

Deep FAVIB: Deep Learning-Based Forward-Aware Quantization via Information Bottleneck Method

Matthias Hummert , Shayan Hassanpour , Dirk Wübben , and Armin Dekorsy 

Department of Communications Engineering,

University of Bremen,

28359 Bremen, Germany

Email: {hummert, hassanpour, wuebben, dekorsy}@ant.uni-bremen.de

Abstract—We focus on a (generic) *joint source-channel coding* problem, appearing in a broad variety of real-world application. Explicitly, a *noisy* observation from a user/source signal should be compressed, ahead of getting forwarded over an *error-prone* and *rate-limited* channel to a remote processing unit. The design problem shall be formulated in a fashion that the impacts of the forward link are taken into account. Aligned with the *Information Bottleneck (IB)* method, we consider the *Mutual Information (MI)* as the fidelity criterion, and work out a data-driven approach to tackle the underlying design problem based upon a finite sample set. For that, we derive a tractable *variational lower-bound* of the objective functional, and present a general *learning architecture* which can be used to optimize the given lower-bound by standard training of the encoder and decoder *Deep Neural Networks*. This approach that is, principally, based upon the (generative) latent variable models, extends the concepts of *Variational AutoEncoder (VAE)* and *Deep Variational Information Bottleneck (Deep VIB)* for (remote) source coding to the context of joint source-channel coding. We validate the effectiveness of our approach by several numerical simulations over typical transmission scenarios.

Index Terms—Information bottleneck, variational autoencoder, deep learning, quantization, joint source-channel coding, 6G

I. INTRODUCTION

The *Information Bottleneck (IB)* method for data compression, first introduced in [1], borrows ideas from the seminal work of Shannon [2] on lossy source coding. Motivated by the fact that, in many real-world clustering applications, it is much easier to determine a *relevant/target* variable (whose information shall be retained) than figuring out the suitable distortion measure, the IB's design formulation applies an intuitive twist on the *single-letter* characterization of the *Rate-Distortion* function. Explicitly, instead of upper-bounding the average distortion, it lower-bounds the (relevant) information between the target and the latent variable. Interested readers are referred to [3] for an overview of conceptual ideas, linked to the applicational scope of the IB method.

Aside from theoretical studies, the practical usage of the IB method (in modern transmission systems) recently has gained an increasing momentum. Those applications include (but are not limited to the) Analog-to-Digital (A/D) converters at receiver front-ends [4], discrete channel decoding schemes [5]–[9], semantic/task-oriented communications [10]–[13], and efficient construction of Polar Codes [14], [15].

In recent years, by exploiting the capabilities of Deep Learning and Neural Networks (NNs), novel data-driven approaches have been devised to efficiently optimize the IB trade-off (see, e.g., [16]–[18]). These methods that work based upon a finite sample set drastically expand the scope of applications of the IB principle as, unlike the conventional algorithms, they do not require the full prior statistical knowledge of input variables. Moreover, they can efficiently handle high dimensional (and even continuous) data.

As the main contribution, we present with *Deep Forward-Aware Vector Information Bottleneck (Deep FAVIB)* the data-driven, sample-based counterpart of the *Forward-Aware Vector Information Bottleneck (FAVIB)* algorithm, first presented in [19]. With this, we directly extend the concept of Deep VIB [16] for (noisy) source coding to the context of joint source-channel coding by taking into account an *error-prone* forward channel (FC). Furthermore, we perform an in-depth analysis of Deep FAVIB and as the main result, we show performance on par with FAVIB. This, indeed, is quite remarkable, as we use solely a finite sample set, without knowing the joint statistics of the input signals. FAVIB, on the other hand, requires the full statistical knowledge. This indicates that Deep FAVIB can be utilized in a much wider scope of scenarios, where the full statistical knowledge of the input signals is not available.

The rest of the paper is organized as follows: In Section II, we introduce our system model and the pertinent design problem. This will be followed by presenting our deep variational approach in Section III. Thereupon, in Section IV, we provide several simulation results to corroborate the effectiveness of our approach. Finally, in Section V, we conclude this paper by a short wrap-up, containing the salient points.

II. SYSTEM OVERVIEW

Notation: The random variable, \mathbf{b} , with the probability mass function $p(\mathbf{b})$ accepts certain realizations, b , from its domain, \mathcal{B} , and with boldface, the random vector, \mathbf{b} , is given. $D_{\text{KL}}(\cdot||\cdot)$, $H(\cdot)$, and $I(\cdot;\cdot)$ are the Kullback-Leibler Divergence (KLD), Shannon's entropy and Mutual Information (MI) [20].

