


Leveraging Variational Autoencoders for Parameterized MMSE Estimation

Michael Baur , *Graduate Student Member, IEEE*, Benedikt Fesl , *Graduate Student Member, IEEE*,
and Wolfgang Utschick , *Fellow, IEEE*

Abstract—In this manuscript, we propose to use a variational autoencoder-based framework for parameterizing a conditional linear minimum mean squared error estimator. The variational autoencoder models the underlying unknown data distribution as conditionally Gaussian, yielding the conditional first and second moments of the estimand, given a noisy observation. The derived estimator is shown to approximate the minimum mean squared error estimator by utilizing the variational autoencoder as a generative prior for the estimation problem. We propose three estimator variants that differ in their access to ground-truth data during the training and estimation phases. The proposed estimator variant trained solely on noisy observations is particularly noteworthy as it does not require access to ground-truth data during training or estimation. We conduct a rigorous analysis by bounding the difference between the proposed and the minimum mean squared error estimator, connecting the training objective and the resulting estimation performance. Furthermore, the resulting bound reveals that the proposed estimator entails a bias-variance tradeoff, which is well-known in the estimation literature. As an example application, we portray channel estimation, allowing for a structured covariance matrix parameterization and low-complexity implementation. Nevertheless, the proposed framework is not limited to channel estimation but can be applied to a broad class of estimation problems. Extensive numerical simulations first validate the theoretical analysis of the proposed variational autoencoder-based estimators and then demonstrate excellent estimation performance compared to related classical and machine learning-based state-of-the-art estimators.

Index Terms—Parameter estimation, variational autoencoder, conditional mean estimator, generative model, inverse problem.

I. INTRODUCTION

GENERATIVE models (GMs) are a class of machine learning (ML) techniques designed to learn data distributions based on samples [1]. Instances of generative models are the

Manuscript received 29 March 2024; revised 2 July 2024 and 31 July 2024; accepted 2 August 2024. Date of publication 7 August 2024; date of current version 21 August 2024. This work was supported in part by Bavarian Ministry of Economic Affairs, Regional Development, and Energy within the project 6G Future Lab Bavaria, and in part by the Federal Ministry of Education and Research of Germany, under Project 16KISK002. The associate editor coordinating the review of this article and approving it for publication was Dr. Xiao Fu. (Corresponding author: Michael Baur.)

The authors are with the TUM School of Computation, Information and Technology, Technical University of Munich, 80333 Munich, Germany (e-mail: mi.baur@tum.de; benedikt.fesl@tum.de; utschick@tum.de).

Digital Object Identifier 10.1109/TSP.2024.3439097

Gaussian mixture model (GMM) [2, Ch. 9], variational autoencoder (VAE) [3], [4], generative adversarial network (GAN) [5], and, more recently, the score-based model [6]. A trained GM can generate unseen samples from the data distribution and often enables likelihood evaluation by providing a so-called generative prior. The generative prior is an approximation of the true data distribution and can be leveraged to solve sophisticated tasks such as inverse problems [7], dealing with recovering not directly observable parameters based on their noisy observations. Exemplarily, [8] uses a GAN for image reconstruction by adopting a compressed sensing (CS) framework and using the GAN as a generator. The approach is further extended in [9] to MRI images. Other image processing-related examples solving inverse problems with GMs involve phase retrieval [10] and blind image deconvolution [11]. In the context of wireless communications, generative priors find application in channel estimation (CE) [12], [13], [14], [15], [16], where the channel is estimated based on a noisy pilot observation, representing another instance of an inverse problem.

For the solution of an estimation task, a frequentist framework assumes the data to be deterministic and commonly constrains the estimator class to be unbiased in search for a minimum variance unbiased estimator [17]. In opposition, if a Bayesian approach is considered, it is well-known in estimation theory that the conditional mean estimator (CME) delivers minimum mean squared error (MMSE) estimates [18, Ch. 10]. Therefore, a (parameterized) Bayesian estimator's goal should be approximating the CME. Moreover, a fundamental aspect of the Bayesian framework is modeling the data as a random variable (RV), enabling the incorporation of a prior distribution into the estimation process. The result is an excellent estimation performance if the prior distribution accurately models the data, e.g., in the form of a generative prior. Therefore, GMs and Bayesian inference can be ideally combined to perform estimation tasks due to the distribution modeling abilities of the GMs.

A well-known GM that can be used for directly approximating the CME is the GMM [12]. However, connections between the CME and other GMs are yet to be discovered in the literature. Exemplarily, the GAN-based estimator from [13], [14], [15] is used in a CS-fashioned way and the score-based approach from [16] requires an iterative posterior sampling process, causing a massive computational complexity. A closely related GM to the GMM is the VAE. Both GMs maximize a lower bound to the data log-likelihood and introduce an artificial

latent space. Nevertheless, the GMM utilizes a discrete latent space, which limits its expressiveness. On the contrary, a VAE uses a continuous latent space, resulting in a better representation ability and a more flexible architectural design. The VAE intrinsically makes no assumption about the data distribution and was shown to work well in domains where it is traditionally challenging to derive statistical data models, e.g., in image processing [19].

In this work, we propose a VAE-parameterized estimator, combining a GM and classical estimation theory, with the following *contributions*:

- We model the analytically intractable data distribution as conditionally Gaussian (CG) with the help of the VAE, yielding conditional first and second moments to parameterize conditional linear minimum mean squared error (LMMSE) estimators given the latent representation and noisy observations. The conditional LMMSE estimators are mean squared error (MSE)-optimal and analytically tractable in closed-form due to the CG likelihood model.
- Since the VAE inherently makes no assumptions about the data distribution, the proposed estimation framework works independently of the adopted data distribution.
- We introduce a low-complexity estimator version based on a maximum a posteriori (MAP) estimate requiring only one neural network (NN) forward pass (MAP-VAE estimator). In contrast to many existing GM-based estimation frameworks for inverse problems, e.g. [13], [14], [15], [16], this procedure allows for a computationally efficient approximation of the CME, which is high-performing and robust, as a consequence of the VAE serving as a generative prior. Compared to classical algorithms for inverse problems such as least squares (LS) or approximate message passing (AMP), the proposed MAP-VAE estimator achieves significant performance gains.
- Three estimator variants are proposed, differing in the availability of ground-truth data during their training and estimation phases. The *VAE-real* variant is particularly noteworthy as it requires no access to samples of ground-truth data in either the training or estimation phases.
- We rigorously derive a bound on the performance gap between the MAP-VAE estimator and the CME, allowing for an interpretable estimation procedure. The bound connects the training objective of the VAE with the resulting estimation performance and reveals that the proposed estimator entails a bias-variance tradeoff that is well-known in the estimation literature.
- As an application example, we consider CE, offering a low-complexity implementation due to the structural properties of the CE problem. Our extensive numerical simulations first validate the theoretical analysis and then underline the superiority of the proposed VAE-based estimator variants compared to the baseline methods under various system configurations.

Moreover, we provide the following extensions in this manuscript compared to the preliminary results in [20]. The analyses in Sections III-A and III-D enhance the theoretical foundation of the VAE-based estimator's MSE-optimality. We

provide a more general treatment by providing a scheme for linear inverse problems of which multiple-input multiple-output (MIMO)-CE is a special instance. We make the training of the VAE signal-to-noise ratio (SNR)-independent, meaning that we use a single trained VAE for every SNR value, in opposition to [20], where an individual VAE is trained for every SNR value. Finally, the numerical simulations in this manuscript are more comprehensive.

The structure of this manuscript is as follows. Section II discusses the signal model and the general problem formulation and provides background information about the VAE. In Section III, we introduce the VAE-based estimator and its three variants, followed by the derivation and interpretation of the error bound between the proposed estimator and the CME. We discuss CE as an application example in Section IV. Numerical simulation results are presented in Section V, and we conclude this manuscript in Section VI.

Notation: We denote vectors and matrices as lower-case and upper-case bold-faced symbols, respectively. Element-wise multiplication is denoted as \odot , the all-zeros vector as $\mathbf{0}$, and the all-ones vector as $\mathbf{1}$. The vectorization operation $\text{vec}(\mathbf{G}) \in \mathbb{C}^{g_1 g_2}$ stacks the columns of $\mathbf{G} \in \mathbb{C}^{g_1 \times g_2}$ into a vector. The Kronecker product of two matrices $\mathbf{B} \in \mathbb{C}^{b_1 \times b_2}$ and $\mathbf{D} \in \mathbb{C}^{d_1 \times d_2}$ is $(\mathbf{B} \otimes \mathbf{D}) \in \mathbb{C}^{b_1 d_1 \times b_2 d_2}$.

II. PRELIMINARIES

A. Signal Model and Problem Formulation

We consider the generic linear inverse problem

$$\mathbf{y} = \mathbf{A}\mathbf{h} + \mathbf{n} \quad (1)$$

with the observation matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$ and additive noise $\mathbf{n} \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}, \mathbf{\Sigma})$. It is assumed that the matrix \mathbf{A} and the noise covariance $\mathbf{\Sigma}$ are given. The task is to recover \mathbf{h} based on \mathbf{y} . The design of \mathbf{A} is characteristic of the problem to be solved, e.g., in CE, \mathbf{A} represents the pilot allocation [12], [13], [14], [15]. For further examples, we refer to [21].

For the solution of (1), we aim to estimate \mathbf{h} based on the noisy observation \mathbf{y} . In the Bayesian framework, \mathbf{h} is a RV with an unknown prior $p(\mathbf{h})$. The goal is to minimize the MSE

$$\mathbb{E} \left[\|\mathbf{h} - \hat{\mathbf{h}}\|^2 \right] = \int \left[\int \|\mathbf{h} - \hat{\mathbf{h}}\|^2 p(\mathbf{h} | \mathbf{y}) d\mathbf{h} \right] p(\mathbf{y}) d\mathbf{y} \quad (2)$$

with the estimate $\hat{\mathbf{h}} \in \mathbb{C}^N$. For minimizing the MSE, minimizing the inner integral is sufficient due to $p(\mathbf{y}) \geq 0$. The minimizer is the well-known CME

$$\mathbb{E}[\mathbf{h} | \mathbf{y}] = \arg \min_{\hat{\mathbf{h}}} \mathbb{E} \left[\|\mathbf{h} - \hat{\mathbf{h}}\|^2 \right] \quad (3)$$

resulting in MSE-optimal estimates, cf. [18, Ch. 10] for details. More generally, the CME is the optimal predictor for all Bregman loss functions, of which the MSE is a special case [22]. Application of Bayes' rule to $p(\mathbf{h} | \mathbf{y})$ yields

$$\mathbb{E}[\mathbf{h} | \mathbf{y}] = \int \mathbf{h} \frac{p_{\mathbf{n}}(\mathbf{y} - \mathbf{A}\mathbf{h}) p(\mathbf{h})}{p(\mathbf{y})} d\mathbf{h}. \quad (4)$$

Note that $p_{\mathbf{n}}$ represents the noise probability density function (PDF). By inspection of (4), it becomes clear why the CME is




Fig. 1. Bayesian network illustrating the VAE decoder distribution $p_\theta(\mathbf{h}|\mathbf{z})$, encoder distribution $q_\phi(\mathbf{z}|\mathbf{y})$, and the known $p(\mathbf{y}|\mathbf{h}) = \mathcal{N}_C(\mathbf{A}\mathbf{h}, \mathbf{\Sigma})$.

difficult to compute. First, it requires access to the unknown and difficult-to-determine prior $p(\mathbf{h})$, necessitating an estimate of $p(\mathbf{h})$. Second, an approximation of the integral in (4) is required since, in general, there exists no closed-form solution. Another approach may involve directly approximating $p(\mathbf{h}|\mathbf{y})$, e.g., with Monte-Carlo sampling methods. Nevertheless, this would, in general, still necessitate calculating an intractable integral over $p(\mathbf{h}|\mathbf{y})$ to yield $E[\mathbf{h}|\mathbf{y}]$. Consequently, such procedures' applicability would be limited, especially in time-sensitive applications.

