Distributed Variational Representation Learning

Iñaki Estella Aguerri   Abdellatif Zaidi

Abstract

The problem of distributed representation learning is one in which multiple sources of information $X_1, \ldots, X_K$ are processed separately so as to learn as much information as possible about some ground truth Y. We investigate this problem from information-theoretic grounds, through a generalization of Tishby's centralized Information Bottleneck (IB) method to the distributed setting. Specifically, K encoders, $K \geq 2$, compress their observations $X_1, \ldots, X_K$ separately in a manner such that, collectively, the produced representations preserve as much information as possible about Y. We study both discrete memoryless (DM) and memoryless vector Gaussian data models. For the discrete model, we establish a single-letter characterization of the optimal tradeoff between complexity (or rate) and relevance (or information) for a class of memoryless sources (the observations $X_1, \ldots, X_K$ being conditionally independent given Y). For the vector Gaussian model, we provide an explicit characterization of the optimal complexity-relevance tradeoff. Furthermore, we develop a variational bound on the complexity-relevance tradeoff which generalizes the evidence lower bound (ELBO) to the distributed setting. We also provide two algorithms that allow to compute this bound: i) a Blahut-Arimoto type iterative algorithm which enables to compute optimal complexity-relevance encoding mappings by iterating over a set of self-consistent equations, and ii) a variational inference type algorithm in which the encoding mappings are parametrized by neural networks and the bound approximated by Markov sampling and optimized with stochastic gradient descent. Numerical results on synthetic and real datasets are provided to support the efficiency of the approaches and algorithms developed in this paper.

I. INTRODUCTION

The problem of extracting a *good* representation of data, i.e., one that makes it easier to extract useful information, is at the heart of the design of efficient machine learning algorithms. One important question, which is often controversial in statistical learning theory, is the choice of a "good" loss function that measures discrepancies between the true values and their estimated fits. There is however numerical evidence that models that are trained to maximize mutual information, or equivalently minimize the error's entropy, often outperform ones that are trained using other criteria, such as mean-square error (MSE), higher-order statistics [2], [3]. Although a complete and rigorous justification of the usage of mutual information as cost function in learning is still awaited, recently a partial explanation appeared in [4] where the authors show that, under some natural data processing property, Shannon's mutual information uniquely quantifies the reduction of prediction risk due to side information. Along the same line of work, Painsky and Wornell [5] show that, for binary classification problems, by minimizing the logarithmic-loss (log-loss) one actually minimizes an upper bound to any choice of loss function that is smooth, proper (i.e., unbiased and Fisher consistent) and convex. Perhaps, this justifies partially why mutual information

(or, equivalently, the corresponding loss function which is the log-loss fidelity measure) is widely used in learning theory and has already been adopted in many algorithms in practice such as the *infomax* criterion [6], the tree-based algorithm of [7] or the well known Chow-Liu algorithm [8] for learning tree graphical models, with applications in genetics [9], image processing [10], computer vision [11] and others. The log-loss measure also plays a central role in the theory of prediction [12, Ch. 09], where it is often referred to as the *self-information* loss function, as well as in Bayesian modeling [13] where priors are usually designed so as to maximize the mutual information between the parameter to be estimated and the observations.

The goal of learning, however, is not merely to learn model parameters accurately for previously seen data. Rather, in essence, it is the ability to successfully apply rules that are extracted from previously seen data to characterize new unseen data. This is often captured through the notion of "generalization error". The generalization capability of a learning algorithm hinges on how sensitive is the output of the algorithm to modifications of the input dataset, i.e., its *stability* [14], [15]. In the context of deep learning, it can be seen as a measure of how much the algorithm overfits the model parameters to the seen data. In fact, efficient algorithms should strike a good balance among their ability to fit training dataset and that to generalize well to unseen data. In statistical learning theory [12], such a dilemma is reflected through that the minimization of the "population risk" (or "test error" in the deep learning literature) amounts to the minimization of the sum of the two terms that are generally difficult to minimize simultaneously, the "empirical risk" on the training data and the generalization error. In order to prevent over-fitting, regularization methods can be employed, which include parameter penalization, noise injection, and averaging over multiple models trained with distinct sample sets. Although it is not yet very well understood how to optimally control model complexity, recent works [16], [17] show that the generalization error can be upper-bounded using the mutual information between the input dataset and the output of the algorithm. This result actually formalizes the intuition that the less information a learning algorithm extracts from the input dataset the less it is likely to overfit; and justifies the use of mutual information also as a regularizer term.

