 2^m that is only non-zero at the index that corresponds to the binary-to-integer conversion of the bit vector. The Tx-NN is used to generate a constellation \mathcal{M} of size $M = 2^m$. In the case of a non-parametrizable mapper, the mapper neural network (Tx-NN) reduces to a real-valued weight matrix \mathbf{W}_m of size 2×2^m and we can generate a

vector of constellation symbols $\mathbf{c} = (c_1, \dots, c_M)^\top \in \mathbb{C}^M$ containing all M modulation symbols of the constellation $\mathcal{M} := \{c_1, \dots, c_M\}$ with

$$\mathbf{c}^\top = (1 \quad \mathbf{j}) \underbrace{\begin{pmatrix} W_{1,1} & \dots & W_{1,M} \\ W_{2,1} & \dots & W_{2,M} \end{pmatrix}}_{=\mathbf{W}_m}. \quad (3)$$

The dot product between a one-hot vector $\tilde{\mathbf{b}}_k$ and \mathbf{c} is then used to select one constellation symbol x_k for transmission:

$$x_k = \tilde{\mathbf{b}}_k^\top \mathbf{c}. \quad (4)$$

AWGN and Wiener phase noise are applied to x_k to simulate a communication channel which is dominated by Gaussian noise and impairments from a non-zero laser linewidth. The standard deviations $\sigma_n = \sqrt{N_0} = \sqrt{\frac{E_s}{\text{SNR}}}$ and $\sigma_\phi = \sqrt{2\pi \frac{\Delta f}{R}}$ are selected so that a channel with SNR = $\frac{E_s}{N_0}$ and linewidth Δf at symbol rate R is simulated. The transmit symbols x_k are affected by complex AWGN $n_k \sim \mathcal{CN}(0, \sigma_n^2)$ and Wiener phase noise $\phi_k = \phi_{k-1} + \Delta\phi_k$ with $\Delta\phi_k \sim \mathcal{N}(0, \sigma_\phi^2)$. The impaired symbols z_k are then further sent through a CPE algorithm to recover and correct the carrier phase. In our system model in Fig. 2, we use either our differentiable BPS or the regular BPS. After CPE, we obtain the phase compensated symbols \hat{x}_k . The demapper neural network (Rx-NN) takes the complex symbols \hat{x}_k and performs demapping to obtain m LLRs \hat{L}_k .

To evaluate the performance of a BICM system we calculate the bitwise mutual information (BMI)¹ as the metric. For a sequence of P bit vectors \mathbf{b}_k ($k \in \{1, \dots, P\}$) and their corresponding LLRs \hat{L}_k , the BMI is approximated as

$$\text{BMI} \approx \mathbb{H}(\mathcal{M}) - \frac{1}{P} \sum_{k=1}^P \sum_{i=1}^m \log_2 \left(1 + e^{(-1)^{b_{k,i}} \hat{L}_{k,i}} \right), \quad (5)$$

where $\mathbb{H}(\mathcal{M})$ is the entropy in bits of the modulation symbols, calculated from the discrete probability of occurrence $p(c_i)$ for each modulation symbol c_i :

$$\mathbb{H}(\mathcal{M}) = -\mathbb{E}[\log_2 p(c_i)]. \quad (6)$$

In case of a uniform distribution of the modulation symbols, (5) further simplifies to

$$\text{BMI} \approx m - \frac{1}{P} \sum_{k=1}^P \sum_{i=1}^m \log_2 \left(1 + e^{(-1)^{b_{k,i}} \hat{L}_{k,i}} \right). \quad (7)$$

Subsequently, instead of a custom implementation, the BCE loss function common to many machine learning frameworks can be used directly as a loss function to optimize the BMI (see [12] for details).

We initialize the Wiener phase noise process with a starting phase $\phi_{\text{start}} \sim \mathcal{U}(-\pi, \pi)$, where $\phi_k = \phi_{\text{start}} + \sum_{k'=0}^{k-1} \Delta\phi_{k'}$. The random initialization of the starting phase helps to improve the robustness of the constellation when BPS is used for CPE without prior compensation using pilots. This allows us to learn a constellation for a pilot-less system which is robust to AWGN and Wiener phase noise.

¹The BMI is often referred to as generalized mutual information (GMI) in the optical communications community. We prefer to use the term BMI due to its easier resemblance with the operational meaning.

A. Parameterizable GCS

Our proposed system model in Fig. 2 contains additional inputs σ_n and σ_ϕ at both Tx-NN and Rx-NN to allow for the optimization of GCS over a range of channel parameters. This parameterization serves multiple purposes: it allows the investigation of changes in the constellation when channel parameters are varied. Additionally, we obtain a mapper and demapper optimized for each set of channel parameters within the range of training parameters. This parameterization is roughly equivalent to training a separate set of mapper and demapper NNs for each individual set of channel parameters. Such individualized training would be significantly more computationally expensive and the bit labeling may be different for different channel parameters preventing us from analyzing the isolated impact of the variation in the parameters on the shaping. With this parametrized geometric constellation shaping (pGCS), each symbol—which corresponds to a particular bit vector—is moved only in a small region defined by the set of channel parameters. The pGCS retains its general structure and changes can be observed with respect to the change in the channel parameters.

B. Trainable Differentiable BPS

Having implemented a differentiable BPS, a natural question arises: Can the differentiable implementation replace the regular BPS not only for the training and optimization, but also for evaluation and implementation? To investigate this possibility, we additionally make the temperature t in the differentiable BPS trainable. To implement a regularization of the temperature parameter t , we choose an unbounded trainable parameter t^* , which we use to calculate $t = \sigma(t^*)$. Here, $\sigma(x) = (1 + e^{-x})^{-1}$ is the sigmoid function and limits t to the range $(0, 1)$. We then use the differentiable BPS and the optimized parameter t^* in the implementation instead of the regular BPS. Since the phase estimate of the differentiable BPS is not limited to the granularity of the test phases φ , the remaining residual phase noise may decrease.