B. VAE Fundamentals

In a parametric approach, the parameterized likelihood model $p_\theta(\mathbf{h})$ approximates the unknown prior $p(\mathbf{h})$. One of the simplest parametric models is assuming a Gaussian prior, parameterized with the sample mean and covariance. The resulting parameterized CME approximation is the well-known LMMSE estimator [18, Ch. 10]. However, assuming a Gaussian prior is restrictive, causing the estimator to perform weakly if the true prior strongly deviates from a Gaussian distribution, which is the case in real-world systems. A way to significantly improve the expressiveness of the likelihood model while preserving the favorable properties of a Gaussian distribution is to let it hold only conditionally so the data is modeled as CG. The CG likelihood model has the form

$$\mathbf{h} | \mathbf{z} \sim p_\theta(\mathbf{h} | \mathbf{z}) = \mathcal{N}_C(\boldsymbol{\mu}_\theta(\mathbf{z}), \mathbf{C}_\theta(\mathbf{z})) \quad (5)$$

with the so-called latent vector $\mathbf{z} \in \mathbb{R}^{N_L}$ such that

$$p_\theta(\mathbf{h}) = \int p_\theta(\mathbf{h} | \mathbf{z}) p(\mathbf{z}) d\mathbf{z} \quad (6)$$

with a fixed $p(\mathbf{z})$. Besides its great properties in terms of expressiveness, the CG model in (5) will be a key aspect for deriving the VAE-parameterized estimator in Section III-A. Since $p_\theta(\mathbf{h} | \mathbf{z})$ is defined according to (5), θ also implicitly parameterizes the intractable posterior

$$p_\theta(\mathbf{z} | \mathbf{h}) = \frac{p_\theta(\mathbf{h} | \mathbf{z}) p(\mathbf{z})}{\int p_\theta(\mathbf{h} | \mathbf{z}) p(\mathbf{z}) d\mathbf{z}}. \quad (7)$$

A blueprint to obtain θ is given by the Bayesian network in Fig. 1 that parameterizes the joint PDF

$$p_\theta(\mathbf{y}, \mathbf{h}, \mathbf{z}) = p(\mathbf{y} | \mathbf{h}) p_\theta(\mathbf{h} | \mathbf{z}) p(\mathbf{z}). \quad (8)$$

Since the system model in (1) sets $p(\mathbf{y} | \mathbf{h}) = \mathcal{N}_C(\mathbf{A}\mathbf{h}, \mathbf{\Sigma})$, only $p_\theta(\mathbf{h} | \mathbf{z})$ must be learned in (8) if $p(\mathbf{z})$ is fixed. By inspection of Fig. 1, the following becomes apparent: all involved RVs are stochastically dependent, whereas \mathbf{y} and \mathbf{z} are conditionally




Fig. 2. Structure of a VAE with CG distributions for $q_\phi(\mathbf{z}|\mathbf{y})$ and $p_\theta(\mathbf{h}|\mathbf{z})$. The encoder and decoder each represent a NN.

independent given \mathbf{h} due to the local Markov property in the Bayesian network. Moreover, the learnable $q_\phi(\mathbf{z} | \mathbf{y})$ in Fig. 1 symbolizes the connection from \mathbf{y} to \mathbf{z} , which we will use to infer a \mathbf{z} based on \mathbf{y} in Section III-A because the true posterior $p_\theta(\mathbf{z} | \mathbf{y})$ is intractable, cf. (7).

The VAE [4] practically realizes the considered Bayesian network, for which an illustration is visible in Fig. 2. For the VAE training, $p_\theta(\mathbf{h})$ is typically decomposed as [23]

$$\log p_\theta(\mathbf{h}) = \mathcal{L}_{\theta, \phi}(\mathbf{h}) + D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{y}) \| p_\theta(\mathbf{z} | \mathbf{h})) \quad (9)$$

with the evidence lower bound (ELBO)

$$\mathcal{L}_{\theta, \phi}(\mathbf{h}) = E_{q_\phi} [\log p_\theta(\mathbf{h} | \mathbf{z})] - D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{y}) \| p(\mathbf{z})) \quad (10)$$

and the non-negative Kullback-Leibler (KL) divergence

$$D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{y}) \| p_\theta(\mathbf{z} | \mathbf{h})) = E_{q_\phi} \left[\log \left(\frac{q_\phi(\mathbf{z} | \mathbf{y})}{p_\theta(\mathbf{z} | \mathbf{h})} \right) \right]. \quad (11)$$

Note that $E_{q_\phi(\mathbf{z}|\mathbf{y})}[\cdot] = E_{q_\phi}[\cdot]$. The variational distribution $q_\phi(\mathbf{z} | \mathbf{y})$ is introduced aiming to approximate the intractable $p_\theta(\mathbf{z} | \mathbf{h})$ as can be seen in (11). Contrary to the VAE from [4], the variational distribution here is conditioned on \mathbf{y} instead of \mathbf{h} since the latter will be inaccessible during the estimation phase after the training. Consequently, a maximization of the ELBO is independent of (7), maximizes $\log p_\theta(\mathbf{h})$, as well as minimizes (11). In summary, a sufficiently trained VAE yields θ for the CG model in (5), as well as an approximation of the intractable posterior in (7) via $q_\phi(\mathbf{z} | \mathbf{y})$.

The remaining distributions in (10) are defined as:

$$q_\phi(\mathbf{z} | \mathbf{y}) = \mathcal{N}(\boldsymbol{\mu}_\phi(\mathbf{y}), \text{diag}(\boldsymbol{\sigma}_\phi^2(\mathbf{y}))), \quad (12)$$

$$p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (13)$$

Moreover, the VAE implements $p_\theta(\mathbf{h} | \mathbf{z})$ and $q_\phi(\mathbf{z} | \mathbf{y})$ as NNs. With these considerations, let us revisit Fig. 2. The encoder takes an observation \mathbf{y} and maps it to $\boldsymbol{\mu}_\phi(\mathbf{y})$ and $\boldsymbol{\sigma}_\phi(\mathbf{y})$ to obtain a reparameterized sample $\mathbf{z} = \boldsymbol{\mu}_\phi(\mathbf{y}) + \boldsymbol{\sigma}_\phi(\mathbf{y}) \odot \boldsymbol{\varepsilon}$. The sample \mathbf{z} is fed into the decoder to obtain $\boldsymbol{\mu}_\theta(\mathbf{z})$ and $\mathbf{C}_\theta(\mathbf{z})$ representing the first and second moments of $p_\theta(\mathbf{h} | \mathbf{z})$.

Due to the CG distributions, the terms in the ELBO can be calculated analytically, which is beneficial for the training of the VAE. The expectation term in (10) can be approximated with a single sample $\tilde{\mathbf{z}} \sim q_\phi(\mathbf{z} | \mathbf{y})$, i.e., $(-E_{q_\phi} [\log p_\theta(\mathbf{h} | \mathbf{z})])$ is replaced by the estimate

$$\log \det(\pi \mathbf{C}_\theta(\tilde{\mathbf{z}})) + (\mathbf{h} - \boldsymbol{\mu}_\theta(\tilde{\mathbf{z}}))^H \mathbf{C}_\theta^{-1}(\tilde{\mathbf{z}}) (\mathbf{h} - \boldsymbol{\mu}_\theta(\tilde{\mathbf{z}})). \quad (14)$$

The KL divergence $D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{y}) \| p(\mathbf{z}))$ in (10) results in

$$\frac{1}{2} (\mathbf{1}^T (-\log \sigma_\phi^2(\mathbf{y}) + \mu_\phi^2(\mathbf{y}) + \sigma_\phi^2(\mathbf{y})) - N_L). \quad (15)$$

By utilizing an expressive decoder NN and $\mathbf{C}_\theta(\mathbf{z})$ parameterization, we assume that a properly trained VAE where (5) holds delivers a $p_\theta(\mathbf{h})$ that well approximates $p(\mathbf{h})$. We will explicitly discuss conditional covariance matrix parameterizations for $\mathbf{C}_\theta(\mathbf{z})$ in Sections III-B and IV.

III. VAE-PARAMETERIZED ESTIMATOR

A. MMSE Estimation With the VAE

After its successful training, the VAE yields $p_\theta(\mathbf{h} | \mathbf{z})$ as CG according to (5). Recall the corresponding Bayesian network in Fig. 1 visualizing the dependencies of the involved RVs, which will be helpful for the following estimator derivation. Starting from (4), the law of total expectation enables reformulating the CME as [24, Sec. 4.3]:

$$\mathbb{E}[\mathbf{h} | \mathbf{y}] = \mathbb{E}_{p_\theta(\mathbf{z} | \mathbf{y})}[\mathbb{E}_\theta[\mathbf{h} | \mathbf{z}, \mathbf{y}] | \mathbf{y}], \quad (16)$$

where the inner expectation is with respect to $p_\theta(\mathbf{h} | \mathbf{z}, \mathbf{y})$. We neglect a possible approximation error between $p_\theta(\mathbf{h})$ and $p(\mathbf{h})$ in (16) as it is irrelevant for the estimator derivation. Since $p(\mathbf{h})$ is anyway inaccessible, an analysis of such an error is only possible empirically in terms of an MSE investigation, which will be done in Section V. Similar to $p_\theta(\mathbf{z} | \mathbf{h})$, $p_\theta(\mathbf{z} | \mathbf{y})$ and $p_\theta(\mathbf{h} | \mathbf{z}, \mathbf{y})$ are also implicitly parameterized by θ due to (7), the fixed prior $p(\mathbf{z})$ in (13) and the model in (1). Indeed, $p_\theta(\mathbf{z} | \mathbf{y})$ is generally inaccessible for the same reason as $p_\theta(\mathbf{z} | \mathbf{h})$, cf. (7). Since the encoder receives \mathbf{y} as input and \mathbf{h} is conditioned on \mathbf{z} according to (5), the training objective in (10) enforces \mathbf{n} and \mathbf{z} to be independent. Then, given (5), we obtain a closed-form expression for the inner expectation in (16) due to the CG property causing \mathbf{y} and \mathbf{h} to be jointly Gaussian given \mathbf{z} . Therefore, $\mathbb{E}_\theta[\mathbf{h} | \mathbf{z}, \mathbf{y}]$ results in [25]:

$$\mu_\theta(\mathbf{z}) + \mathbf{C}_\theta(\mathbf{z})\mathbf{A}^H(\mathbf{A}\mathbf{C}_\theta(\mathbf{z})\mathbf{A}^H + \Sigma)^{-1}(\mathbf{y} - \mathbf{A}\mu_\theta(\mathbf{z})), \quad (17)$$

where the matrix \mathbf{A} and vector \mathbf{y} belong to (1), and $\mu_\theta(\mathbf{z})$, $\mathbf{C}_\theta(\mathbf{z})$, and \mathbf{z} to (5). See Appendix A for a step-by-step derivation of (17).