A growing body of works focuses on developing learning rules and algorithms using information theoretic approaches, e.g., see [18]–[23] and references therein. Most relevant to this paper is the Information Bottleneck (IB) method of [18] which readily and elegantly captures the above mentioned viewpoint of seeking the right balance between data fit and generalization by using the mutual information both as a cost function and as a regularizer term. Specifically, IB formulates the problem of extracting the relevant information that some signal $X \in \mathcal{X}$ provides about another one $Y \in \mathcal{Y}$ that is of interest as that of finding a representation U that is maximally informative about Y (i.e., large mutual information $I(U;Y)$) while being minimally informative about X (i.e., small mutual information $I(U;X)$). In the IB framework, $I(U;Y)$ is referred to as the *relevance* of U and $I(U;X)$ is referred to as the *complexity* of U, where complexity here is measured by the minimum description length (or rate) at which the observation is compressed. Accordingly, the performance of learning with the IB method and the optimal mapping of the data X to U are found by solving the Lagrangian formulation

$$\mathcal{L}^*_{\mathrm{IB},s} := \max_{P_{U|X}} I(U;Y) - sI(U;X), \tag{1}$$

where $P_{U|X}$ is a stochastic map that assigns the observation X to a latent variable U from which Y is inferred, and s is the Lagrange multiplier. There is an implicit tradeoff between complexity and relevance: when the description length (complexity) is low, the largest relevant information captured by U from Y is restricted, and the other way around. The simultaneously achievable relevance-complexity pairs $(I(U;Y), I(U;X))$ define a region characterized in [18] and the resulting optimal mappings $P_{U|X}$ are made to perform at different points of the relevance-complexity region by considering different values of s.

Direct optimization of (1) to obtain the optimal mappings $P_{U|X}$ is generally challenging. Instead, a tight variational bound can be optimized, which results in the optimization of a generalized version of the (2) has the evidence lower bound (ELBO) [24] (and the β-VAE bound [25]) , used, e.g., for variational inference:

$$\max_{P_{U|X}, Q_{Y|U}, Q_U} \mathbb{E}[\log Q_{Y|U}] - s D_{\mathrm{KL}}(P_{X|U} \| Q_U), \tag{2}$$

where $Q_{Y|U}$ and Q_U are variational approximations of the optimal decoder and the latent space, and $D_{\mathrm{KL}}(\cdot \| \cdot)$ is the Kullback-Leiber divergence.

Different optimization methods have been proposed to optimize (2). An iterative Blahut-Arimoto (BA) type algorithm, that converges to a stationary solution of (1) is proposed in [18], which results in an algorithm similar to the expectation maximization (EM) algorithm [26] for problems with discrete [27], [28] and Gaussian data [29], [30]. However, this algorithm generally requires a good estimation of the data distribution or perfect knowledge of it, and becomes too complex for general high-dimensional data. An alternative strategy, which only requires data samples, has been proposed in [31]–[33] to overcome this limitation, based on approximating the bound (2) by parameterizing the encoder, decoder and prior distributions with deep neural networks (DNNs), using Monte-Carlo sampling and optimizing the DNN to maximize it. This approach is essentially that used for variational inference [34] and variational autoencoders (VAE) [24].

The IB approach, has found remarkable applications in supervised and unsupervised learning problems such as classification, clustering and prediction. Perhaps key to the analysis, and theoretical development, of the IB method is its elegant connection with information-theoretic rate-distortion problems, for it is now well known that the IB problem is essentially a remote source coding problem [35], [36] in which the distortion is measured under logarithmic loss. Recent works show that this connection turns out to be useful for a better understanding of deep neural networks (DNN) [37], the emergence of invariance and disentanglement in DNN [38], the minimization of PAC-bayesian bounds on the test error [38], [39]. Other connections, that are more intriguing, exist also with seemingly unrelated problems such as privacy and hypothesis testing [40]–[42] or multiterminal networks with oblivious relays [43]. We close this section by mentioning that the abstract viewpoint of IB seems also instrumental towards a better understanding of the so-called *representation learning* [44], which is an active research area in machine learning that focuses on identifying and disentangling the underlying explanatory factors that are hidden in the observed data in an attempt to render learning algorithms less dependent on feature engineering. We also point out that there exists an extensive literature on building optimal estimators of information quantities (e.g. entropy, mutual information), as well as their Matlab/Python implementations, including in the high-dimensional regime,
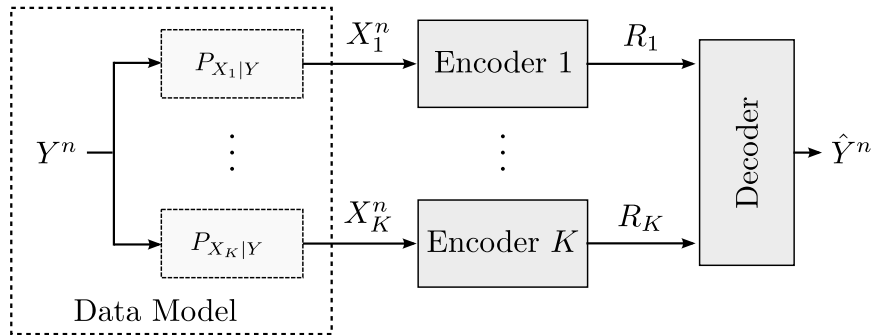
Fig. 1: A model for distributed learning, e.g., multi-view learning.