IV. JOINT GEOMETRIC AND PROBABILISTIC CONSTELLATION SHAPING SYSTEM MODEL

To incorporate the idea of learning a joint geometric and probabilistic shaping of the constellation, we extend our system model in Fig. 2 with a distribution sampler at the transmitter for the training in Fig. 3, as proposed in [13]. This extension changes the operation of the training in terms of how bits and symbols are generated and sent through the channel. With GeoPCS, an NN (denoted by “p-NN”) first generates a vector of logits ℓ_p to represent the probability of occurrence for each constellation point in the log domain. Then, a softmax operation is applied on ℓ_p to get a vector of probabilities \mathbf{p} .

For each training batch, the symbol probabilities \mathbf{p} and the batch size S are passed to a distribution sampler. To find the quantized number of symbols according to the given distribution of symbol probabilities \mathbf{p} , we apply Algorithm 2 from [26]. We convert the symbol indices to their binary representations to get S bit vectors \mathbf{b}_k , which can be passed through the neural mapper to get S complex symbols \mathbf{x} .

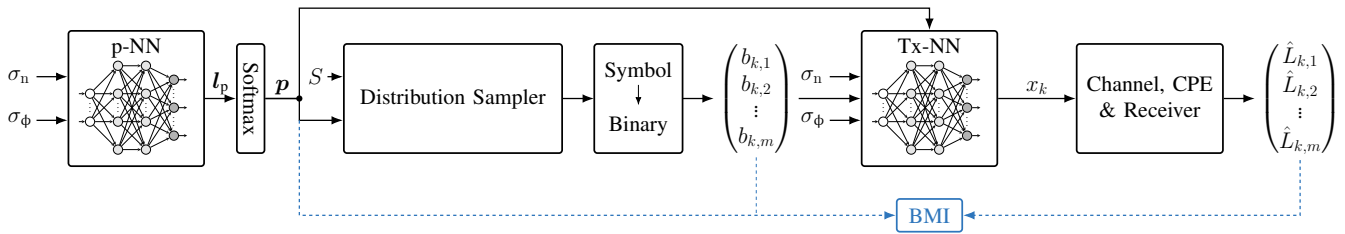


Fig. 3. Parameterizable bitwise auto-encoder system model for GeoPCS

To account for the non-uniform symbol probability, we change the energy normalization operation in the neural mapper to apply \mathbf{p} to scale the contribution of each constellation point according to the symbol probability $p(c_i)$. The normalized constellation is obtained as

$$\mathcal{M}_{\text{norm}} = \left\{ c_i \cdot \left(\sum_{j=1}^M p(c_j) |c_j|^2 \right)^{-\frac{1}{2}} : c_i \in \mathcal{M} \right\}. \quad (8)$$

The remaining operations to obtain the S LLRs vectors are the same as in the case for only geometrical shaping (GS) in Fig. 2. For optimization of this GeoPCS system model, we use a modified loss function based on the BMI in (5). This is motivated by the fact that using the BCE does not correctly incorporate the reduction of the entropy resulting from probabilistic shaping in (5). The entropy is calculated directly from \mathbf{p} and optimization is performed with automatic differentiation and gradient descent. We also extend the approach presented in [13] to use a neural network with one hidden layer with inputs for σ_ϕ and σ_n to allow for parameterization of ℓ_p . In the case parameterization is not required, the neural network for ℓ_p simplifies to a vector of trainable weights.

A. Probability Distributions with Symmetry

In order to learn a logits vector ℓ_p which can be implemented jointly with forward error correction (FEC) in a BICM system akin to the probabilistic amplitude shaping (PAS) scheme [4], [27], we extend the probabilistic shaping to be parameterizable with a symmetry parameter s . This parameter s controls the size of the output dimension of the p-NN to be 2^{m-s} . This limits the parameter s to a range of $[0, m-1]$ to have any probabilistic shaping. To obtain a probability of occurrence for each modulation symbol, the logits vector $\ell_{p,s}$ is repeated s times to form the final logits vector ℓ_p of length 2^m . This introduces an s -fold symmetry in ℓ_p and in turn this results in s bits of the bit vector forming the modulation symbols that are 0 or 1 with equal probability. This property can be inferred from the symbol-to-bit mapping and bit-to-symbol mapping, which convert the symbol index to the binary representation. This property can then be used to combine $m-s$ information bits from a distribution matcher (in the final system) with s parity check bits from an FEC encoder to select modulation symbols corresponding to the probability distribution. For regular and fully flexible probabilistic shaping, the parameter s is set to zero. To perform probabilistic constellation shaping which contains one bit with equal probability distribution the symmetry parameter has to be set to $s=1$. As highlighted in the system models in

TABLE I
SYSTEM PARAMETERS FOR TRAINING TX-NN AND RX-NN

Bits per Symbol m	6
Epochs	1000
Batches per Epoch	10, ..., 500
Batchsize per Epoch	1000, ..., 10000
Optimization Algorithm	Adam [24]
Learning rate	0.001
Temperature parameter t	1.0, ..., 0.001
Test phases ϕ	60
Cumulative sum length N	128
Symbol rate	32 GBaud

[4], [27], a fixed number of bits with an equal probability distribution can be used to assign the approximately uniform distributed parity check bits from FEC.

V. SIMULATION SETUP

We perform GCS and GeoPCS with the system model shown in Fig. 2 and Fig. 3 with the set of hyperparameters given in Table I. The trainable parameters of the NNs are initialized randomly with the Glorot initializer. Validation is performed with the regular BPS algorithm unless stated otherwise. Across epochs, the channel parameters such as SNR and linewidth are fixed. In each batch of every epoch, the input bits are randomly permuted and a new noise realization is used. While a symbol rate of 32 GBaud is quite low for modern single-carrier transmission systems, it is a realistic value for multi-carrier transmission systems [28]–[30]. Additionally, for low cost systems where coherent transmission is relevant, the symbol rate may be higher, but the resulting phase noise is still comparable to our parameterization due to low cost and high linewidth lasers [31].