It remains to solve the intractable outer expectation in (16). To this end, the approximation of $p_\theta(\mathbf{z} | \mathbf{h})$ via $q_\phi(\mathbf{z} | \mathbf{y})$ in (11) comes into play. Although (11) shows the approximation of $p_\theta(\mathbf{z} | \mathbf{h})$ instead of $p_\theta(\mathbf{z} | \mathbf{y})$, the parameter combination that maximizes $p_\theta(\mathbf{h} | \mathbf{z})$ in (10) also maximizes $p_\theta(\mathbf{y} | \mathbf{z})$ since the noise distribution is considered to be known and not subject to optimization, thus permitting the substitution. Consequently, by replacing $p_\theta(\mathbf{z} | \mathbf{y})$ with $q_\phi(\mathbf{z} | \mathbf{y})$ in (16),

$$\mathbb{E}[\mathbf{h} | \mathbf{y}] \approx \mathbb{E}_{q_\phi} [t_\theta(\mathbf{z}, \mathbf{y})], \quad t_\theta(\mathbf{z}, \mathbf{y}) = \mathbb{E}_\theta[\mathbf{h} | \mathbf{z}, \mathbf{y}]. \quad (18)$$

As we can easily obtain samples of $q_\phi(\mathbf{z} | \mathbf{y})$ with the help of the encoder, we can approximate $\mathbb{E}[\mathbf{h} | \mathbf{y}]$ using samples of the form $\mathbf{z}^{(k)} = \mu_\phi(\mathbf{y}) + \varepsilon^{(k)} \odot \sigma_\phi(\mathbf{y})$ where every $\varepsilon^{(k)}$ is a sample from $\mathcal{N}(\mathbf{0}, \mathbf{I})$, $k = 1, \dots, K$. Based on the samples $\mathbf{z}^{(k)}$ we can

approximate the MMSE estimator as a consequence of the law of large numbers [26]:

$$\hat{\mathbf{h}}_{\text{VAE}}^{(K)}(\mathbf{y}) = \frac{1}{K} \sum_{k=1}^K t_\theta(\mathbf{z}^{(k)}, \mathbf{y}), \quad \mathbf{z}^{(k)} \sim q_\phi(\mathbf{z} | \mathbf{y}), \quad (19)$$

where $t_\theta(\mathbf{z}^{(k)}, \mathbf{y})$ is evaluated with (17).

The estimator $\hat{\mathbf{h}}_{\text{VAE}}^{(K)}(\mathbf{y})$ generally becomes better for a large number of samples K . However, a large K is unwanted in a real-time system. It is desirable to reduce the complexity of the estimator as much as possible, which means that K should be low. To this end, we first obtain a MAP estimate for \mathbf{z} , which is $\mu_\phi(\mathbf{y})$ at the encoder output due to the Gaussianity of $q_\phi(\mathbf{z} | \mathbf{y})$, cf. (12). The MAP estimate is subsequently passed in a single step through the decoder to evaluate $t_\theta(\mathbf{z}, \mathbf{y})$. Consequently, we define the MAP-VAE estimator

$$\hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y}) = \hat{\mathbf{h}}_{\text{VAE}}^{(1)}(\mathbf{y}) = t_\theta(\mathbf{z}^{(1)} = \mu_\phi(\mathbf{y}), \mathbf{y}) \quad (20)$$

based on the MAP estimate $\mathbf{z}^{(1)} = \mu_\phi(\mathbf{y})$ from $q_\phi(\mathbf{z} | \mathbf{y})$. In Section III-D, we rigorously analyze the performance gap between the MAP-VAE estimator and the CME. Furthermore, in Section V-A, we compare $\hat{\mathbf{h}}_{\text{VAE}}^{(K)}(\mathbf{y})$ and $\hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y})$ for different K , demonstrating that their estimation quality is nearly identical. Unless otherwise stated, we use the MAP-VAE estimator in (20) for the numerical simulations.

B. Covariance Matrix Parameterization

According to (5), the VAE aims to learn a full covariance matrix for $\mathbf{h} | \mathbf{z}$. However, learning such a full matrix requires learning a large number of parameters, resulting in huge NNs. It is also known that covariances usually exhibit problem-specific structures, which can be exploited to drastically reduce the number of parameters to learn.

In this work, we consider equidistantly sampled wide-sense stationary (WSS) random processes covering a broad class of applications in signal processing:

- sensor array processing with equidistantly spaced sensors, e.g., for beamforming [27] or speech processing [28]
- CE in the spatial and time-frequency domain [12], [29]
- times series analysis in financial markets [30]

As a result of the WSS assumption, the covariance matrix is Toeplitz structured. The parameterization of a Toeplitz matrix is possible with an oversampled discrete fourier transform (DFT) matrix as demonstrated in [31], [32], [33]. However, if the covariance matrix dimensions are large, a circulant matrix can asymptotically approximate the Toeplitz covariance matrix [34]. By reasonably assuming that the VAE finds latent conditions that preserve the structural properties of the second moments [35], we can choose

$$\mathbf{C}_\theta(\mathbf{z}) = \mathbf{F}_N^H \text{diag}(\mathbf{c}_\theta(\mathbf{z})) \mathbf{F}_N, \quad \mathbf{c}_\theta(\mathbf{z}) \in \mathbb{R}_+^N, \quad (21)$$

parameterizing a circulant matrix, where $\mathbf{F}_N \in \mathbb{C}^{N \times N}$ is a DFT matrix. By choosing a covariance matrix parameterization in accordance with the structure of the actual covariance, assuming $p_\theta(\mathbf{h} | \mathbf{z})$ well-approximates the true distribution is reasonably motivated.

Circulant matrices have the advantage that they allow for a low-complexity and memory-efficient implementation and have already been used in previous work, cf. [29]. This can be seen in (21) since a positive and real-valued vector $\mathbf{c}_\theta(\mathbf{z})$ suffices to parameterize a full covariance matrix. Due to the DFT matrix, (21) is furthermore straightforwardly invertible in $\mathcal{O}(N \log N)$ time (by using FFTs), motivating its usage in the proposed VAE-based estimation framework.

C. Variants of VAE-Based Estimators

We present three possible estimator variants that leverage the VAE. All three estimators have in common that the VAEs can be trained offline before application. The estimators differ in their ground-truth data knowledge during the training and evaluation phase. A comprehensive overview of all variants with their losses will be shown in Section IV in Table I.

1) *VAE-genie*: To determine the full potential of our method, we assume $\mathbf{n} = \mathbf{0}$ in (1) for the encoder input while (17) is still evaluated with a non-zero noise realization. VAE-genie is supposed to exhibit the best estimation results among all variants because the $\mu_\theta(\mathbf{z})$ and $\mathbf{C}_\theta(\mathbf{z})$ at the decoder are inferred with the ground-truth data at the encoder and its latent representation. Although VAE-genie even has the potential to outperform the CME, as the ground-truth data acts as side information, this estimator is not applicable in practice, since it requires ground-truth knowledge during the evaluation phase. Instead, it can be a suitable benchmark result in a scenario where the optimal estimator is unknown and inaccessible. VAE-genie requires ground-truth data knowledge during the training and evaluation phase.

2) *VAE-noisy*: This estimator version directly relates to Fig. 2. The encoder receives the noisy observation \mathbf{y} as input with $\mathbf{n} \neq \mathbf{0}$. VAE-noisy only requires ground-truth data access during the training phase to compute (14) for its loss. During the evaluation phase, the mean $\mu_\phi(\mathbf{y})$ is obtained based on the noisy observation \mathbf{y} to compute (17), which is the reason for the name of this estimator. We expect that VAE-noisy delivers worse estimation quality than VAE-genie as VAE-genie has ground-truth knowledge in the evaluation phase. VAE-noisy is, in contrast, applicable in practice.

3) *VAE-real*: Similar to VAE-noisy, this estimator variant also receives \mathbf{y} as encoder input. The change compared to VAE-noisy happens at the decoder in Fig. 2 where VAE-real learns first and second moments for $p_\theta(\mathbf{y} | \mathbf{z})$ instead of $p_\theta(\mathbf{h} | \mathbf{z})$. However, to efficiently compute $\mathbb{E}_\theta[\mathbf{h} | \mathbf{z}, \mathbf{y}]$ we require a CG model for \mathbf{h} and not \mathbf{y} . As long as $\mathbb{E}[\mathbf{n}] = \mathbf{0}$, which is the case in (1), the mean of $\mathbf{y} | \mathbf{z}$ is $\mathbf{A}\mu_\theta(\mathbf{z})$. A simple workaround can determine the conditional covariance of $\mathbf{y} | \mathbf{z}$. While the VAE decoder continues to output $\mathbf{C}_\theta(\mathbf{z})$, e.g., according to (21), the matrix $\mathbf{A}\mathbf{C}_\theta(\mathbf{z})\mathbf{A}^H + \Sigma$ is used as covariance for $p_\theta(\mathbf{y} | \mathbf{z})$. Consequently, in (14), VAE-real replaces $\mu_\theta(\mathbf{z})$ with $\mathbf{A}\mu_\theta(\mathbf{z})$ and $\mathbf{C}_\theta(\mathbf{z})$ with $\mathbf{A}\mathbf{C}_\theta(\mathbf{z})\mathbf{A}^H + \Sigma$ during the training. This way, the decoder forces to substitute only the desired part, the conditional covariance $\mathbf{C}_\theta(\mathbf{z})$, which is used for the computation of (17). It should be noted that no ground-truth data is needed by VAE-real, neither during training nor during

evaluation. VAE-real is the most realistic estimator variant since noisy observations can be utilized to train the VAE. In contrast, access to ground-truth data during the training phase is usually related to a considerable additional effort and may sometimes be impractical.

D. MSE-Optimality and Conditional Bias-Variance Tradeoff

In this section, we provide a theoretical analysis of the introduced MAP-VAE estimator. Before establishing a bound on the difference between the MAP-VAE estimator and the CME, let us denote the decoder NN functions as

$$f_{\theta,1} : \mathbb{R}^{N_L} \rightarrow \mathbb{C}^N, \mathbf{z} \mapsto \mu_\theta(\mathbf{z}), \quad (22)$$

$$f_{\theta,2} : \mathbb{R}^{N_L} \rightarrow \mathcal{C}_+^N, \mathbf{z} \mapsto \mathbf{C}_\theta(\mathbf{z}), \quad (23)$$

where \mathcal{C}_+^N is the set of all $N \times N$ positive semi-definite (PSD) matrices (we consider the case of a full covariance matrix as this trivially includes all parameterized covariances discussed in Section III-B). In this section, we assume $\mathbf{y} = \mathbf{h} + \mathbf{n}$ to analyze the theoretical properties independent of \mathbf{A} .

Theorem 1: Consider a decorrelated observation $\mathbf{y} = \mathbf{h} + \mathbf{n}$ with $\mathbf{n} \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}, \varsigma^2 \mathbf{I})$ and let (5) and (16) hold. Further, assume the decoder neural network functions are Lipschitz continuous, i.e., for $i = \{1, 2\}$ and $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{N_L}$,

$$\|f_{\theta,i}(\mathbf{a}) - f_{\theta,i}(\mathbf{b})\|_2 \leq L_i \|\mathbf{a} - \mathbf{b}\|_2. \quad (24)$$

Then, the expected Euclidean distance between the CME (16) and the MAP-VAE estimator (20) is upper bounded as

$$\mathbb{E} \left[\left\| \mathbb{E}[\mathbf{h} | \mathbf{y}] - \hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y}) \right\|_2 \right] \leq (C_1 L_1 + C_2 L_2) \cdot \left(\sqrt{\text{tr}(\mathbf{C}_{p_\theta(\mathbf{z} | \mathbf{y})})} + \sqrt{\mathbb{E} \left[\left\| \mu_{p_\theta(\mathbf{z} | \mathbf{y})} - \mu_\phi(\mathbf{y}) \right\|_2^2 \right]} \right) \quad (25)$$

with the SNR-dependent factors

$$C_1 = \sqrt{\mathbb{E} \left[\frac{\varsigma^4}{(\xi_{\min}(\mathbf{y}) + \varsigma^2)^2} \right]}, \quad C_2 = \sqrt{\frac{N}{\varsigma^2}}. \quad (26)$$

where $\xi_{\min}(\mathbf{y})$ is the smallest eigenvalue of $\mathbf{C}_\theta(\mu_\phi(\mathbf{y}))$.

Proof: See Appendix B.

1) *MSE-Optimality*: Theorem 1 shows the expected distance of the MAP-VAE estimator to the CME only depends on the first two moments of the posterior distribution $p_\theta(\mathbf{z} | \mathbf{y})$, which is approximated by $q_\phi(\mathbf{z} | \mathbf{y})$. In particular, the bound is smaller if the first moments of $q_\phi(\mathbf{z} | \mathbf{y})$ match those of $p_\theta(\mathbf{z} | \mathbf{y})$, which can reasonably be assumed to be the case after successfully training the VAE, being a mild assumption as no restrictions to higher moments apply.