A. System Model and Problem Formulation

In Fig. 1 the general setup is depicted. A discrete-valued data source $x \in \mathcal{X}$ generates modulated symbols, mapped




Fig. 1. System model: Symbols x are transmitted over an access channel, quantized, forwarded via an imperfect channel, and reconstructed by decoder \hat{x} .

by any suitable modulation scheme. Afterwards, an access channel $p(y|x)$ will introduce signal distortion, leading to the received signal $y \in \mathcal{Y}$. A quantizer $p(z|y)$ will then compress/quantize the signal y to another signal $z \in \mathcal{Z}$ with a fixed cardinality $|\mathcal{Z}| = N$, yielding a clustered y . After quantization, a rate-limited Forward Channel (FC) $p(t|z)$ with capacity R will introduce further distortions, yielding $t \in \mathcal{T}$. Finally, a decoder $p(\hat{x}|t)$ will reconstruct the transmitted symbol \hat{x} . This forms a joint source-channel coding scheme as, through the compression, the quantizer also takes into account the FC impacts.

B. Conventional IB Method

To start the discussion we begin by designing the quantizer $p(z|y)$ with the conventional IB method. To that end, we need to know the joint statistics of the input signals. The quantizer is designed such that, through the compression, the information about the source signal x is mostly preserved.

1) *Perfect Forwarding*: Presuming a *perfect/error-free* FC, and by utilizing the IB method, we maximize the MI between x and z , i.e., $I(x; z)$ under the constraint that the MI between y and z , i.e., $I(y; z)$ is limited to the capacity R of the FC¹. Mathematically, we write

$$p^*(z|y) = \operatorname{argmax}_{p(z|y): I(y; z) \leq R} I(x; z), \quad (1)$$

with $0 \leq R \leq \log_2 |\mathcal{Z}|$ providing an upper-bound to $I(y; z)$. This constrained optimization problem can be reformulated with the *Lagrange Method of Multipliers* [22] leading to

$$p^*(z|y) = \operatorname{argmax}_{p(z|y)} I(x; z) - \lambda I(y; z), \quad (2)$$

where $\lambda \geq 0$ is directly related to the capacity R of the FC. The objective function in (2) sets a basic trade-off between reconstruction $I(x; z)$ and compression $I(y; z)$. Maximizing the MI between x and z will ensure maximal reconstruction capabilities for the decoder, while minimizing the MI between y and z will ensure maximal compression w.r.t. y . In other words, we want to compress the observed signal y to a variable z such that the relevant information about the source x is mostly preserved. The stationary solution of (2) has been derived for each $(y, z) \in \mathcal{Y} \times \mathcal{Z}$ in [1] as

$$p^*(z|y) = \frac{p(z)}{\omega(y, \lambda)} \exp\left(-\lambda^{-1} D_{\text{KL}}(p(x|y) || p(x|z))\right), \quad (3)$$

where $\omega(y, \lambda)$ is a normalization function, ensuring a correct conditional distribution. An iterative algorithm has also been presented in [1] which, in principle, performs the *Fixed-Point Iterations* [23] on the implicit solution (3). For non-zero λ values, usually, a *soft/stochastic* quantizer is achieved.

¹Refer to [21] for the (asymptotic) remote source coding formulation with the *Logarithmic Loss* distortion.

2) *Imperfect Forwarding*: For an *Imperfect/error-prone* FC, we extend the design problem to take its effects into account

$$p^*(z|y) = \operatorname{argmax}_{p(z|y): I(y; z) \leq R} I(x; t). \quad (4)$$

Applying the *Lagrange Method of Multipliers* [22], we get

$$p^*(z|y) = \operatorname{argmax}_{p(z|y)} \underbrace{I(x; t) - \lambda I(y; z)}_{\mathcal{L}_{\text{FAVIB}}}. \quad (5)$$

Now, the quantizer is designed such that the information about the source signal x after the FC is mostly preserved, with the same constraint on compression rate $I(y; z)$ w.r.t. the FC capacity, R . The stationary solution of (5) has been derived for each $(y, z) \in \mathcal{Y} \times \mathcal{Z}$ in [19] as

$$p^*(z|y) = \frac{p(z)}{\omega(y, \lambda)} \exp\left(-\lambda^{-1} \sum_{t \in \mathcal{T}} p(t|z) D_{\text{KL}}(p(x|y) || p(x|t))\right), \quad (6)$$

where $\omega(y, \lambda)$ is again a normalization function. Similar to the original IB scenario, an iterative algorithm, the *Forward-Aware (Vector) Information Bottleneck (FAVIB)*, has been presented in [19] that performs the *Fixed-Point Iterations* [23] on the derived implicit solution (6), together with its proof of convergence to a stationary point of the objective functional $\mathcal{L}_{\text{FAVIB}}$. Like before, for non-zero λ values, usually, a *soft/stochastic* quantizer is achieved.