We implemented our system in the PyTorch machine learning framework [32] and our source code is accessible in [33]. Selection of an appropriate range of temperature values for the softmax_t operation during the optimization is crucial for the stability of the training. Starting with a temperature $t=1$ and then decreasing it to $t=0.001$ as the training of the NNs progresses yielded stable optimization in our setup. We attribute this behavior to the fact that for low temperature values, only a few of the distance measures around the minimum contribute to the phase estimate. Therefore, most of the other test phases are scaled with a value close to zero; subsequently, also the contribution to the gradients is small. Therefore, a rather wide inclusion of values at the beginning of the training helps to boost the training, to avoid getting stuck

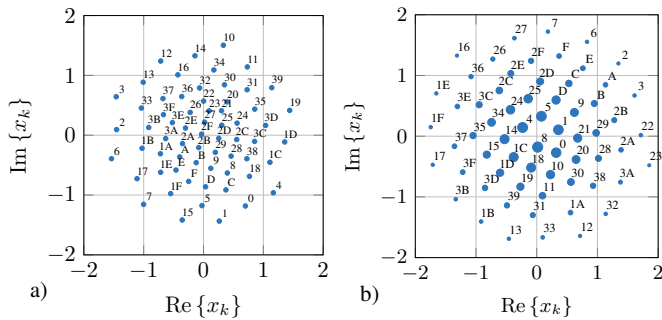


Fig. 4. Bit-labeled constellation with a) GCS and b) GeoPCS, trained using the differentiable BPS algorithm at 17 dB SNR and 100 kHz laser linewidth. The size of the constellation points in b) indicates the probability of occurrence. Bit labels are obtained by converting the bit vectors to their hexadecimal representation, e.g., $(0, 1, 1, 1, 1, 1) \equiv 1F$

in a local minimum and finally to reach a good distribution of constellation points that are optimized in later training epochs.

Training and validation for the shaped constellations has been performed with random phase initialization of the Wiener phase noise process. In the case of QAM, the Wiener phase noise has been initialized with zero and BPS only operates on the first quadrant, since otherwise the phase offset cannot be estimated correctly due to the rotational ambiguities in the QAM constellation.

VI. RESULTS

We performed optimization of the constellation shape with both GCS and the GeoPCS system models. Both approaches have been trained with a fixed channel parameterization to get a general idea about the shape and bit labeling of the learned constellation mapper and demapper. Additionally, we also trained with varying channel parameters as additional inputs to the NNs. This allows us to conduct a further investigation of the effects of CPE on learned communication systems and we obtain an estimate of the expected performance gain for optimized GCS and GeoPCS.

A. Fixed Channel Parametrization GCS and GeoPCS

In this approach, the training is performed with fixed channel parameters, which are not passed as inputs to Tx-NN and Rx-NN. The optimization of the NNs is performed at a fixed operating point, and we expect the performance on different channel parameterizations to be worse. In Fig. 4a), we show a GS constellation for an SNR = 17 dB and laser linewidth $\Delta f = 100$ kHz. A pronounced feature is the introduced asymmetry in the constellation in the lower left corner. Multiple constellation points are placed further separated from the other constellation points than their counterparts in other corners. This positioning increases the robustness and supports the operation of the BPS algorithm by resolving possible ambiguities of constellation points. This asymmetry comes at the expense of more closely packed constellation points at the center of the constellation and thus offers lower robustness to AWGN compared to a constellation optimized for AWGN only. In Fig. 4b), we show a GeoPCS constellation with the

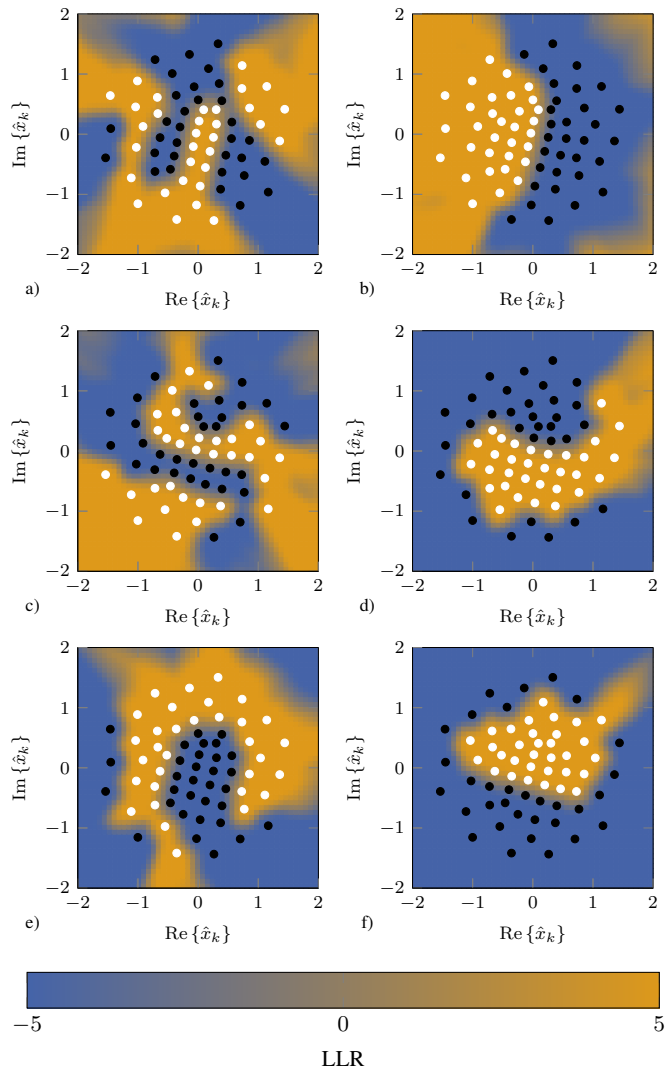


Fig. 5. Neural demapper decision regions for each of the bits for constellation in Fig. 4a). Sub-figure a) corresponds to the least-significant bit and sub-figure f) to the most significant bit of the binary labels.

same fixed channel parameters as the GCS constellation. For the human eye, the induced asymmetry is hard to spot, but the constellation contains less rotational symmetry compared to a square QAM.