Moreover, the smaller the variances of $p_\theta(\mathbf{z} | \mathbf{y})$, the better the CME approximation. Intuitively, this means that the less stochastic a mapping from the observation to the latent space is, the better the MAP-VAE estimator performs. Let us consider the following setup to motivate the encoder variances to become small after training. Assume the input data is compressible onto a lower-dimensional manifold, i.e., a lossless compression mapping exists from \mathbb{C}^N to \mathbb{R}^{N_L} . In particular, this is known

to be fulfilled for *natural signals*, e.g., images or audio signals, wireless channels (especially in mmWave systems), or, in general, signals that exhibit a sparse representation through a dictionary. Then, a deterministic mapping exists into the latent space that can be learned by the VAE. In other words, there is no necessity for a stochastic mapping, and the variances in (25) can be chosen as zero without performance loss. This holds without restriction for the VAE-genie variant, where the encoder input is noiseless. For the VAE-noisy and VAE-real variants, although the latent encoding is trained to be stochastically independent of the noise, finding a deterministic mapping may be more intricate, especially in the low SNR regime, yielding a possibly higher encoder variance. We elaborate on this hypothesis in more detail for the example of channel estimation in MIMO systems in Section IV and show through simulations in Section V-A that the VAE's encoder variances are indeed converging towards zero during the training process.

Concluding the above discussion, the bound in Theorem 1 establishes a connection between the training of the VAE, purely relying on likelihood maximization, and the resulting MSE performance. Moreover, the impact of the latent dimension on the estimation performance is better interpretable. Thus, by a successful training of a well-designed VAE, the resulting parameterized estimator converges to the CME, thereby achieving a low MSE. We validate this proposition also through numerical results in Section V-A.

2) *Conditional Bias-Variance Tradeoff*: In addition to the above insights about the connection of the VAE's training and the resulting estimation performance, the constants C_1 and C_2 in (26) have a reciprocal behavior over the SNR and, in particular, are vanishing in the high and low SNR, respectively. That is,

$$\lim_{\zeta^2 \rightarrow 0} C_1 = 0, \quad \lim_{\zeta^2 \rightarrow \infty} C_2 = 0. \quad (27)$$

Interestingly, this can be interpreted as a *conditional bias-variance tradeoff* since $C_2 L_2$ in (25), addressing the contribution of the conditional covariances, vanishes in low SNR; moreover, $C_1 L_1$, attributed to the conditional means, vanish in high SNR, cf. (55). Thus, the parameterized conditional covariance quality is less critical in the low SNR regime, as the parameterized LMMSE estimator relies more on the conditional first moment and vice versa in the high SNR regime. Consequently, the respective error terms have less impact on the bound to the CME. The entailment of such a *conditional bias-variance tradeoff* is a highly desirable property of the proposed estimator as it serves as a regularization for the estimation performance and allows for great interpretability. Moreover, the analysis holds without restriction for all discussed estimator variants in Section III-B and all parameterized conditional covariance matrices since we made no assumptions about their structural properties.

IV. EXAMPLE APPLICATION: CHANNEL ESTIMATION

In this work, we consider MIMO CE as an application example. In a MIMO communications system, the transmitter with N_{tx} antennas sends N_{p} pilots to the receiver with N_{rx}

antennas for estimating the channel matrix $\mathbf{H} \in \mathbb{C}^{N_{\text{rx}} \times N_{\text{tx}}}$. More precisely, the noisy observations

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathbf{N} \in \mathbb{C}^{N_{\text{rx}} \times N_{\text{p}}} \quad (28)$$

are obtained at the receiver with the pilot matrix $\mathbf{X} \in \mathbb{C}^{N_{\text{tx}} \times N_{\text{p}}}$ and noise matrix \mathbf{N} . After vectorizing (28), the relation to (1) becomes apparent. Consequently, $\mathbf{y} = \text{vec}(\mathbf{Y})$, $\mathbf{h} = \text{vec}(\mathbf{H})$, $\mathbf{A} = (\mathbf{X}^T \otimes \mathbf{I})$, and $\mathbf{n} = \text{vec}(\mathbf{N})$. Further, $M = N_{\text{rx}}N_{\text{p}}$ and $N = N_{\text{rx}}N_{\text{tx}}$. We investigate the uplink of a communications system where the mobile terminal (MT) transmits to the base station (BS) with $N_{\text{tx}} < N_{\text{rx}}$. However, the proposed framework can also be applied to the downlink since \mathbf{A} has a comparable structure.

We assume that the BS and MT are both equipped with a uniform linear array (ULA) with half-wavelength spacing. Note that a different array structure or antenna spacing can be straightforwardly reflected by the VAE's parameterized covariance at the decoder output. Furthermore, we consider the fully determined case of (28), i.e., $N_{\text{p}} = N_{\text{tx}}$. We utilize DFT pilots, resulting in a unitary \mathbf{X} , which results in a unitary \mathbf{A} . Moreover, we assume $\Sigma = \zeta^2 \mathbf{I}$ with given ζ^2 . Therefore, we perform a LS estimate of (1) to interpret it as a denoising task relating directly to the theoretical analysis in the previous section. The underdetermined case involving a wide \mathbf{A} is investigated in [32] and the uniform rectangular array (URA) case at the BS in [33] covering more advanced setups. The works [32], [33] demonstrate a superior performance of the VAE-based estimators, highlighting the framework's versatile applicability under various system configurations.

Due to the common WSS assumptions in wireless communications [36, Sec. 2.6], the transmit- and receive-side covariance matrices at the BS and MT side, respectively, are Toeplitz structured, which are approximated by circulant matrices as explained in Section III-B. When additionally assuming uncorrelated scattering in the vicinity of the transmitter and receiver, we can decompose the channel covariance matrix (CCM) into the Kronecker product of the transmit- and receive-side circulant-structured covariance matrices [37]:

$$\mathbf{C}_{\theta}(\mathbf{z}) = \mathbf{Q}^H \text{diag}(c_{\theta}(\mathbf{z}))\mathbf{Q}, \quad c_{\theta}(\mathbf{z}) \in \mathbb{R}_+^N \quad (29)$$

where $\mathbf{Q} = (\mathbf{F}_{N_{\text{tx}}} \otimes \mathbf{F}_{N_{\text{rx}}})$. In (29), $\mathbf{C}_{\theta}(\mathbf{z})$ is a block-circulant matrix, possessing the same favorable attributes regarding memory efficiency and low-complexity as an ordinary circulant matrix. For a single-input multiple-output (SIMO) system, which implies $N_{\text{tx}} = 1$, (29) simplifies to (21).

A. Training Loss and Network Architecture

In principle, we train a VAE with the loss in (10), and, after the training, perform CE as described in Section III-A. Indeed, we can simplify (10) as a result of the circulant parameterization. Let $\mathbf{h}_Q = \mathbf{Q}\mathbf{h}$, then the negative decoder likelihood in (14) can be expressed as

$$N \log \pi + \mathbf{1}^T (\lambda_{\theta}(\tilde{\mathbf{z}}) \odot |\mathbf{h}_Q - \mathbf{Q}\boldsymbol{\mu}_{\theta}(\tilde{\mathbf{z}})|^2 - \log \lambda_{\theta}(\tilde{\mathbf{z}})) \quad (30)$$

TABLE I
OVERVIEW OF THE PROPOSED VAE-BASED ESTIMATOR VARIANTS

Variant	VAE Encoder Input	Training Loss $\mathcal{L}_{\theta,\phi}$ (Relates to One Batch Element)
VAE-genie	ground-truth channel	$\mathbf{1}^T [\lambda_{\theta}(\tilde{z}) \odot h_Q - Q\mu_{\theta}(\tilde{z}) ^2 - \log \lambda_{\theta}(\tilde{z}) - \log \sigma_{\phi}(\mathbf{h}) + 0.5(\mu_{\phi}^2(\mathbf{h}) + \sigma_{\phi}^2(\mathbf{h}))]$
VAE-noisy	noisy observation	$\mathbf{1}^T [\lambda_{\theta}(\tilde{z}) \odot h_Q - Q\mu_{\theta}(\tilde{z}) ^2 - \log \lambda_{\theta}(\tilde{z}) - \log \sigma_{\phi}(\mathbf{y}) + 0.5(\mu_{\phi}^2(\mathbf{y}) + \sigma_{\phi}^2(\mathbf{y}))]$
VAE-real	noisy observation	$\mathbf{1}^T [(c_{\theta}(\tilde{z}) + \varsigma^2 \mathbf{1})^{-1} \odot y_Q - Q\mu_{\theta}(\tilde{z}) ^2 + \log(c_{\theta}(\tilde{z}) + \varsigma^2 \mathbf{1}) - \log \sigma_{\phi}(\mathbf{y}) + 0.5(\mu_{\phi}^2(\mathbf{y}) + \sigma_{\phi}^2(\mathbf{y}))]$




Fig. 3. Detailed illustration of the different layers constituting our VAE implementation. The real and imaginary parts of the input are stacked as convolutional channels (CCs) and processed. The colored arrows represent different layers or layer compositions. Purple \Rightarrow stands for a 1×1 CL, orange \Rightarrow for a block of a convolutional layer (CL), batch normalization (BN) layer, and ReLU activation function, gray \Rightarrow for a reshaping layer (RL), green \Rightarrow for a linear layer (LL), and red \Rightarrow for a block of a transposed CL, BN layer, and ReLU activation function.

with $\lambda_{\theta}(z) = c_{\theta}^{-1}(z)$ and the element-wise absolute value $|\cdot|$. Eq. (14) reduces the numerical complexity during the training process because it avoids the inversion of a full covariance matrix compared to (30). What is more, we utilize the LS estimate of (1) (or \mathbf{h} for VAE-genie) multiplied with Q as encoder input. Thus, the encoder input is effectively transformed into the angular or beamspace domain [38, Sec. 7.3], which is known to be sparse or highly compressible in massive MIMO systems, especially in mmWave systems. This validates the hypothesis of having a deterministic compression mapping that can be learned through the encoder in Section III-D. Therefore, a performance of the MAP-VAE estimator close to the CME can be expected, which is also seen later in Section V.

Combining every aspect from this section, the reformulated training objective that VAE-noisy is supposed to minimize is:

$$\mathcal{L}_{\theta,\phi} = \mathbf{1}^T [\lambda_{\theta}(\tilde{z}) \odot |h_Q - Q\mu_{\theta}(\tilde{z})|^2 - \log \lambda_{\theta}(\tilde{z}) - \log \sigma_{\phi}(\mathbf{y}) + 0.5(\mu_{\phi}^2(\mathbf{y}) + \sigma_{\phi}^2(\mathbf{y}))]. \quad (31)$$

The argument and constants are omitted for brevity and \tilde{z} is a sample from $q_{\phi}(z|\mathbf{y})$. Since VAE-genie has the ground-truth channel as encoder input, the training loss for this variant replaces $\mu_{\phi}(\mathbf{y})$ and $\sigma_{\phi}(\mathbf{y})$ with $\mu_{\phi}(\mathbf{h})$ and $\sigma_{\phi}(\mathbf{h})$, respectively. Apart from that, the training loss is identical to (31). For the training of VAE-real, in (31), h_Q is replaced with $\mathbf{y}_Q = Q\mathbf{A}^H\mathbf{y}$, and $\lambda_{\theta}(z)$ with $(c_{\theta}(z) + \varsigma^2\mathbf{1})^{-1}$. Table I summarizes the proposed estimator variants with an overview regarding the respective encoder input and training loss. In each case, the training loss refers to a single batch element.