III. DEEP FAVIB

Here, we develop our new deep learning approach, named Deep FAVIB, to approximately tackle the design problem (5), when instead of the joint statistics $p(x, y)$, solely a (finite) sample set $\{x_m, y_m\}_{m=1}^M$ (where M is the total number of samples) is available. This approach which generalizes the *Deep Variational Information Bottleneck (Deep VIB)* [16], is based upon the (generative) latent variable models, specifically, the well-known concept of *Variational Auto-Encoder (VAE)* [24], [25].

A. Variational Lower-Bound

The starting point to develop a data-driven design approach is to introduce a tractable *Variational Lower-Bound (VLB)* on the objective functional $\mathcal{L}_{\text{FAVIB}}$ in (5). We start by introducing variables A and B for reconstruction and compression. More specifically, let A be a lower-bound on $I(x; t)$, i.e., $I(x; t) \geq A$, and B an upper-bound on $I(y; z)$, i.e., $I(y; z) \leq B$. Then, it applies

$$\mathcal{L}_{\text{FAVIB}} = I(x; t) - \lambda I(y; z) \geq A - \lambda B = \mathcal{L}^{\text{VLB}}. \quad (7)$$




Fig. 2. Learning architecture for Deep FAVIB, including two DNNs for encoder and decoder, a Gumbel sampler, and a softmax/argmax unit

For relevant information $I(x; t)$, the following holds true

$$I(x; t) = \underbrace{H(x)}_{\geq 0} - H(x|t) \quad (8a)$$

$$\geq \underbrace{\sum_{t \in \mathcal{T}} p(t) D_{\text{KL}}(p(x|t) || q(x|t))}_{\geq 0} + \sum_{x \in \mathcal{X}, t \in \mathcal{T}} p(x, t) \log q(x|t) \quad (8b)$$

$$\geq E_{x,t} \{\log q(x|t)\} = A, \quad (8c)$$

where we introduce the proxy posterior $q(x|t)$, replacing the perfect decoder $p(x|t)$. From (8a) to (8b), the non-negativity of entropy $H(x)$, and, from (8b) to (8c), the non-negativity of KLD (a.k.a. information inequality) has been applied [20]. Further, for the compression rate $I(y; z)$, we can follow

$$I(y; z) = \sum_{y \in \mathcal{Y}, z \in \mathcal{Z}} p(y, z) \log \frac{p(z|y)}{r(z)} - \underbrace{D_{\text{KL}}(p(z) || r(z))}_{\geq 0} \quad (9a)$$

$$\leq \mathbb{E}_{y,z} \left\{ \log \frac{p(z|y)}{r(z)} \right\} = B, \quad (9b)$$

in which we introduce $r(z)$ as an arbitrary prior for the latent variable z . From (9a) to (9b), the non-negativity of KLD has been applied. Finally, for the VLB we get with (7)

$$\mathcal{L}^{\text{VLB}} = E_{x,t \sim p(x,t)} \{\log q(x|t)\} - \lambda E_{y,z \sim p(y,z)} \left\{ \log \frac{p(z|y)}{r(z)} \right\}. \quad (10)$$

To design the quantizer/encoder $p(z|y)$ and the approximate decoder $q(x|t)$, next, we introduce parameterized distributions, which are realized via Deep Neural Networks (DNNs). It holds

$$\begin{aligned} \mathcal{L}^{\text{DNN}} &= E_{x,t \sim p(x,t)} \{\log q_\phi(x|t)\} - \lambda E_{y,z \sim p(y,z)} \left\{ \log \frac{p_\theta(z|y)}{r_\psi(z)} \right\} \\ &= \underbrace{E_{t \sim p(t)} \{ E_{x \sim p(x|t)} \{\log q_\phi(x|t)\} \}}_{\text{reconstruction}} \\ &\quad - \lambda \underbrace{E_{y \sim p(y)} \{ D_{\text{KL}}(p_\theta(z|y) || r_\psi(z)) \}}_{\text{regularization}}, \end{aligned} \quad (11)$$

with weights θ , ψ , and ϕ . Note that we want to maximize (11). We observe that maximizing the relevant information corresponds to minimizing the *cross-entropy* loss (which is the standard reconstruction loss for classification, following the *Maximum-Likelihood* learning rule [26]), averaged over t . The counterpart term for compression rate, on the other hand, acts as a regularization since the quantizer $p_\theta(z|y)$ should match the prior $r_\psi(z)$ via KLD, averaged over y .