To investigate the performance of our E2E system optimized with GCS, we take a look at the learned neural demapper and the decision regions for individual bits as depicted in Fig. 5. For better visualization, the demapper decision regions are shown for LLRs between $[-5, 5]$.

With the chosen neural demapper architecture, our system is able to accurately partition the constellation into two sets of the same size for each bit. The decision regions shown in Fig. 5 are an extension of the 1D LLR plots shown in [34] to the full 2D region. A 1D graph would not be enough to characterize the decision regions of the neural demapper, since $L_{k,i}$ depends on both $\text{Re}\{\hat{x}_k\}$ and $\text{Im}\{\hat{x}_k\}$.

Finally, we compare the attainable performance of the different approaches in Fig. 6. For the systems trained on

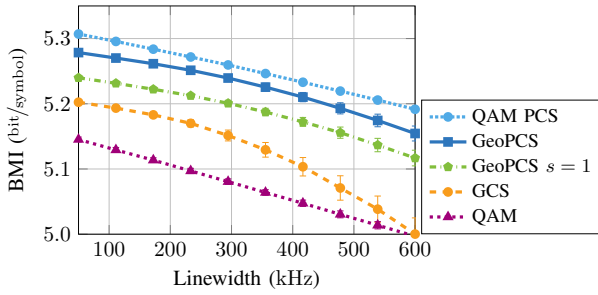


Fig. 6. Performance comparison between constellations optimized at SNR = 17 dB with GCS and GeoPCS at an SNR = 17 dB and a laser linewidth $\Delta f = 100$ kHz.

a fixed set of channel parameters, the GeoPCS system has the best performance, outperforming our previously presented work on GCS [1] by 0.05 bit/symbol at the channel parameters both systems were trained on and by even more in the region of higher laser linewidths. The GeoPCS constellation with introduced symmetry also outperforms the GCS constellation. All shaped constellations outperform a square QAM constellation paired with a neural demapper, which was trained on the same channel parameters as the shaped constellation systems. For a fair comparison, the validation of the square QAM was not initialized with a random start phase, since QAM is rotationally symmetric and the BPS algorithm wouldn't be able to disambiguate carrier phases if the phase difference equals a multiple of $\pi/2$; this means that phase slips (sometimes referred to as cycle slips) do not play a role in our evaluation. To investigate the performance of probabilistically shaped QAM, we perform an E2E optimization to learn the parameter λ of the Maxwell-Boltzmann distribution for the fixed channel parameters. We apply a similar approach as for GeoPCS, but instead of the probability of occurrence for each symbol, the NN returns the parameter λ , which is then used with the Maxwell-Boltzmann distribution and known constellation symbols to obtain the probability distribution. In Fig. 6, we report the results of this validation as “QAM PCS”. For the fixed channel parameterization we see a significant gain compared to square QAM and a similar performance as our proposed schemes. This difference in performance is most likely caused by Gray labeling in the square QAM and “QAM PCS” which is not achieved by the trained geometric constellations, but rather an almost Gray labeling is learned.

B. GCS with Varying Channel Parameterization

To optimize GCS for a range of varying channel parameterizations, we train the pGCS system on channel parameterizations drawn from $\mathcal{U}(\sigma_{n,\min}, \sigma_{n,\max})$ and $\mathcal{U}(\sigma_{\phi,\min}, \sigma_{\phi,\max})$. The parameters are also provided to the NNs, such that the system is optimized for each set of provided channel parameters and learns how to change the constellation to obtain a better constellation and demapper for the given condition. In Fig. 7, we show the obtained transmit constellation for training the system on a channel with SNR between 14 dB and 24 dB and the laser linewidth Δf between 50 kHz and 600 kHz. The constellation is depicted for a fixed SNR = 18 dB

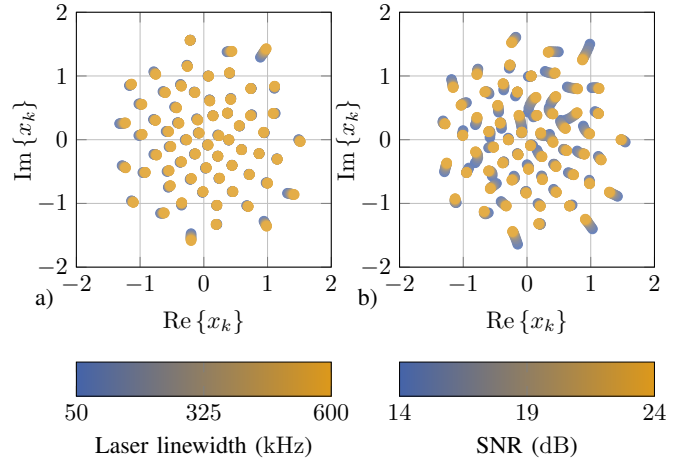


Fig. 7. pGCS constellation for $M = 64$ at a) fixed SNR = 18 dB and varying laser linewidth, and b) fixed laser linewidth $\Delta f = 100$ kHz and varying SNR

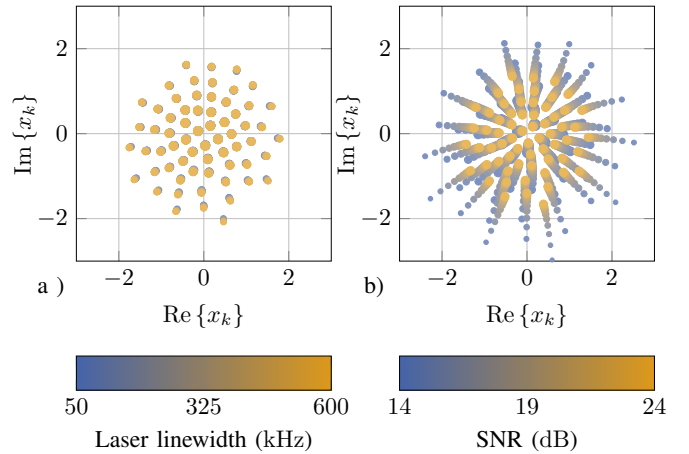


Fig. 8. pGeoPCS constellation for $M = 64$ at a) fixed SNR = 18 dB and varying laser linewidth, and b) fixed laser linewidth $\Delta f = 100$ kHz and varying SNR