We briefly describe our VAE implementation at this point. The simulation code with the corresponding architectures is also publicly available.¹ Fig. 3 illustrates the VAE implementation. The arrow colorings in Fig. 3 symbolize different layers

or layer compositions. On the left, it is visible that the real and imaginary parts of the encoder input are stacked as convolutional channels (CCs). As a first block, the purple arrow \Rightarrow represents a 1×1 CL that maps to a higher number of CCs, which is different for every system configuration. Subsequently, three orange arrows \Rightarrow follow, representing a block of a CL, a BN layer, and a ReLU activation function. In each CL, the CC amount at the output is multiplied by a factor of 1.75. After a reshaping layer (RL) and linear layer (LL), symbolized by the gray arrow \Rightarrow and green arrow \Rightarrow , respectively, we arrive at the latent space. The reparameterized sample z is fed into the decoder, which is a symmetrically flipped version of the encoder. The red arrows \Rightarrow symbolize blocks of a transposed CL, a BN layer, and a ReLU activation function. At the output, we have a sample with three CCs that is fed into an RL and LL to produce the decoder outputs. We use exponential functions to enforce strictly positive values as it is required for σ_{ϕ} and c_{θ} .

The number of CCs, kernel size, and latent dimension are different for every system configuration and are found by a random search over the hyperparameter space by searching for the combination that yields the highest value for (30) [39]. We perform the random search with the help of the *Tune* package [40]. We use 2D CLs in the MIMO case and 1D CLs in the SIMO case. A batch size of 128, a learning rate of $7 \cdot 10^{-4}$ in combination with Adam [23], and a stride of two in the CLs are used. Overall, we found the estimation performance is robust regarding the selected VAE architecture as long as it contains enough layers and model parameters for the considered problem. We implement the NNs with PyTorch and refer the reader to the simulation code for further details. Additionally, we experimented with BN and its variants to determine how we can achieve the best performance [41], [42], [43]. We achieve the best performance with BN as is proposed in [41]. The only

¹<https://github.com/tum-msv/vae-estimator>.

important point is to consider a large enough batch size to limit the variance of the stochastic gradient. We additionally use the method of free bits during the training as described in [23].

B. Computational Complexity

In this section, we discuss the computational complexity of the proposed estimator. The procedure to determine $\hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y})$ can be split into two parts. The first step is a forward pass through the VAE to acquire $\boldsymbol{\mu}_\theta(\mathbf{z})$ and $\mathbf{C}_\theta(\mathbf{z})$. The second step is the evaluation of $t_\theta(\mathbf{z}, \mathbf{y})$ in (17) with given $\boldsymbol{\mu}_\theta(\mathbf{z})$ and $\mathbf{C}_\theta(\mathbf{z})$. The computational complexity of the first step is tied to the VAE architecture in Fig. 3. Since all layers exhibit a different complexity, we need a complexity bound for which two aspects are relevant. First, a CL requires $\mathcal{O}(RN)$ time, with R being the product of the number of parameters divided by the stride in the CL. Second, the final LL requires $\mathcal{O}(N^2)$ time. The remaining layers exhibit less complexity than the CLs and final LL. Although R should be increased if N grows, R arguably does not show more than linear growth in N . In conclusion, utilizing $\mathcal{O}(N^2)$ as complexity bound per layer is reasonable. For the D layers of the VAE forward pass, this makes an overall complexity of $\mathcal{O}(DN^2)$.

We come to the second step of obtaining $\hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y})$, which is the evaluation of $t_\theta(\mathbf{z}, \mathbf{y})$. In principle, the inversion of $\mathbf{A}\mathbf{C}_\theta(\mathbf{z})\mathbf{A}^H + \varsigma^2\mathbf{I}$ dominates the complexity. Let us inspect $\mathbf{A}\mathbf{C}_\theta(\mathbf{z})\mathbf{A}^H = \mathbf{A}\mathbf{Q}^H \text{diag}(\mathbf{c}_\theta(\mathbf{z}))\mathbf{Q}\mathbf{A}^H$ in more detail. If we assume to have unitary pilots and set $\tilde{\mathbf{A}} = \mathbf{A}\mathbf{Q}^H$ we can show that $\tilde{\mathbf{A}}\tilde{\mathbf{A}}^H = \mathbf{I}$ and $\tilde{\mathbf{A}}^H\tilde{\mathbf{A}} = \mathbf{I}$ holds, so $\tilde{\mathbf{A}}$ is unitary. Hence, the inverse of $\mathbf{A}\mathbf{C}_\theta(\mathbf{z})\mathbf{A}^H$ is $\tilde{\mathbf{A}} \text{diag}(\mathbf{c}_\theta^{-1}(\mathbf{z}))\tilde{\mathbf{A}}^H$. We can therefore simplify the estimate $t_\theta(\mathbf{z}, \mathbf{y})$ as in (17) to

$$\boldsymbol{\mu}_\theta(\mathbf{z}) + \mathbf{Q}^H \text{diag}(\mathbf{1} + \mathbf{c}_\theta(\mathbf{z}) \odot \varsigma^{-2}\mathbf{1})\mathbf{Q}(\mathbf{A}^H\mathbf{y} - \boldsymbol{\mu}_\theta(\mathbf{z})) \quad (32)$$

whose complexity is $\mathcal{O}(N \log N)$ due to multiplying with \mathbf{Q} . As can be seen from our elaborations above, the evaluation of the VAE requires $\mathcal{O}(DN^2)$ time, which outweighs the evaluation time of $\mathcal{O}(N \log N)$ for (32). Additionally, many potentials exist to reduce the VAE complexity, e.g., with pruning [44]. Moreover, the computations in the VAE are highly parallelizable due to the CLs, which mitigates the $\mathcal{O}(DN^2)$ complexity.

C. Channel Models

We consider different channel models in this work to validate the proposed methods. The 3rd Generation Partnership Project (3GPP) defines an urban macrocell spatial channel model which computes the transmit-side CCM as [45]:

$$\mathbf{C}_{\delta, \text{tx}} = \int_{-\pi}^{\pi} g_{\text{tx}}(\vartheta; \boldsymbol{\delta}) \mathbf{a}_{\text{tx}}(\vartheta) \mathbf{a}_{\text{tx}}^H(\vartheta) d\vartheta. \quad (33)$$

The vector $\mathbf{a}_{\text{tx}}(\vartheta)$ denotes the transmit array steering vector, which is $[1, \exp(j\pi \sin(\vartheta)), \dots, \exp(j\pi(N_{\text{tx}} - 1) \sin(\vartheta))]^H$ in the case of a ULA. Analogously, the receive-side CCM $\mathbf{C}_{\delta, \text{rx}}$ is obtained. The function $g_{\text{tx}}(\cdot; \boldsymbol{\delta})$ describes an angular power spectrum relating to the involved propagation clusters and is parameterized by the vector $\boldsymbol{\delta}$, following a prior distribution

$p(\boldsymbol{\delta})$ that accounts for the involved path gains and angles. More precisely, $g_{\text{tx}}(\cdot; \boldsymbol{\delta})$ is a mixture of Laplace densities whose standard deviations represent the angular spreads, cf. [29] for more details. The CCM for \mathbf{h} in (1) is determined as $\mathbf{C}_\delta = (\mathbf{C}_{\delta, \text{tx}} \otimes \mathbf{C}_{\delta, \text{rx}})$, under the assumption of uncorrelated scattering [37]. For a large number of sub-paths per propagation cluster, which is a common assumption for sub-6 GHz frequency bands, a CG channel distribution is well-motivated by the central limit theorem. A channel realization can, thus, be obtained with $\mathbf{h} | \boldsymbol{\delta} \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}, \mathbf{C}_\delta)$ [38, Sec. 4.2]. Accordingly, a correlated Rayleigh fading model is enforced that only holds conditionally, meaning that every channel is individually associated with a different set of path gains and angles contained in $\boldsymbol{\delta}$ representing different propagation clusters. Note that \mathbf{C}_δ is different for every channel realization, causing $p(\mathbf{h})$ to be non-Gaussian.

The QuaDRiGa channel simulator allows for the simulation of realistic channels with spatial consistency and time evolution [46], [47]. MIMO channel matrices are modeled as a superposition of in total L propagation paths such that $\mathbf{H} = \sum_{\ell=1}^L \mathbf{G}_\ell \exp(-2\pi j f_c \tau_\ell)$ where the carrier frequency is denoted as f_c and the delay of the ℓ -th path as τ_ℓ . The entries of the matrix \mathbf{G}_ℓ represent the complex-valued gain between every antenna pair caused by the path loss, antenna radiation pattern, and polarization. We use version 2.6.1 of QuaDRiGa to simulate channels at a frequency of 6 GHz in an urban macrocell scenario. We place the BS at a height of 25 m, and it covers a sector of 120° . Twenty percent of the MTs are outdoors 1.5 m above the ground at a distance between 35 and 500 m from the BS. The remaining eighty percent are situated indoors at different floor levels. We consider a line of sight (LOS) propagation environment, where $L = 37$. We equip the BS with “3GPP-3D” antennas and the MTs with omnidirectional antennas. After generation, the channels are post-processed to normalize the path gain. Compared to the 3GPP channel model, which is fully stochastic, the QuaDRiGa simulator is of a geometric nature. QuaDRiGa determines channel realizations by a geometric simulation in a randomized and approximately realistic BS environment. The QuaDRiGa model enables us to highlight that the proposed framework works independently of the adopted channel model.

D. Related Channel Estimators

This section presents related channel estimators as baselines for the numerical simulations in Section V-B. In the case of the 3GPP channel model from Section IV-C, we have access to the true CCM \mathbf{C}_δ . This allows us to evaluate a genie covariance-based estimator (genie-cov) [29], which is given by the LMMSE formula

$$\hat{\mathbf{h}}_{\text{genie-cov}}(\mathbf{y}) = \mathbf{C}_\delta \mathbf{A}^H (\mathbf{A}\mathbf{C}_\delta \mathbf{A}^H + \boldsymbol{\Sigma})^{-1} \mathbf{y}. \quad (34)$$

This estimator uses utopian genie knowledge to acquire \mathbf{C}_δ .

A practical estimator can be based on the sample covariance matrix $\hat{\mathbf{C}} = \frac{1}{T_r} \sum_{i=1}^{T_r} \mathbf{h}_i \mathbf{h}_i^H$ for T_r samples in the training dataset. The corresponding estimator reads as

$$\hat{\mathbf{h}}_{\text{global-cov}}(\mathbf{y}) = \hat{\mathbf{C}} \mathbf{A}^H (\hat{\mathbf{C}} \mathbf{A}^H + \boldsymbol{\Sigma})^{-1} \mathbf{y}. \quad (35)$$

LS estimation is another comparison method we investigate in our simulations. An LS estimate can be obtained as $\hat{\mathbf{h}}_{\text{LS}}(\mathbf{y}) = \mathbf{A}^H \mathbf{y}$.

CS-based CE techniques are another prominent topic in the literature. Especially regarding millimeter waves, CS algorithms are potentially interesting candidates [48]. This work compares the proposed estimators with the AMP algorithm [49], [50]. As a dictionary for AMP, we use a two times over-sampled DFT matrix.

We also want to compare the proposed estimators with current ML-based channel estimators. A recently proposed method exploits structural information of the MMSE estimator to design a neural network-based estimator for the SIMO signal model [29]. The derivation leads to a convolutional neural network with ReLU activation function, so we call this estimator CNN. The extension of [29] to the MIMO case is proposed in [51], to which we also refer in our simulations.

The last comparison method in this section, also recently proposed, is based on a GMM [12]. The idea is to fit a GMM to the underlying channel distribution and parameterize a channel estimator with the help of the GMM, representing an estimator based on a generative prior. We fit a GMM with 128 mixture components for all simulations and a restriction on the fitted covariances such that they are block-circulant.