B. Learning Architecture & Implementation Details

In general, we need to design a *stochastic* encoder $p_\theta(z|y)$ and decoder $q_\phi(x|t)$ via DNNs. To estimate the gradients of \mathcal{L}^{DNN} , the conventional approach of utilizing the reparameterization trick [24] to enable the Monte-Carlo sampling and subsequently, replacing the expectation terms with their (empirical) estimates can be exploited here as well. Since our main focus will be on discrete latent spaces, the *Gumbel-Softmax/Concrete Distribution* [27], [28] can be employed to do the trick, i.e., reparameterizing the underlying categorical distribution. The overall learning architecture of our proposed data-driven approach, i.e., the Deep FAVIB, has been depicted in Fig. 2.

The received signal $y \in \mathcal{Y}$ is generally complex-valued. To serve as an input of the Neural Network (NN)-Encoder (with weights θ) the received signal $\mathbf{y}_{\text{real}} \in \mathbb{R}^2$ is stacked into a 2D vector containing real and imaginary parts, separately. This is necessary as NNs cannot handle complex numbers straightforwardly². \mathbf{y}_{real} serves as input for the encoder NN, whose output $\pi \in (0, 1)^N$ forms the categorical distribution of the discrete latent z . To sample from the corresponding concrete variable, N i.i.d. samples from $\text{Gumbel}(0, 1)$ distribution are generated and stacked into the vector $\mathbf{g} \in \mathbb{R}^N$. The sum signal $\log(\pi) + \mathbf{g}$ is then fed into a softmax/argmax unit. If we use argmax, one of the N entries is set to 1, and the remaining $N-1$ entries become 0. This is known as *one-hot encoding*. For *training*, argmax is not a suitable choice as gradients cannot be calculated through this function. To circumvent this issue, the softmax is applied. The i -th entry of the sample vector \mathbf{z}_{samp} , is calculated as ($i = 1$ to N)

$$z_{\text{samp},i} = \frac{\exp\left(\left(\log(\pi_i) + g_i\right)/\tau\right)}{\sum_{j=1}^N \exp\left(\left(\log(\pi_j) + g_j\right)/\tau\right)} \in [0, 1], \quad (12)$$

where $\tau > 0$ is a hyper-parameter, named temperature. The smaller the τ , the closer behaves softmax like argmax. τ is an important tradeoff factor. The smaller τ , the more steep the softmax becomes and gradients change rapidly, and hence a poor local optima may be found. If τ is too large, the softmax will differ too much from argmax. During *deployment*, only the argmax is used, yielding one latent realization z that is then transmitted over the FC and the decoder NN $q_\phi(x|t)$ with parameters ϕ is used to recover \hat{x} . The decoder consists of a standard Feed-Forward NN. This overall chain extends the

²The problem arises when updating the gradients, as complex derivatives are not always straightforward to calculate.

VAE architecture. Specifically, in a traditional VAE we go from x to z and back to x , but in our chain, we have a noisy version y of x , as the input signal to the encoder. Additionally, we have an *error-prone* FC, disturbing the latent variable z , and consequently, the input signal to the decoder is also a noisy version t of z .

C. Neural Networks & Supervised Learning

NNs are nonlinear functions with trainable parameters, here θ and ϕ . These parameters are adapted w.r.t. a loss function, here $-\mathcal{L}^{\text{DNN}}$ (11). For a given data set $\{x_m, y_m\}_{m=1}^M$, we can update the parameters by some form of *Stochastic Gradient Descent* with back-propagation. To train the decoder, samples t_m are needed. These t_m are generated through the data set indirectly. They are formed via a markov chain from y_m to t_m , where the quantizer influences how z_m and subsequently t_m is formed. The goal is to adapt the weights for both encoder and decoder NNs together to minimize the loss function. The prior $r_\psi(z)$ although having its own trainable parameters (probabilities of different clusters), does not depend on a NN.

IV. NUMERICAL RESULTS

We compare our proposed scheme's performance, with FAVIB, as, ideally, Deep FAVIB should perform close to FAVIB, by maximizing a lower-bound. FAVIB is guaranteed to converge to a stationary solution [19]. To avoid poor local optima, FAVIB is run 100 times and the best result is shown.