and varying laser linewidth in a color change in Fig. 7a). For a high laser linewidth, the constellation exhibits a few points at the outside of the constellation at the top right and bottom right corners, which move outwards compared to the constellation optimized for a lower laser linewidth. The change in the transmit constellation diagram is small and in stark contrast to the change in the transmit constellation diagram for varying SNR depicted in Fig. 7b). In this diagram, the laser linewidth $\Delta f = 100$ kHz is fixed and the color indicates the SNR. For low SNR, the same constellation points as for varying laser linewidth are moved outwards compared to the transmit constellation for high SNR. Additionally, for low SNR, more points at the center of the constellation are moved closer together to allow for a separation of other constellation points at the outer edge of the constellation. This “sacrifices” the information contained in at least one bit carried by the constellation and results in a more Gaussian-like shaping of the complete transmit constellation diagram.

The results of training the parameterized GeoPCS

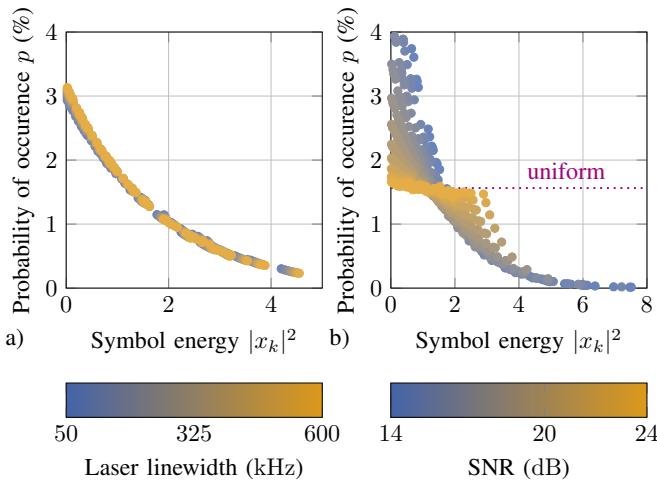


Fig. 9. Learned PCS for a) varying laser linewidth at SNR = 18 dB, and b) varying SNR at laser linewidth $\Delta f = 100$ kHz

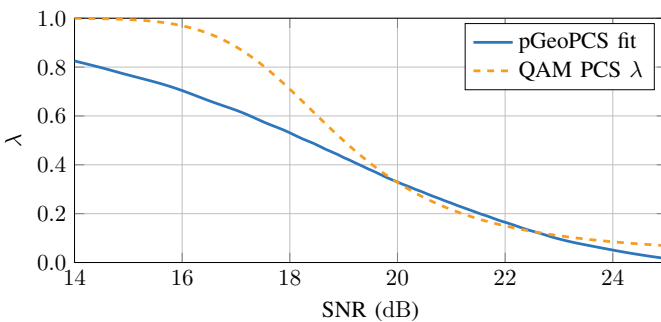


Fig. 10. Numerical fit of learned PCS to parameter λ of the Maxwell-Boltzmann distribution over SNR

(pGeoPCS) are depicted in Fig. 8 and Fig. 9. The transmit constellation is shown in Fig. 8a) for varying laser linewidth and shown in Fig. 8b) for varying SNR. The change in the probability of occurrence for the transmit symbols over the symbol energy is shown in Fig. 9a) for varying laser linewidth and for varying SNR in Fig. 9b). Contrary to the pGCS system, a varying laser linewidth does not change any features of the constellation, neither the geometric constellation shape in Fig. 8a) nor the probabilistic constellation shape in Fig. 9a). This result is somewhat surprising, as this means that the learned transmit constellation is only changed for varying SNR and is independent of the laser linewidth. For a change in SNR, a behaviour similar to that of the AWGN channel is observed [13]. The probabilistic shaping approaches the shape of the well-known Maxwell-Boltzmann distribution.

In Fig. 10, we depict the value λ for the normalized Maxwell-Boltzmann probability mass function

$$P_{MB}(c_i) = \frac{e^{-\lambda|c_i|^2}}{\sum_{v=1}^M e^{-\lambda|c_v|^2}}, \quad (9)$$

which has the closest numerical fit to the learned symbol distribution. The fit to the Maxwell-Boltzmann distribution is very accurate with a maximum observed Kullback-Leibler divergence of 0.0002. In addition, we also display the learned

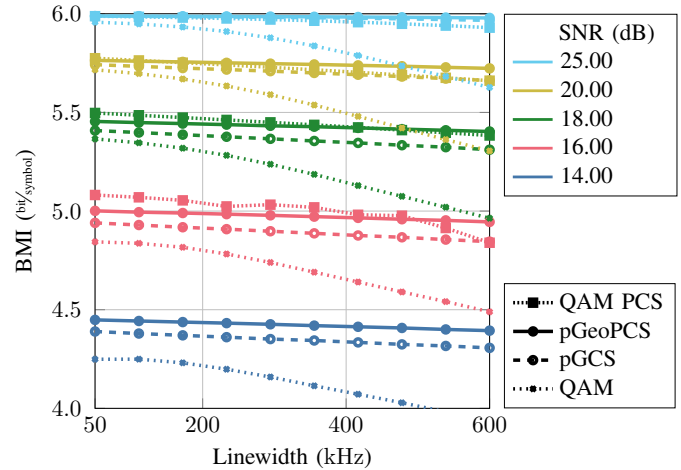


Fig. 11. Performance comparison between constellations optimized with GCS and GeoPCS with varying channel parameters.

parameter λ of the parameterized “QAM PCS” constellation. In our training, we limited the λ to the range of $[0, 1]$ by applying a sigmoid function at the output of the neural network. In the resulting plot, a similar shape of the learned λ parameter over a changing SNR can be observed. The λ parameter learned by the probabilistically shaped QAM differs from the fitted λ parameter from the pGeoPCS constellation especially for low SNR. For a higher λ , a stronger probabilistic shaping is performed. This difference may indicate a reason for the performance difference seen in Fig. 11.