V. SIMULATION RESULTS

This section presents the CE results based on numerical simulations. We create 200,000 channel realizations for every system configuration in the upcoming section representing a randomly sampled realistic BS environment. The channels are divided into $T_r = 180,000$ training, $T_v = 10,000$ validation, and $T_e = 10,000$ test samples. The channels are normalized such that $E[\|\mathbf{h}\|^2] = N$. In our experiments, we calculate the normalized mean squared error (NMSE) as $\frac{1}{T_e N} \sum_{i=1}^{T_e} \|\mathbf{h}_i - \hat{\mathbf{h}}_i\|^2$ for the test dataset, where we denote the i -th channel realization and corresponding estimate as \mathbf{h}_i and $\hat{\mathbf{h}}_i$, respectively. Accordingly, we define the SNR as N_{tx}/σ^2 . We train the VAEs for a range of SNR values between -19 and 39 dB. The proposed estimators are, therefore, SNR-independent. During the training of VAE-noisy and VAE-real, we sample new realizations \mathbf{n} after every epoch. We train the VAEs until (30) does not improve for 100 consecutive epochs on the validation dataset. If not stated otherwise, $N_L = 16$ for one propagation cluster and $N_L = 32$ in all other cases.

A. Numerical Convergence Analysis

At first, we illustrate the training progress of the VAE-noisy variant for the 3GPP SIMO signal model with three propagation clusters at an SNR of 10 dB in Fig. 4. The VAE-genie and VAE-real variants exhibit a similar behavior, so we only display VAE-noisy here. It can be observed that most of the training progress happens in the first 20 epochs. Interestingly, an increase of the REC term from (30) coincides with a decrease of the NMSE, which indicates that the VAE learns to properly model the data. Moreover, this validates the theoretical analysis in Theorem 1




Fig. 4. Training of the VAE-noisy variant for the 3GPP channel model (SIMO case) with three propagation clusters at an SNR of 10 dB. ELBO refers to the complete training loss in (31), REC to the negative of (30), and KL to (15). REC is plotted including the in (30) omitted constants.




Fig. 5. Normalized MSE for different numbers of training samples at an SNR of 10 dB for the 3GPP channel model (SIMO case) with three propagation clusters and 128 antennas at the receiver. The dotted lines display the achieved result with the complete training dataset of 180,000 samples.

that showed a smaller gap to the CME and, thus, a lower NMSE for a VAE that better matches the first moments of the posterior distributions, which is achieved during a successful training.

Further, we investigate two critical quantities of the model selection process: the training dataset's size and the latent space's dimensionality. Regarding the size of the training dataset, a larger size is likely to lead to better estimation results. Fig. 5 provides insights into this matter. We display the estimation results of the test dataset for the three proposed variants of VAE-based estimators depending on the size of the training dataset. The 3GPP channel model (SIMO case) with three propagation clusters and 128 antennas at the receiver is used in Fig. 5. As dotted lines, we display the attained estimation result for the complete training dataset of 180,000 samples. We also show the estimation performance of the genie-cov estimator in blue. It is visible that the most progress is reached from 10^2 to 10^4 training samples. More than 10^4 training samples only lead to minor NMSE improvements for all three types of VAE-based estimators.

The influence of the dimensionality of the latent space on the estimation result is less apparent than the size of the training dataset. We illustrate the NMSE for dimensionalities in the range $[4, 32]$ for the 3GPP SIMO channel model with 128 antennas at the receiver in Fig. 6 by considering one and three propagation clusters. The NMSE is nearly constant for the case with one propagation cluster. In contrast, the NMSE decreases from dimensionality 4 to 16 for the three propagation clusters case and saturates for larger dimensional latent spaces. In practice, the operator must select an ample enough latent space to obtain a desirable performance.




Fig. 6. Normalized MSE for different sizes of the latent space at an SNR of 10 dB for the 3GPP channel model (SIMO case) with one or three propagation clusters and 128 antennas at the receiver.




Fig. 8. Normalized MSE for different numbers of samples K drawn in the latent space for the evaluation of $\hat{\mathbf{h}}_{\text{VAE}}^{(K)}(\mathbf{y})$ at an SNR of 10 dB for the 3GPP channel model (SIMO case) with three propagation clusters and 128 antennas at the receiver. The dashed lines represent the estimate $\hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y})$ in (20).




Fig. 7. Trace of encoder variances and NMSE over training epochs on the 3GPP channel model (SIMO case) with 128 antennas, one propagation cluster, 10 dB SNR, and $N_L = 4$. The NMSE curves are displayed as dashed.

Theorem 1 in Section III-D described that the convergence of the MAP-VAE estimator to the CME depends on the vanishing of $\text{tr}(\mathbf{C}_{p\theta}(z|\mathbf{y}))$. To this end, we analyze the trace of the encoder variance, i.e., the summed variance σ_ϕ^2 of q_ϕ , which optimally is a good approximation of $\text{tr}(\mathbf{C}_{p\theta}(z|\mathbf{y}))$, over the training epochs on the validation dataset in Fig. 7. We evaluate VAE-genie and VAE-noisy with $N_L = 4$ on the 3GPP channel model (SIMO case) with 128 antennas and one propagation cluster at an SNR of 10 dB. It is visible that both encoder variances are decreasing in a comparable way as the NMSE, indicating that lower traces improve the NMSE. The noise variance detrimentally influences VAE-noisy’s encoder variance trace since it is always higher than that of VAE-genie. In conclusion, since the wireless channel data is expected to be sparse or compressible in the angular domain, cf. Section IV, the VAE indeed aims to find a less stochastic encoder mapping during training. This is in agreement with the argumentation in Section III-D and the observation of a decreasing NMSE of the parameterized estimator during training in Fig. 4.

As pointed out in Section III-A, we approximate the CME with the MAP-VAE estimator by only forwarding the latent mean vector $\mu_\phi(z)$ to approximate the outer expectation in (16). It is interesting to see the NMSE performance for different numbers of K samples from $q_\phi(z|\mathbf{y})$ to compute $\hat{\mathbf{h}}_{\text{VAE}}^{(K)}(\mathbf{y})$ from (19). Fig. 8 provides such an analysis by showing the NMSE for different numbers of latent samples. As dashed lines, we display the MAP-VAE estimator, which only uses the single sample $\mu_\phi(z)$ for the input of the VAE’s decoder. We observe that VAE-real benefits the most from more samples. For VAE-genie and VAE-noisy, there are only slight improvements

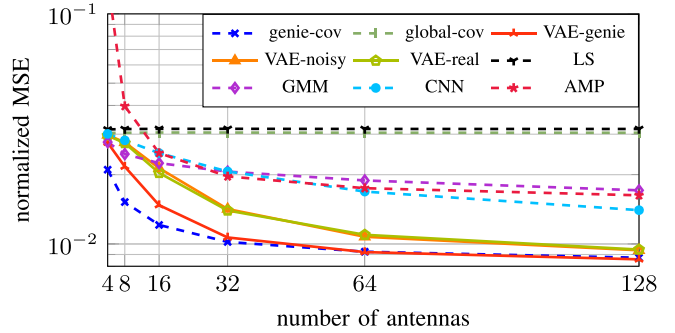


Fig. 9. Normalized MSE for the 3GPP channel model (SIMO case) with three propagation clusters for different numbers of antennas at the receiver at an SNR of 15 dB. The proposed methods are displayed with solid linestyles.

present. Interestingly, only taking the mean value, representing the MAP-VAE estimator, delivers an estimation performance of about $K = 64$ samples for VAE-real. The excellent performance of the MAP-VAE estimator is a supporting argument for the theoretical analysis of the estimator in Theorem 1 that predicts a small distance from the CME if the VAE is well-trained.

B. Normalized MSE Results

We begin with an NMSE investigation for the 3GPP channel model (SIMO case) with three propagation clusters and varying numbers of antennas at the receiver at an SNR of 15 dB in Fig. 9. The illustration shows that the proposed VAE-based methods need a sufficiently large amount of antennas to develop their full potential. From 16 antennas on, the VAE-based methods outperform the baselines and exhibit increasing performance gains if more antennas are considered. All other baselines perform significantly worse than the proposed methods in the large antenna regime. It is also visible that VAE-genie converges to the genie-cov curve. Surprisingly, the VAE-real variant is almost on par with the VAE-noisy variant, although no ground-truth data is available for its training, underlining the strong performance of the VAE as a generative prior even in cases with imperfect training data.

A massive amount of antennas is significant for prospective communications systems. Hence, we investigate the large




Fig. 10. Normalized MSE for the 3GPP channel model (SIMO case) with one propagation cluster and 128 antennas at the receiver. The proposed methods are displayed with solid linestyle.




Fig. 11. Normalized MSE for the QuaDRiGa channel model (SIMO case) with LOS channels and 128 antennas at the receiver. The proposed methods are displayed with solid linestyle.

antenna regime in more detail in the following. We inspect the NMSE performance for the 3GPP channel model with one propagation cluster and 128 antennas at the receiver in the SIMO case over the SNR in Fig. 10. The proposed estimators outperform the baseline methods and achieve about 10 dB advantage compared to LS over the whole SNR range.

To highlight the independence of the adopted channel model, we show simulation results for the QuaDRiGa channel model (SIMO case) in Fig. 11. This time, we cannot display the genie-cov curve as the true CCM is unavailable. As can be seen in the plot, all VAE-based estimators show superior NMSE results. The advantages are not as pronounced as in the previous figures but still noticeable.

Fig. 12 shows the NMSE performance for the 3GPP channel model (MIMO case) with one propagation cluster, 32 antennas at the receiver, and 4 antennas at the transmitter. The qualitative behavior of the curves is similar to Fig. 10, where also one propagation cluster is considered. CNN and GMM show the worst NMSE among the ML-based methods for SNRs larger than -5 dB. In this case as well, VAE-noisy and VAE-real show comparable performance. Compared to LS, the VAE-based methods attain a performance gain between 6 and 13 dB. In Fig. 13, we illustrate estimation results for the 3GPP channel model (MIMO case) with three propagation clusters, 32 antennas at the receiver, and 4 antennas at the transmitter. As in the




Fig. 12. Normalized MSE for the 3GPP channel model (MIMO case) with one propagation cluster, 32 antennas at the receiver, and 4 antennas at the transmitter. The proposed methods are displayed with solid linestyle.




Fig. 13. Normalized MSE for the 3GPP channel model (MIMO case) with three propagation clusters, 32 antennas at the receiver, and 4 antennas at the transmitter. The proposed methods are displayed with solid linestyle.

previous figure, the proposed estimators clearly outperform the baselines. However, the performance gaps in Fig. 12 are more noticeable than in Fig. 13. The performance gain compared to LS shrinks to a range from 2 to 11 dB.

In summary, the VAE-based methods exhibit immense performance gains for large antenna arrays, i.e., larger equal 16 antennas and all considered numbers of propagation clusters, significantly outperforming the baseline methods. The strong performance for arrays with many antennas is likely due to the circulant approximation to the Toeplitz CCM, which becomes better for large arrays. VAE-genie lies almost on the genie-cov estimator, and the performance of VAE-noisy and VAE-real is almost identical for all considered scenarios. Moreover, all ML-based methods use genie knowledge during the training phase in the form of ground-truth channel training data, except for VAE-real, which is trained and evaluated solely based on noisy pilot observations. From this point of view, the strong estimation results of VAE-real are even more meaningful. What is more, the parameterization choices for $p_{\theta}(\mathbf{h} | \mathbf{z})$ and $q_{\phi}(\mathbf{z} | \mathbf{y})$ excellently fulfill their purpose in accordance with Theorem 1 as a result of the VAE-based approaches' strong estimation results. This becomes apparent when comparing the VAE to the global-cov or GMM results, either adopting a Gaussian or GMM prior, respectively, highlighting the superiority of the VAE-based prior.