We use a rather simple scenario to make comparisons with FAVIB. Explicitly we use Quadrature Phase Shift Keying (QPSK) modulation, an Additive White Gaussian Noise (AWGN) access channel with noise variance σ_n^2 , and regarding FC, we consider an N -ary symmetric model with error probability e . Remember that z is a clustered version of y , containing N clusters. Specifically, with probability $1 - e$, no cluster flip/change occurs and with probability e , one of the $N - 1$ remaining clusters is chosen. i.e. if $e = 0$, $t = z$ is true. For training, we use $M = 1\text{e}6$ samples, a batch size of 10000 and a maximum of 50000 epochs. We apply *Early Stopping*, to track the lowest overall loss, and save the best weights. The learning rate is set to 10^{-5} and Adam [29] is used for optimization. The model parameters are λ , σ_n^2 , τ and e . The NN configurations of the encoder and decoder can be seen in Table I. Both NNs follow a Feed-forward structure with multiple hidden layers with *Rectified Linear Unit (ReLU)* activation functions. The output activation function of the encoder is linear to form the log probabilities $\log(\pi)$ and the output activation function of the decoder is a softmax to distinguish between the symbols in the modulation alphabet \mathcal{X} . Training is needed only once for each specific setup. Afterwards, we simply deploy the trained NNs. Nevertheless, for each parameter sweep in the upcoming plots, a new training is conducted.

A. Temperature τ and Quantization Borders

To investigate the effects of temperature τ on the performance of Deep FAVIB, we vary τ over a certain range and calculate the obtained relevant information $I(x;t)$ for

TABLE I
CONFIGURATION FOR ENCODER NN, DECODER NN, AND PRIOR

Name	# of Hidden Layers	Width of Layers	# of Weights
$p_\theta(z y)$	3	300, 200, 100	82816
$q_\phi(x t)$	3	300, 200, 100	85602
$r_\psi(z)$	0	0	N




Fig. 3. Relevant information $I(x;t)$ versus temperature τ for different error probabilities e . $N = 16$, $\lambda = 0.01$, and $\sigma_n^2 = 0.4$.

different error probabilities e . We set $N = 16$, $\lambda = 0.01$, and $\sigma_n^2 = 0.4$. The pertinent results have been depicted in Fig. 3. In general, we observe that lower temperature values τ lead to better performance for high error probabilities e , and vice versa. As a consequence, the temperature value τ has to be chosen carefully depending on the scenario, as it can have a direct influence on the quality of outcome. For upcoming simulations, we fix $\tau = 0.1$ if $e > 0$ as this yields the best result. If $e = 0$, we fix $\tau = 1$, as high values of τ lead to better performance in this case.

To understand the dynamics of compression based on the quality of forward channel, in Figs. 4 and 5, we present the obtained quantization regions of unreliable ($e = 0.25$) as well as fully reliable forwarding ($e = 0$), respectively. The horizontal and vertical axes show the real and imaginary parts of the received signal y . In color, with arbitrary numbering, the quantization regions are shown. We also provide the individual contribution $i(z)$ of each cluster z (a.k.a. partial MI) to the overall relevant information $I(x;t)$ in Table II. Note that $I(x;t) = \sum_z i(z)$ holds.

We set $N = 16$, $\sigma_n^2 = 0.4$, and $\lambda = 0$ for both figures (i.e., during training, only reconstruction loss is considered). The difference comes in the error probabilities and, consequently, the chosen temperature values. In Fig. 4, $e = 0$ (fully reliable forwarding), and accordingly, we choose $\tau = 1$. In Fig. 5, we set $e = 0.25$ (highly unreliable forwarding), and based on our previous investigations, we choose $\tau = 0.10$. During our simulations we observed that low temperature values, usually, lead to usage of less clusters and vice versa. In Fig. 4, we observe that all 16 clusters are used. Explicitly, 4 large corner clusters are formed, as well as 10 "in-between" clusters and 2 clusters around the origin. Referring to Table II, it is observed that the 4 corner clusters $\{9, 10, 11, 12\}$ carry the most

TABLE II
PARTIAL MI $i(z)$ OF EACH CLUSTER FOR FIGS. 4 AND 5

Fig.	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	\sum
4	0.002	0.047	0.035	0.046	0.048	0.017	0.001	0.050	0.020	0.304	0.298	0.270	0.255	0.039	0.033	0.050	1.518
5	≈ 0	0.221	0.218	≈ 0	≈ 0	≈ 0	0.010	0.009	0.215	≈ 0	0.221	0.013	≈ 0	≈ 0	0.009	≈ 0	0.917