In Fig. 11, we compare the performance of our different approaches to constellation shaping with varying channel parameters. The performance is shown in BMI over a varying laser linewidth. Colors indicate SNR and different lines and markers indicate the system. The NNs of all depicted systems used σ_n and σ_ϕ as input to optimize the mapper and demapper for the channel parameters used for validation. In this comparison, applying GeoPCS provides a consistent performance gain over GCS. But it is also clearly visible that for higher SNR, this shaping gain disappears. In comparison, the results of “QAM PCS”, which is a parameterized version of the probabilistically shaped square QAM (using a Maxwell-Boltzmann distribution), show stability issues for low SNR. For the lowest value of 14 dB, the results are not shown since the performance was impacted by many cycle slips. In the higher SNR regime, a similar picture as to the fixed training point can be seen: for low laser linewidths, “QAM PCS” outperforms the pGeoPCS scheme, where for higher laser linewidths pGeoPCS is more robust and shows a better performance. We conclude from these results that the probabilistically shaped square QAM has advantages in the low laser linewidth regime if the shaping parameter λ is chosen correctly. For high laser linewidths, our proposed schemes employing geometric shaping show better robustness. We have to emphasize that the square QAM constellations have a better starting condition, their phase noise realizations are initialized with zero phase offset, therefore mandating the use of regular pilot symbols to retain this phase reference.

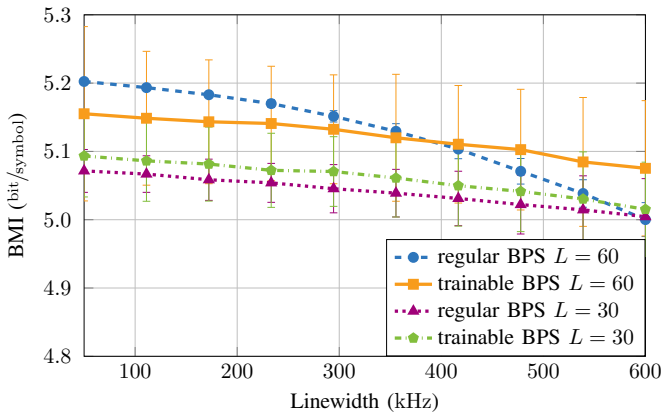


Fig. 12. Performance comparison between trainable differentiable BPS and regular BPS trained on a channel with SNR = 17 dB and laser linewidth $\Delta f = 100$ kHz.

C. Trainable Differentiable BPS

For the trainable BPS, our assumption is that a lower number of test phases could be used for the final system since the differentiable BPS can interpolate between the test phases. Therefore, we train and evaluate GCS systems with and without a trainable differentiable BPS for a number of test phases $L = 30$ and $L = 60$. All systems are trained on a fixed SNR = 17 dB and a fixed laser linewidth $\Delta f = 100$ kHz. The performance comparison between the system with trainable differentiable BPS and the system with regular BPS is depicted in Fig. 12. For $L = 60$, the regular BPS has a consistently better performance compared to the trainable BPS. We attribute this to the phase discontinuity at 0 and 2π . If the actual phase error ϕ_k is close to the discontinuity, we obtain similar values for the cumulative sum $D_{k,\ell}$ for the test phases $\varphi_1 = 0$ and $\varphi_L = \frac{L-1}{L}2\pi$ in the BPS algorithm. In that case, the softmin_t will return high values for both these test phases even when the temperature t is very low. Subsequent multiplication and summation with the test phase vector φ returns a phase error estimate $\hat{\phi}_k$ around π which is significantly different from the actual phase error leading to incorrect phase correction and consequently large errors in the LLRs. This issue does not occur with an argmin , since the issue of phase error estimates of consecutive symbols oscillating between 0 and 2π is corrected by the phase unwrapping step after obtaining the phase error estimates.

The incorrect phase corrections in some cases leads to higher performance variation of the trainable BPS between runs, which is indicated by the error bars. For $L = 30$ the trainable BPS has a slightly better performance. In this case, the ability to perform interpolation with the softmin_t operation between test phases might reduce residual phase noise compared to the regular BPS, but high performance variations between runs are still observed. We therefore suggest further investigation into increasing the stability and robustness of the differentiable BPS to avoid problems at phase errors close to the discontinuity.

TABLE II
PARAMETERS FOR THE DP SSFM

Transmitter SNR	24 dB
RRC rolloff	0.1
RRC filter length	1000
RRC oversampling	2 samples/symbol
Simulation oversampling	16 samples/symbol
Baud rate	32 GBaud
Optical carrier wavelength	1550 nm
Span length	80 km
Number of spans	4
SSFM steps per span	80
Attenuation α_{dB}	0.2 dB km ⁻¹
Dispersion D	17 ps nm ⁻¹ km ⁻¹
Kerr parameter γ	1.2 W ⁻¹ km ⁻¹
EDFA noise figure	5 dB

D. Validation with Dual Polarization Split-Step Fourier Method Simulation

In order to verify our results reported in Fig. 11 on a channel impaired by chromatic dispersion, we additionally perform validation with an implementation of a dual polarization split-step Fourier method (DP-SSFM). This method simulates the effects described by the coupled non-linear Schrödinger equation. We do not include other channel variations, like polarization mode dispersion (PMD) or random polarization rotation. This would entail investigating joint optimization of constellation shaping and blind equalization methods [35]. Similarly to a geometrically shaped constellation with prominent asymmetric features, a constellation which is shaped geometrically to assist a blind equalizer may exhibit other asymmetric features. The study of interaction between the optimization of geometric constellation shaping and equalization is an interesting topic, but is outside the scope of this work. Applying the DP-SSFM introduces some inter-symbol interference and the results of this validation may give an indication if increased tolerance for phase noise come at a cost of higher sensitivity to inter-symbol interference.