VI. CONCLUSION

This manuscript presents a novel estimation technique based on the VAE. The idea is to tractably model the underlying data distribution as CG via a VAE, yielding a powerful generative prior. The CG modeling allows us to parameterize the MSE-optimal CME under the VAE framework. We propose three estimator variants, of which we find the VAE-real variant particularly appealing as it does not require access to ground-truth data during training or evaluation. We provide theoretical analysis that quantifies the error gap between the proposed MAP-VAE estimator and the CME and relates the training process of the VAE to the estimation performance, supporting the strong estimation capabilities of the proposed VAE-based estimators. Our extensive CE simulations highlight that the proposed methods attain excellent performance for various system configurations. In future works, we want to investigate the effects of regularization terms in the training objective and analyze other (especially wide) observation matrices. Moreover, we plan to consider nonlinear system models, e.g., quantized systems [52], to broaden the application horizon.

APPENDIX

A. Derivation of (17)

For the parameterized joint PDF of \mathbf{y} and \mathbf{h} given \mathbf{z} ,

$$p_{\theta}(\mathbf{y}, \mathbf{h} | \mathbf{z}) = p(\mathbf{y} | \mathbf{h}, \mathbf{z}) p_{\theta}(\mathbf{h} | \mathbf{z}) = p(\mathbf{y} | \mathbf{h}) p_{\theta}(\mathbf{h} | \mathbf{z}) \quad (36)$$

with $p(\mathbf{y} | \mathbf{h}) = \mathcal{N}_{\mathbb{C}}(\mathbf{A}\mathbf{h}, \Sigma)$ and $p_{\theta}(\mathbf{h} | \mathbf{z})$ as in (5). The dependency of $p(\mathbf{y} | \mathbf{h}, \mathbf{z})$ from \mathbf{z} can be dropped due to the Bayesian network in Fig. 1 showing \mathbf{y} is conditionally independent of \mathbf{z} given \mathbf{h} . Since the multiplication of two Gaussian distributions is again Gaussian, $p_{\theta}(\mathbf{y}, \mathbf{h} | \mathbf{z}) =$

$$\mathcal{N}_{\mathbb{C}} \left(\begin{bmatrix} \mathbf{A}\boldsymbol{\mu}_{\theta} \\ \boldsymbol{\mu}_{\theta} \end{bmatrix}, \begin{bmatrix} \Sigma^{-1} & -\Sigma^{-1}\mathbf{A} \\ -\mathbf{A}^H\Sigma^{-1} & \mathbf{A}^H\Sigma^{-1}\mathbf{A} + \mathbf{C}_{\theta}^{-1} \end{bmatrix}^{-1} \right) \quad (37)$$

after some algebraic reformulations, which shows that \mathbf{y} and \mathbf{h} are jointly Gaussian given \mathbf{z} . We omit the \mathbf{z} -argument here for notational brevity. For the derivation of $E_{\theta}[\mathbf{h} | \mathbf{z}, \mathbf{y}]$, we are interested in the conditional $p_{\theta}(\mathbf{h} | \mathbf{z}, \mathbf{y})$. Using standard results for jointly Gaussian distributions, the conditional is again Gaussian with the mean vector $E_{\theta}[\mathbf{h} | \mathbf{z}, \mathbf{y}] =$

$$\boldsymbol{\mu}_{\theta} + (\mathbf{A}^H\Sigma^{-1}\mathbf{A} + \mathbf{C}_{\theta}^{-1})^{-1} \mathbf{A}^H\Sigma^{-1}(\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_{\theta}) \quad (38)$$

and covariance matrix $(\mathbf{A}^H\Sigma^{-1}\mathbf{A} + \mathbf{C}_{\theta}^{-1})^{-1}$. Application of the matrix inversion lemma to (38) and further algebraic reformulations yield (17), concluding the derivation.

B. Proof of Theorem 1

Proof: Let us define the variables

$$\Psi = f_{\theta,2}(\mathbf{z}) - f_{\theta,2}(\boldsymbol{\mu}_{\phi}), \quad \psi = f_{\theta,1}(\mathbf{z}) - f_{\theta,1}(\boldsymbol{\mu}_{\phi}), \quad (39)$$

$$\Gamma = f_{\theta,2}(\boldsymbol{\mu}_{\phi}) + \varsigma^2 \mathbf{I}, \quad (40)$$

and denote $E_{\mathbf{z}}[\cdot] := E_{p_{\theta}(\mathbf{z}|\mathbf{y})}[\cdot]$ for notational convenience. First, let us rewrite $\mathbf{T}(\mathbf{z}) = \mathbf{C}_{\theta}(\mathbf{z})(\mathbf{C}_{\theta}(\mathbf{z}) + \varsigma^2 \mathbf{I})^{-1}$ as

$$\mathbf{T}(\mathbf{z}) = (\mathbf{C}_{\theta}(\mathbf{z}) + \varsigma^2 \mathbf{I})^{-1} \mathbf{C}_{\theta}(\mathbf{z}) \quad (41)$$

$$= (\Gamma + \Psi)^{-1} f_{\theta,2}(\mathbf{z}) \quad (42)$$

$$= (\Gamma^{-1} - \Gamma^{-1}\Psi(\Gamma + \Psi)^{-1}) f_{\theta,2}(\mathbf{z}) \quad (43)$$

$$= \Gamma^{-1} f_{\theta,2}(\boldsymbol{\mu}_{\phi}) + \Gamma^{-1}\Psi (\mathbf{I} - (\Gamma + \Psi)^{-1} f_{\theta,2}(\mathbf{z})) \quad (44)$$

by using the push-through identity in (41) and the matrix inversion lemma in (43). Using this result, we rewrite the CME $E[\mathbf{h} | \mathbf{y}]$ from (16) in terms of the MAP-VAE estimator $\hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y})$ from (20) and an additive error term as

$$\begin{aligned} E[\mathbf{h} | \mathbf{y}] &= E_{\mathbf{z}}[f_{\theta,1}(\mathbf{z}) + \mathbf{T}(\mathbf{z})(\mathbf{y} - f_{\theta,1}(\mathbf{z}))] \quad (45) \\ &= \underbrace{f_{\theta,1}(\boldsymbol{\mu}_{\phi}) + \Gamma^{-1} f_{\theta,2}(\boldsymbol{\mu}_{\phi})(\mathbf{y} - f_{\theta,1}(\boldsymbol{\mu}_{\phi}))}_{=\hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y})} \\ &\quad + E_{\mathbf{z}}[(\mathbf{I} - \Gamma^{-1} f_{\theta,2}(\boldsymbol{\mu}_{\phi}))\psi] \\ &\quad + E_{\mathbf{z}}[\Gamma^{-1}\Psi(\mathbf{I} - (\Gamma + \Psi)^{-1} f_{\theta,2}(\mathbf{z}))(\mathbf{y} - f_{\theta,1}(\mathbf{z}))]. \quad (46) \end{aligned}$$

We further note that we can simplify the term

$$\mathbf{I} - (\Gamma + \Psi)^{-1} f_{\theta,2}(\mathbf{z})(\mathbf{y} - f_{\theta,1}(\mathbf{z})) = \mathbf{y} - E[\mathbf{h} | \mathbf{y}, \mathbf{z}] \quad (47)$$

$$= E[\mathbf{n} | \mathbf{y}, \mathbf{z}]. \quad (48)$$

Thus, we get an upper bound on the expected Euclidean distance between the MAP-VAE estimator and the CME $\varepsilon = E[\|E[\mathbf{h} | \mathbf{y}] - \hat{\mathbf{h}}_{\text{VAE}}(\mathbf{y})\|_2]$ as

$$\begin{aligned} \varepsilon &\leq E[\|\mathbf{I} - \Gamma^{-1} f_{\theta,2}(\boldsymbol{\mu}_{\phi})\|_2 E_{\mathbf{z}}[\|\psi\|_2]] \\ &\quad + E[\|\Gamma^{-1}\|_2 \cdot \|E_{\mathbf{z}}[\Psi E[\mathbf{n} | \mathbf{y}, \mathbf{z}]]\|_2] \quad (49) \\ &\leq E\left[\frac{\varsigma^2}{\xi_{\min} + \varsigma^2} E_{\mathbf{z}}[\|\psi\|_2]\right] + \frac{1}{\varsigma^2} E[E_{\mathbf{z}}[\|\Psi E[\mathbf{n} | \mathbf{y}, \mathbf{z}]\|_2]] \quad (50) \end{aligned}$$

where we used the reformulation from (48) and the bounds on the spectral norms

$$\|\mathbf{I} - \Gamma^{-1} f_{\theta,2}(\boldsymbol{\mu}_{\phi})\|_2 = \frac{\varsigma^2}{\xi_{\min} + \varsigma^2}, \quad \|\Gamma^{-1}\|_2 \leq \frac{1}{\varsigma^2}, \quad (51)$$

together with the triangle and Cauchy-Schwarz inequalities. Note that ξ_{\min} is a function of \mathbf{y} and the outer expectations are with respect to $p(\mathbf{y})$ if not denoted otherwise. Thus, we employ Hölder's inequality for both summands in (50) to get

$$\begin{aligned} \varepsilon &\leq C_1 \sqrt{E[E_{\mathbf{z}}[\|\psi\|_2^2]]} \\ &\quad + \frac{1}{\varsigma^2} \sqrt{E[E_{\mathbf{z}}[\|\Psi E[\mathbf{n} | \mathbf{y}, \mathbf{z}]\|_2^2]]} \sqrt{E[E_{\mathbf{z}}[\|\Psi\|_2^2]]} \quad (52) \end{aligned}$$

$$\leq C_1 \sqrt{E[E_{\mathbf{z}}[\|\psi\|_2^2]]} + \frac{1}{\varsigma^2} \sqrt{E[\|\mathbf{n}\|_2^2]} \sqrt{E[E_{\mathbf{z}}[\|\Psi\|_2^2]]} \quad (53)$$

$$= C_1 \sqrt{E[E_{\mathbf{z}}[\|\psi\|_2^2]]} + C_2 \sqrt{E[E_{\mathbf{z}}[\|\Psi\|_2^2]]} \quad (54)$$

where we used Jensen's inequality in combination with the law of total expectation in (53), and with C_1, C_2 in (26). By resubstituting the variables in (39), we get

$$\varepsilon \leq C_1 \sqrt{E[E_{\mathbf{z}}[\|f_{\theta,1}(\mathbf{z}) - f_{\theta,1}(\boldsymbol{\mu}_{\phi})\|_2^2]]}$$

$$+ C_2 \sqrt{\mathbb{E}[\mathbb{E}_z[\|f_{\theta,2}(z) - f_{\theta,2}(\mu_\phi)\|_2^2]]} \quad (55)$$

$$\leq (C_1 L_1 + C_2 L_2) \sqrt{\mathbb{E}[\mathbb{E}_z[\|z - \mu_\phi\|_2^2]]} \quad (56)$$

by using the Lipschitz continuity (24). By defining the first and second moments of the posterior distribution $p_\theta(z | \mathbf{y})$ as $\bar{\mu} = \mathbb{E}_z[z]$ and $\bar{C} = \mathbb{E}_z[(z - \bar{\mu})(z - \bar{\mu})^H]$, we write

$$\sqrt{\mathbb{E}[\mathbb{E}_z[\|z - \mu_\phi\|_2^2]]} = \sqrt{\mathbb{E}[\mathbb{E}_z[\|z - \bar{\mu} + \bar{\mu} - \mu_\phi\|_2^2]]} \quad (57)$$

$$\leq \sqrt{\mathbb{E}[\mathbb{E}_z[(\|z - \bar{\mu}\|_2 + \|\bar{\mu} - \mu_\phi\|_2)^2]]} \quad (58)$$

$$\leq \sqrt{\text{tr}(\bar{C})} + \sqrt{\mathbb{E}[\|\bar{\mu} - \mu_\phi\|_2^2]} \quad (59)$$

since only one summand in (58) depends on \mathbf{y} or \mathbf{z} , respectively. Plugging (59) in (56) yields (25), completing the proof. \square