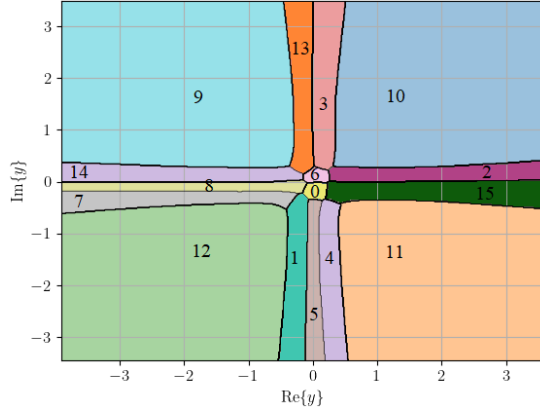


Fig. 4. Obtained quantization regions, equiprobable QPSK signaling over an AWGN access channel with $\sigma_n^2 = 0.4$, $N = 16$ clusters, temperature $\tau = 1$ and $e = 0$ (fully reliable forwarding).




Fig. 5. Obtained quantization regions, equiprobable QPSK signaling over an AWGN access channel with $\sigma_n^2 = 0.4$, $N = 16$ clusters, temperature $\tau = 0.1$, and $e = 0.25$ (unreliable forwarding).

information (roughly, 74% of the overall MI). The clusters in the center $\{0,6\}$, carry almost no information ($i(\cdot) \approx 0$) and act as an erasure.

In Fig. 5, we observe that not all $N = 16$ clusters are utilized. Explicitly, 4 corner cluster, 4 "in-between" clusters and one cluster around the origin are formed. The remaining 7 clusters are not used. Referring to Table II, we see that the 4 corner cluster $\{1, 2, 8, 10\}$ carry 95% of the MI. The center cluster $\{4\}$ carries no information ($i(4) \approx 0$) and acts as an erasure.

In comparison, in case of small e (Fig. 4), the forward channel capacity is large enough to allow the quantizer to actively utilize (almost) all clusters to generate a finer decomposition at each quadrant. On the contrary, in case of large e (Fig. 5), the quantizer cannot afford the luxury of such fine decompositions. The information is best conveyed with a lower number of active clusters. Overall, this behaviour can be interpreted as a form of channel coding, as the FC has direct impact on the engendered quantization regions. This highlights

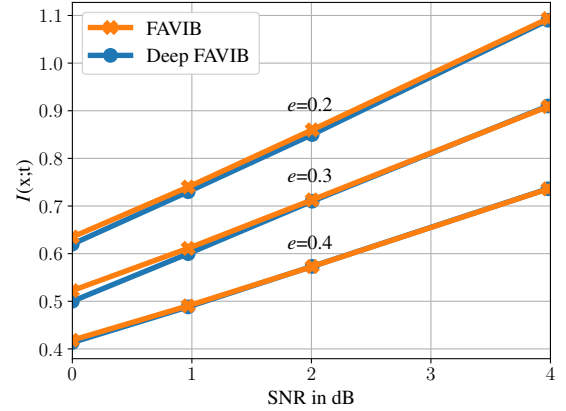


Fig. 6. Relevant information $I(x;t)$ versus SNR of AWGN access channel, $N = 32$ clusters, $\lambda = 0.01$, $\tau = 0.1$, and different error probabilities e .

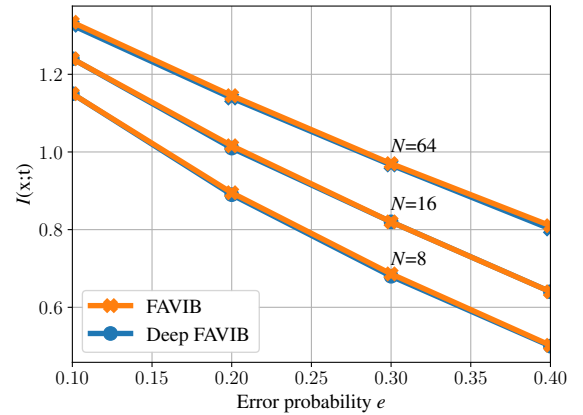


Fig. 7. Relevant information $I(x;t)$ versus error probability e of FC, $\sigma_n^2 = 0.4$, $\lambda = 0.01$, $\tau = 0.1$, and different number of clusters N .

the connection to joint source-channel coding.