First, we apply a root-raised-cosine (RRC) pulse shaping and upsampling to the simulation rate, to the transmit symbols x_k . Then, AWGN and Wiener phase noise are added to the transmit samples to simulate transmitter impairments. The same Wiener phase noise realization is used on both polarizations. The impaired transmit samples are then sent through the DP-SSFM and a different Wiener phase noise is applied on the receiver side. Afterwards we downsample the received signal and apply digital chromatic dispersion compensation. The matched filter follows the dispersion compensation and we apply the BPS algorithm separately for both polarizations. In this validation we use the regular BPS. In the case of square QAM, we apply a genie-aided phase-slip compensation on every symbol to remove any phase ambiguity by a multiple of $\pi/2$. We use the trained neural receivers to obtain LLRs and compute the performance in terms of BMI.

We selected system parameters reported in Table II. To find the best launch power we performed a sweep across the launch power for each constellation shaping method. For this validation, we use the models trained in a parameterizable fashion for the channel with AWGN, Wiener phase noise

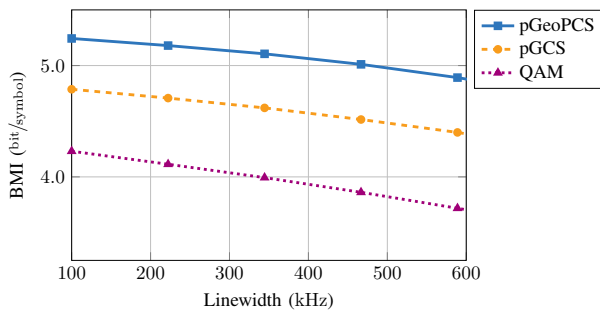


Fig. 13. Performance comparison between square QAM, pGCS, and pGeoPCS constellations on a DP-SSFM channel. The shown laser linewidth is the sum of the laser linewidth at the transmitter and receiver, corresponding to $\sigma_{\phi, \text{tot}}$

and the BPS algorithm as presented in Sec. IV and Sec. III, and use them in the validation without any re-training. One challenge is correctly selecting the input parameters σ_{ϕ} and σ_n for the NNs. Since we apply i.i.d. Wiener phase noise at the transmitter and receiver, we assume the resulting phase noise coming from the Wiener phase processes can be described with $\sigma_{\phi, \text{tot}} = \sqrt{2}\sigma_{\phi}$ and we use $\sigma_{\phi, \text{tot}}$ as input parameter for the NNs. For AWGN, we have multiple sources to account for, first we apply AWGN at the transmitter, but also the ASE noise contribution of the EDFAs is modeled as AWGN. Additionally, the noise contribution of the non-linearity and interaction between non-linearity and chromatic dispersion is not accounted for in the parameterization of the NNs. Therefore we perform a sweep of the σ_n parameter, which is only used as input to the NNs, between 17 dB and 24 dB to find the best operation point for the parameterizable mappers and demappers.

In Fig. 13, we show the performance of the different constellations in terms of their BMI over a varying laser linewidth. The launch power has been optimized between 0 dBm and 3 dBm in steps of 1 dBm and for all constellations, a launch power of 2 dBm has shown the best performance. As for the σ_n input parameter for the NNs, the best performance for all constellations has been observed for a parameter corresponding to an SNR ≈ 18 dB. Validation with the DP-SSFM model confirms our observations made on the linear channel model in Fig. 6 and in Fig. 11. The parametrizable GeoPCS constellation outperforms both pGCS and QAM constellations. We computed the E2E SNR in our simulations to be between 19 dB and 20 dB. Comparing the results in Fig. 13 with results obtained on the linear channel, we observe a significant impact by additional non-linear impairments and residual inter-symbol interference.

VII. CONCLUSION

We presented and analyzed bitwise auto-encoders for joint optimization of geometric and probabilistic shaping for Wiener phase noise channels with carrier phase estimation. We used a differentiable carrier phase estimation (CPE) algorithm which allows for efficient end-to-end optimization using gradient-descent. With the proposed system, optimization of GCS leads to a constellation robust to phase noise and supports the

operation of the CPE. A further extension to joint geometric and probabilistic constellation shaping provides additional shaping gain compared to GCS. In this work, we showed that joint geometric and probabilistic shaping can be applied to communication systems impaired by AWGN and Wiener phase noise with BPS as the carrier phase estimation algorithm. With a parameterizable mapper, demapper and probabilistic shaper, we are able to outperform square QAM by more than 0.1 bit/symbol for low SNR values. Comparison with gradient-free methods and a study of their computational complexity and convergence speed is planned future work.

ACKNOWLEDGMENT

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101001899). The authors acknowledge support by the state of Baden-Württemberg through bwHPC.

REFERENCES

- [1] A. Rode, B. Geiger, and L. Schmalen, "Geometric constellation shaping for phase-noise channels using a differentiable blind phase search," in *Optical Fiber Communications Conference (OFC)*, Mar. 2022, p. Th2A.32.
- [2] A. Rode and L. Schmalen, "Optimization of geometric constellation shaping for Wiener phase noise channels with varying channel parameters," in *European Conference on Optical Communication (ECOC)*, Sep. 2022, p. Mo3D.5.
- [3] F. Kschischang and S. Pasupathy, "Optimal nonuniform signaling for Gaussian channels," *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 913–929, May 1993.
- [4] G. Böcherer, P. Schulte, and F. Steiner, "Probabilistic shaping and forward error correction for fiber-optic communication systems," *Journal of Lightwave Technology*, vol. 37, no. 2, pp. 230–244, Jan. 2019.
- [5] G. Forney and L.-F. Wei, "Multidimensional constellations. I. Introduction, figures of merit, and generalized cross constellations," *IEEE Journal on Selected Areas in Communications*, vol. 7, no. 6, pp. 877–892, Aug. 1989.
- [6] G. Forney, "Multidimensional constellations. II. Voronoi constellations," *IEEE Journal on Selected Areas in Communications*, vol. 7, no. 6, pp. 941–958, Aug. 1989.
- [7] P. Larsson, "Golden angle modulation," *IEEE Wireless Communications Letters*, vol. 7, no. 1, pp. 98–101, Feb. 2018.
- [8] R. T. Jones, T. A. Eriksson, M. P. Yankov, and D. Zibar, "Deep learning of geometric constellation shaping including fiber nonlinearities," in *European Conference on Optical Communication (ECOC)*, Sep. 2018.
- [9] B. Karanov, M. Chagnon, F. Thouin, T. A. Eriksson, H. Bülow, D. Lavery, P. Bayvel, and L. Schmalen, "End-to-end deep learning of optical fiber communications," *Journal of Lightwave Technology*, vol. 36, no. 20, pp. 4843–4855, Oct. 2018.
- [10] K. Gümüş, A. Alvarado, B. Chen, C. Häger, and E. Agrell, "End-to-end learning of geometrical shaping maximizing generalized mutual information," in *Optical Fiber Communication Conference (OFC)*, Mar. 2020, p. W3D.4.
- [11] H. Dzieciol, G. Liga, E. Sillekens, P. Bayvel, and D. Lavery, "Geometric shaping of 2-D constellations in the presence of laser phase noise," *Journal of Lightwave Technology*, vol. 39, no. 2, pp. 481–490, Jan. 2021.
- [12] S. Cammerer, F. Ait Aoudia, S. Dörner, M. Stark, J. Hoydis, and S. ten Brink, "Trainable communication systems: Concepts and prototype," *IEEE Transactions on Communications*, vol. 68, no. 9, pp. 5489–5503, Sep. 2020.
- [13] V. Aref and M. Chagnon, "End-to-end learning of joint geometric and probabilistic constellation shaping," in *Optical Fiber Communications Conference (OFC)*, Mar. 2022, p. W4I.3.
- [14] M. Stark, F. Ait Aoudia, and J. Hoydis, "Joint learning of geometric and probabilistic constellation shaping," in *IEEE Global Communications Conference (GLOBECOM)*, Dec. 2019.