REFERENCES

- [1] L. Ruthotto and E. Haber, "An introduction to deep generative modeling," *GAMM-Mitteilungen*, vol. 44, no. 2, 2021, Art. no. e202100008.
- [2] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, U.K.: Springer, 2006.
- [3] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014.
- [4] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. 2nd Int. Conf. Learn. Representation*, 2014.
- [5] I. Goodfellow et al., "Generative adversarial networks," *Commun. ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [6] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019.
- [7] G. Ongie et al., "Deep learning techniques for inverse problems in imaging," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 1, pp. 39–56, May 2020.
- [8] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed sensing using generative models," in *Proc. 34th Int. Conf. Mach. Learn.*, vol. 70, 2017, pp. 537–546.
- [9] A. Jalal et al., "Robust compressed sensing MRI with deep generative priors," in *Proc. Adv. Neural Inf. Process. Syst.*, 2021.
- [10] P. Hand, O. Leong, and V. Voroninski, "Phase retrieval under a generative prior," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018.
- [11] M. Asim, F. Shamshad, and A. Ahmed, "Blind image deconvolution using deep generative priors," *IEEE Trans. Comput. Imag.*, vol. 6, pp. 1493–1506, 2020.
- [12] M. Koller, B. Fesl, N. Turan, and W. Utschick, "An asymptotically MSE-optimal estimator based on Gaussian mixture models," *IEEE Trans. Signal Process.*, vol. 70, pp. 4109–4123, 2022.
- [13] E. Balevi and J. G. Andrews, "Wideband channel estimation with a generative adversarial network," *IEEE Trans. Wireless Commun.*, vol. 20, no. 5, pp. 3049–3060, May 2021.
- [14] E. Balevi, A. Doshi, A. Jalal, A. Dimakis, and J. G. Andrews, "High dimensional channel estimation using deep generative networks," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 1, pp. 18–30, Jan. 2021.
- [15] A. S. Doshi, M. Gupta, and J. G. Andrews, "Over-the-air design of GAN training for mmWave MIMO channel estimation," *IEEE J. Sel. Areas Inf. Theory*, vol. 3, no. 3, pp. 557–573, Sep. 2022.
- [16] M. Arvinte and J. I. Tamir, "MIMO channel estimation using score-based generative models," *IEEE Trans. Wireless Commun.*, vol. 22, no. 6, pp. 3698–3713, Jun. 2023.
- [17] T. Diskin, Y. C. Eldar, and A. Wiesel, "Learning to estimate without bias," *IEEE Trans. Signal Process.*, vol. 71, pp. 2162–2171, 2023.
- [18] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [19] S. Zhao, J. Song, and S. Ermon, "Towards a deeper understanding of variational autoencoding models," 2017, *arXiv:1702.08658*.
- [20] M. Baur, B. Fesl, M. Koller, and W. Utschick, "Variational autoencoder leveraged MMSE channel estimation," in *Proc. 56th Asilomar Conf. Signals, Syst., Comput.*, 2022, pp. 527–532.
- [21] M. Rani, S. B. Dhok, and R. B. Deshmukh, "A systematic review of compressive sensing: Concepts, implementations and applications," *IEEE Access*, vol. 6, pp. 4875–4894, 2018.
- [22] A. Banerjee, X. Guo, and H. Wang, "On the optimality of conditional expectation as a Bregman predictor," *IEEE Trans. Inf. Theory*, vol. 51, no. 7, pp. 2664–2669, Jul. 2005.
- [23] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *Found. Trends Mach. Learn.*, vol. 12, no. 4, pp. 307–392, 2019.
- [24] D. P. Bertsekas and J. N. Tsitsiklis, *Introduction to Probability*, 2nd ed. Nashua, NH, USA: Athena Scientific, 2008.
- [25] J. Yang et al., "Compressive sensing by learning a Gaussian mixture model from measurements," *IEEE Trans. Image Process.*, vol. 24, no. 1, pp. 106–119, Jan. 2015.
- [26] M. Loève, *Probability Theory I*, 4th ed. New York, NY, USA: Springer, 1977.
- [27] D. Fuhrmann, "Application of Toeplitz covariance estimation to adaptive beamforming and detection," *IEEE Trans. Signal Process.*, vol. 39, no. 10, pp. 2194–2198, Oct. 1991.
- [28] Y. Ephraim, D. Malah, and B.-H. Juang, "On the application of hidden Markov models for enhancing noisy speech," *IEEE Trans. Acoust.*, vol. 37, no. 12, pp. 1846–1856, Dec. 1989.
- [29] D. Neumann, T. Wiese, and W. Utschick, "Learning the MMSE channel estimator," *IEEE Trans. Signal Process.*, vol. 66, no. 11, pp. 2905–2917, Jun. 2018.
- [30] A. J. Heckens, S. M. Krause, and T. Guhr, "Uncovering the dynamics of correlation structures relative to the collective market motion," *J. Statist. Mech. Theory Exp.*, vol. 2020, no. 10, 2020, Art. no. 103402.
- [31] B. Fesl et al., "Channel estimation based on Gaussian mixture models with structured covariances," in *Proc. 56th Asilomar Conf. Signals, Syst. Comput.*, 2022, pp. 533–537.
- [32] M. Baur, N. Turan, B. Fesl, and W. Utschick, "Channel estimation in underdetermined systems utilizing variational autoencoders," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2024, pp. 9031–9035.
- [33] M. Baur, B. Böck, N. Turan, and W. Utschick, "Variational autoencoder for channel estimation: Real-world measurement insights," in *Proc. 27th Int. Work. Smart Antennas*, Piscataway, NJ, USA: IEEE Press, 2024, pp. 117–122.
- [34] R. M. Gray, "Toeplitz and circulant matrices: A review," *Found. Trends Commun. Inf. Theory*, no. 3, pp. 155–239, 2006.
- [35] B. Böck, M. Baur, N. Turan, D. Semmler, and W. Utschick, "A statistical characterization of wireless channels conditioned on side information," 2024, *arXiv:2406.04282*.
- [36] X. Yin and C. Xiang, *Propagation Channel Characterization, Parameter Estimation, and Modeling for Wireless Communications*. Hoboken, NJ, USA: Wiley, 2016.
- [37] J. Keramoal, L. Schumacher, K. Pedersen, P. Mogensen, and F. Frederiksen, "A stochastic MIMO radio channel model with experimental validation," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 6, pp. 1211–1226, Aug. 2002.
- [38] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [39] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *J. Mach. Learn. Res.*, vol. 13, no. 10, pp. 281–305, 2012.
- [40] R. Liaw et al., "TUNE: A research platform for distributed model selection and training," 2018, *arXiv:1807.05118*.
- [41] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [42] T. Salimans and D. P. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2016.
- [43] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," in *Adv. Neural Inf. Process. Syst. Deep Learn. Symp.*, 2016.
- [44] S. Anwar, K. Hwang, and W. Sung, "Structured pruning of deep convolutional neural networks," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 13, no. 3, pp. 1–18, 2017.
- [45] "Spatial channel model for multiple input multiple output (MIMO) simulations (Release 16)," 3rd Generation Partnership Project (3GPP), Tech. Rep. 25.996 V16.0.0, 2020.
- [46] S. Jaeckel, L. Raschkowski, K. Borner, and L. Thiele, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE Trans. Antennas Propag.*, vol. 62, no. 6, pp. 3242–3256, Jun. 2014.
- [47] S. Jaeckel et al., "QuaDRiGa - quasi deterministic radio channel generator, user manual and documentation," Fraunhofer Heinrich Hertz Institute, Berlin, Germany, Tech. Rep. v2.6.1, 2021.

- [48] S. A. Busari, K. M. S. Huq, S. Mumtaz, L. Dai, and J. Rodriguez, "Millimeter-Wave Massive MIMO communication for future wireless systems: A survey," *IEEE Commun. Surv. Tut.*, vol. 20, no. 2, pp. 836–869, 2nd Quart. 2018.
- [49] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. Motivation and construction," in *Proc. IEEE Inf. Theory Work. Inf. Theory*, Piscataway, NJ, USA: IEEE Press, 2010, pp. 1–5.
- [50] A. Maleki, L. Anitori, Z. Yang, and R. G. Baraniuk, "Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP)," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4290–4308, Jul. 2013.
- [51] B. Fesl, N. Turan, M. Koller, and W. Utschick, "A low-complexity MIMO channel estimator with implicit structure of a convolutional neural network," in *Proc. 22nd Int. Work. Signal Process. Adv. Wirel. Commun.*, 2021, pp. 11–15.
- [52] B. Fesl, N. Turan, B. Bock, and W. Utschick, "Channel estimation for quantized systems based on conditionally gaussian latent models," *IEEE Trans. Signal Process.*, vol. 72, pp. 1475–1490, 2024.



Michael Baur (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. (with high distinction) degrees in electrical engineering and information technology from the Technische Universität München (TUM), München, Germany, in 2017 and 2020, respectively. In 2020, he joined the Chair of Signal Processing with the TUM, where he is currently a Research Assistant and a Ph.D. candidate. In 2019, he was a Visiting Scholar with the Department of Electronic Engineering, Tsinghua University, Beijing, China. His research interests

include generative models, statistical signal processing, and estimation theory with a focus on applications in wireless communications.



Benedikt Fesl (Graduate Student Member, IEEE) received the B.Sc. (with distinction) and M.Sc. (with high distinction) degrees in electrical engineering and information technology from the Technische Universität München (TUM), München, Germany, in 2018 and 2020, respectively. In 2020, he joined the Chair of Signal Processing with the TUM, where he is currently a Research Assistant and a Ph.D. candidate. In 2019, he was a Visiting Scholar with the Department of Electrical Engineering and Computer Science, University of California Irvine, Irvine, CA,

USA. His research interests include statistical signal processing, estimation theory, machine learning, generative modeling, and signal processing in quantized systems.



Wolfgang Utschick (Fellow, IEEE) received the Diploma and Ph.D. degrees in electrical engineering (with distinction) from the Technical University of Munich (TUM), Germany, in 1993 and 1998, respectively. Before that, he has completed several years of certified industrial training programs. Since 2002, he has been an Extraordinarius Associate Professor and then an Ordinarius Full Professor in methods of signal processing with TUM. Since 2011, he regularly teaches courses in Singapore as a TUM Asia Faculty Member and as a regular Guest

Professor with Singapore Institute of Technology. In 2021, he became a Core Member of the newly founded Munich Data Science Institute. From 2017 to 2022, he worked two terms of office as the elected Dean of the Department of Electrical and Computer Engineering, TUM. He served on the Board of Deans for TUM's newly founded School of Communication, Information, and Technology. He is a Member of the VDE (German Commission for Electrical, Electronic, and Information Technologies), where he currently serves as the Deputy Head of the Specialist Group KT 1 for Information and System Theory of the ITG (German Information Technology Society). Since 2016, he has been the Director of the Special Research Group Mobility & Transport of the Bavarian Science Forum BayWISS. In 2024, he was elected as a DFG (German Research Foundation) Review Board Member for Electrical Engineering and Information Technology. He teaches courses on signal processing, stochastic processes, optimization theory, and machine learning in wireless communications and other application areas. He has authored or co-authored many technical articles in international journals and conference proceedings. He has edited several books and is the Founder and an Editor of the book series titled *Foundations in Signal Processing, Communications and Networking* (Springer). He holds several patents in multiantenna signal processing. He has been a Principal Investigator in several research projects funded by the German Research Foundation (DFG) and coordinated a German DFG Priority Program on Communications Over Interference Limited Networks. He is a Principal Investigator of the 6G Future Laboratory of the Free State of Bavaria and the 6G-life Consortium of the German Federal Ministry of Education and Research. From 2016 to 2022, he chaired the German Chapter of the IEEE Signal Processing Society. He was a Member of the IEEE Signal Processing Society Technical Committee on Signal Processing for Communications and Networking, and also served as an Associate Editor for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS and IEEE TRANSACTIONS ON SIGNAL PROCESSING.