B. End-to-End Rate for Different SNR and Error Probabilities

We evaluate the overall MI $I(x;t)$ against various noise levels σ_n^2 and error probabilities e for $\lambda = 0.01$, $N = 32$ and $\tau = 0.1$. The obtained results are shown in Fig. 6.

As expected, the relevant information $I(x;t)$ decreases, the lower the SNR becomes, as less information can be passed through the system overall. The trends of Deep FAVIB and FAVIB are nearly linear but flatten out for low SNR. This is consistent across all error probabilities e shown and is strictly related to the capacity of AWGN channel and its relation to the noise variance σ_n^2 and the SNR [19]. Furthermore, both algorithms yield lower relevant information $I(x;t)$, the higher e becomes, as naturally, the less reliable the forward channel is, the less information can be carried to the decoder. As a major result, we observe that the performance of Deep FAVIB is quite close to FAVIB. This shows that a single training

round for Deep FAVIB is sufficient to deliver quite comparable results with the best out of 100 reruns of FAVIB.

In Fig. 7, we show the relevant information $I(x; t)$ against the error probability e of the FC for a fixed noise variance of $\sigma_n^2 = 0.4$. We set $\lambda = 0.01$ and vary N . As expected, the higher the error probability e becomes, the less MI is available as the scheme becomes less reliable overall. Furthermore, naturally, as the quantizer output cardinality N grows, higher relevant information values can be achieved as more available clusters lead to an easier forwarding of information. By comparing the performances of both algorithms, it is again observed that Deep FAVIB, with a single training round, comes on par with the best result out of 100 reruns of FAVIB for all evaluated number of clusters N . This clearly substantiates the promising performance of our novel Deep FAVIB approach.

V. SUMMARY

In this work, we presented the *Deep FAVIB*, a data-driven, deep variational approach to address the IB-based design problem for a generic joint source-channel coding setup that appears in real-world applications. The main advantage of Deep FAVIB is the fact that, it is trained by a (finite) sample set, thereby obviating the need for prior statistical knowledge of the input signals. Other State-of-the-Art (SoTA) algorithms like FAVIB require the prior knowledge of the joint statistics of input signals. Deep FAVIB that is based upon (generative) latent variable models directly generalizes the well-known concepts of *Variational Auto-Encoder (VAE)* [24] and *Deep Variational Information Bottleneck (Deep VIB)* [16] to the context of joint source-channel coding by incorporating an *error-prone* forward channel into the design problem. We further substantiated the performance of Deep FAVIB by an in-depth analysis and comparison with FAVIB. As the main result, we showed performance on par with FAVIB, which is a remarkable result. It was shown that, a single training round for Deep FAVIB is sufficient to come on par with the best result out of 100 reruns of FAVIB. This directly motivates the usage of Deep FAVIB in real-world scenarios where the joint statistics of input signals is not available.

ACKNOWLEDGEMENT

This work was partly funded by the German ministry of education and research (BMBF) under grants 16KISK016 (Open6GHub) and 16KISK068 (6G-TakeOff).