- [15] V. Oliari, B. Karanov, S. Goossens, G. Liga, O. Vassilieva, I. Kim, P. Palacharla, C. Okonkwo, and A. Alvarado, "Hybrid geometric and probabilistic shaping: Is it really necessary?" in *Advanced Photonics Congress*, Jul. 2022, p. SpTu1J.4.
- [16] M. Mazur, J. Schröder, A. Lorences-Riesgo, T. Yoshida, M. Karlsson, and P. A. Andrekson, "Overhead-optimization of pilot-based digital signal processing for flexible high spectral efficiency transmission," *Optics Express*, vol. 27, no. 17, pp. 24 654–24 669, Aug. 2019.
- [17] T. Pfau, S. Hoffmann, and R. Noe, "Hardware-efficient coherent digital receiver concept with feedforward carrier recovery for M -QAM constellations," *Journal of Lightwave Technology*, vol. 27, no. 8, pp. 989–999, Apr. 2009.
- [18] O. Jovanovic, M. P. Yankov, F. Da Ros, and D. Zibar, "End-to-end learning of a constellation shape robust to variations in SNR and laser linewidth," in *European Conference on Optical Communication (ECOC)*, Sep. 2021.
- [19] F. Ait Aoudia and J. Hoydis, "End-to-end learning of communications systems without a channel model," in *Asilomar Conference on Signals, Systems, and Computers*, Oct. 2018, pp. 298–303.
- [20] K. Yan, F. Yang, C. Pan, J. Song, F. Ren, and J. Li, "Genetic algorithm aided gray-APSK constellation optimization," in *International Wireless Communications and Mobile Computing Conference (IWCMC)*, Jul. 2013, pp. 1705–1709.
- [21] O. Jovanovic, M. P. Yankov, F. Da Ros, and D. Zibar, "Gradient-free training of autoencoders for non-differentiable communication channels," *Journal of Lightwave Technology*, vol. 39, no. 20, pp. 6381–6391, Oct. 2021.
- [22] O. Jovanovic, F. Da Ros, D. Zibar, and M. P. Yankov, "Geometric constellation shaping for fiber-optic channels via end-to-end learning," *arXiv*, Nov. 2022, arXiv:2211.04311.
- [23] C. J. Maddison, A. Mnih, and Y. W. Teh, "The concrete distribution: A continuous relaxation of discrete random variables," in *International Conference on Learning Representations (ICLR)*, Mar. 2017.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [25] A. Rode and L. Schmalen, "Geometric constellation shaping with low-complexity demappers for Wiener phase-noise channels," in *Optical Fiber Communications Conference (OFC)*, Mar. 2023, p. Th3E.3.
- [26] G. Böcherer and B. C. Geiger, "Optimal quantization for distribution synthesis," *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 6162–6172, Nov. 2016.
- [27] L. Schmalen, "Probabilistic constellation shaping: Challenges and opportunities for forward error correction," in *Optical Fiber Communications Conference (OFC)*, Mar. 2018.
- [28] H. Sun *et al.*, "800G DSP ASIC design using probabilistic shaping and digital sub-carrier multiplexing," *Journal of Lightwave Technology*, vol. 38, no. 17, pp. 4744–4756, Sep. 2020.
- [29] B. Geiger, E. Sillekens, F. Ferreira, R. Killely, L. Galdino, and P. Bayvel, "Record 2.29 Tb/s GS-256QAM transmission using a single receiver," in *European Conference on Optical Communication (ECOC)*, 2022, p. We3D.3.
- [30] D. Welch *et al.*, "Digital subcarrier multiplexing: Enabling software-configurable optical networks," *Journal of Lightwave Technology (early access)*, 2022.
- [31] A. Shahpari, R. M. Ferreira, R. S. Luis, Z. Vujicic, F. P. Guiomar, J. D. Reis, and A. L. Teixeira, "Coherent access: a review," *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 1050–1058, Feb. 2017.
- [32] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [33] A. Rode, S. Chimmalgi, and B. Geiger, 2022, https://github.com/kit-cel/e2e_opt_diff_bps, source code to be published upon acceptance of this paper.
- [34] O. Shental and J. Hoydis, "'Machine LLRning': Learning to softly demodulate," in *IEEE Global Communications Conference (GLOBECOM)*, Dec. 2019.
- [35] V. Lauinger, F. Buchali, and L. Schmalen, "Blind equalization and channel estimation in coherent optical communications using variational autoencoders," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 9, pp. 2529–2539, Sep. 2022.