REFERENCES

- [1] N. Tishby, F. C. Pereira, and W. Bialek, "The Information Bottleneck Method," in *37th Annual Allerton Conf. on Comm., Control, and Computing, Monticello, IL, USA*, September 1999.
- [2] C. E. Shannon, "Coding Theorems for a Discrete Source with a Fidelity Criterion," *IRE Int. Convention Record*, part 4, vol. 7, pp. 142–163, March 1959.
- [3] A. Zaidi, I. Estella-Aguerri, and S. Shamai (Shitz), "On the Information Bottleneck Problems: Models, Connections, Applications and Information Theoretic Views," *Entropy*, vol. 22, no. 2, Art. no. 151, January 2020.
- [4] G. Zeitler, A. C. Singer, and G. Kramer, "Low-Precision A/D Conversion for Maximum Information Rate in Channels with Memory," *IEEE Transact. on Comms.*, vol. 60, no. 9, pp. 2511–2521, September 2012.
- [5] F. J. C. Romero and B. M. Kurkoski, "LDPC Decoding Mappings That Maximize Mutual Information," *IEEE Jnl. on Selected Areas in Comms.*, vol. 34, no. 9, pp. 2391–2401, September 2016.
- [6] M. Stark, J. Lewandowsky, and G. Bauch, "Information-Optimum LDPC Decoders with Message Alignment for Irregular Codes," in *IEEE Global Comms. Conf.*, Abu Dhabi, UAE, December 2018.
- [7] M. Stark, L. Wang, G. Bauch, and R. D. Wesel, "Decoding Rate-Compatible 5G-LDPC Codes With Coarse Quantization Using the Information Bottleneck Method," *IEEE Open Jnl. of the Comms. Society*, vol. 1, pp. 646–660, May 2020.
- [8] T. Monsees, O. Griebel, M. Herrmann, D. Wübben, A. Dekorsy, and N. Wehn, "Minimum-Integer Computation Finite Alphabet Message Passing Decoder: From Theory to Decoder Implementations towards 1 Tb/s," *Entropy*, vol. 24, no. 10, Art. no. 19, October 2022.
- [9] T. Monsees, D. Wübben, A. Dekorsy, O. Griebel, M. Herrmann, and N. Wehn, "Finite-Alphabet Message Passing Using Only Integer Operations for Highly Parallel LDPC Decoders," in *IEEE 23rd Int. Workshop on Signal Proc. Advances in Wireless Comm.*, July 2022.
- [10] J. Shao, Y. Mao, and J. Zhang, "Learning Task-Oriented Comm. for Edge Inference: An Information Bottleneck Approach," *IEEE Jnl. on Selected Areas in Comms.*, vol. 40, no. 1, pp. 197–211, January 2022.
- [11] F. Pezone, S. Barbarossa, and P. Di Lorenzo, "Goal-Oriented Communication for Edge Learning Based on the Information Bottleneck," in *IEEE Int. Conf. on Acoustics, Speech and Signal Proc.*, Singapore, Singapore, May 2022.
- [12] D. Gündüz, Z. Qin, I. E. Aguerri, H. S. Dhillon, Z. Yang, A. Yener, K. K. Wong, and C.-B. Chae, "Beyond Transmitting Bits: Context, Semantics, and Task-Oriented Communications," *IEEE Jnl. on Selected Areas in Comms.*, vol. 41, no. 1, pp. 5–41, January 2023.
- [13] E. Beck, C. Bockelmann, and A. Dekorsy, "Semantic Information Recovery in Wireless Networks," *Sensors*, vol. 23, no. 14, Art. no. 6347, July 2023.
- [14] I. Tal and A. Vardy, "How to Construct Polar Codes," *IEEE Transact. on Information Theory*, vol. 59, no. 10, pp. 6562–6582, October 2013.
- [15] M. Stark, A. Shah, and G. Bauch, "Polar Code Construction Using the Information Bottleneck Method," in *IEEE Wireless Comms. and Networking Conf. Workshops*, Barcelona, Spain, April 2018.
- [16] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, "Deep Variational Information Bottleneck," in *Int. Conf. on Learning Representations*, Toulon, France, April 2017.
- [17] A. Zaidi and I. Estella-Aguerri, "Distributed Deep Variational Information Bottleneck," in *IEEE Int. Workshop on Signal Proc. Advances in Wireless Comms.*, Atlanta, GA, USA, May 2020.
- [18] Q. Wang, C. Boudreau, Q. Luo, P.-N. Tan, and J. Zhou, "Deep Multi-View Information Bottleneck," in *SIAM Int. Conf. on Data Mining*, Calgary, Alberta, Canada, May 2019.
- [19] S. Hassanpour, T. Monsees, D. Wübben, and A. Dekorsy, "Forward-Aware Information Bottleneck-Based Vector Quantization for Noisy Channels," *IEEE Transact. on Comm.*, vol. 68, no. 12, pp. 7911–7926, December 2020.
- [20] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. John Wiley & Sons, 2006.
- [21] P. Harremoës and N. Tishby, "The Information Bottleneck Revisited or How to Choose a Good Distortion Measure," in *IEEE Int. Symposium on Information Theory*, Nice, France, June 2007.
- [22] D. Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods*. Academic Press, 1982.
- [23] J. H. Mathews and K. D. Fink, *Numerical Methods Using MATLAB*, 4th ed. Pearson Prentice Hall, 2004.
- [24] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *arXiv 2013*, arXiv:1312.6114.
- [25] B. Ghojogh, A. Ghodsi, F. Karray, and M. Crowley, "Factor Analysis, Probabilistic Principal Component Analysis, Variational Inference, and Variational Autoencoder: Tutorial and Survey," *arXiv 2021*, arXiv:2101.00734.
- [26] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [27] E. Jang, S. Gu, and B. Poole, "Categorical Reparameterization with Gumbel-Softmax," *arXiv 2016*, arXiv:1611.01144.
- [28] C. J. Maddison, A. Mnih, and Y. W. Teh, "The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables," *arXiv 2016*, arXiv:1611.00712.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," no. arXiv:1412.6980. *arXiv*, 2017.