e.g., [31]–[33], [45]–[48] and references therein.

In this paper, we study the distributed learning problem as depicted in Figure 1. Here, Y is the signal to be predicted and $(X_1, \ldots, X_K)$ are K views of it that could each be relevant to understand one or more aspects of Y. The observations (or views) could be distinct or redundant. We make the assumption that $(X_1, \ldots, X_K)$ are independent given Y. This assumption holds in many practical scenarios. For example, the reader may think of $(X_1, \ldots, X_K)$ as being independent noisy versions of Y. The model generalizes the aforementioned centralized learning setting to the distributed setting. By opposition to the centralized case, however, encoders should strike a suitable tradeoff between data fit and generalization *collectively*, not individually.

A. Example: Multi-view Learning

In many data analytics problems, data is collected from various sources of information or feature extractors; and is intrinsically *heterogenous*. For example, an image can be identified by its color or texture features; and a document may contain text and images. Conventional machine learning approaches concatenate all available data into one big row vector (or matrix) on which a suitable algorithm is then applied. Treating different observations as a single source might cause overfitting and is not physically meaningful because each group of data may have different statistical properties. Alternatively, one may partition the data into groups according to samples homogeneity, and each group of data be regarded as a separate *view*. This paradigm, termed *multi-view learning* [49], [50], has received growing interest; and various algorithms exist, sometimes under references such as *co-training* [51]–[54], *multiple kernel learning* [54] and *subspace learning* [55]. By using distinct encoder mappings to represent distinct groups of data, and jointly optimizing over all mappings to remove redundancy, multiview learning offers a degree of flexibility that is not only desirable in practice but is likely to result in better learning capability. Actually, as Vapnik shows in [56], local learning algorithms produce less errors than global ones. Viewing the problem as that of function approximation, the intuition is that it is usually non-easy to find a unique function that holds good predictability properties in the entire data space.

Besides, the distributed learning of Figure 1 clearly finds application in all those scenarios in which learning is performed collaboratively but distinct learners either only access subsets of the data (e.g., due to physical constraints) or they access independent noisy versions of the entire dataset. Two such examples are Google Goggles and Siri in which the locally collected data is processed on clouds.

B. Main Contributions

We extend Tishby's centralized IB method to the distributed learning setting shown in Figure 1. We study the performance of this model in the *information plane*; and we characterize optimal complexity-relevance tradeoffs for a class of discrete memoryless sources as well as for Vector Gaussian sources. Furthermore, we develop a variational bound on the optimal information-rate that can be seen as a generalization of IB method, the ELBO and the β-VAE criteria [25], [57], [58] to the distributed setting. We also provide two algorithms which are trained by optimizing this bound. Specifically, the main contributions of this paper are:

- In Section III-A, we establish a single-letter characterization of optimal tradeoffs between complexity and relevance for the distributed learning model of Figure 1 for a class of discrete sources for which the observations are independent conditionally on the target source. In doing so, we exploit the connection with the distributed Chief Executive Officer (CEO) source coding problem under logarithmic-loss distortion measure studied in [59].

- In Section III-B, we consider the problem of maximizing relevance under a constraint on the sum complexity. We derive a variational bound which generalizes the IB cost and the ELBO to the distributed setting and involves a novel regularization term to account for the joint complexity of the model.

- In Section III-C, we study a memoryless vector Gaussian data model. For this model, we find an explicit analytic characterization of optimal tradeoffs between complexity and relevance. The result generalizes the Gaussian information bottleneck projections [29], [30] to the distributed learning scenario.

- In Section IV we develop two algorithms that allow to optimize the derived variational bound: : 1) a Blahut-Arimoto [60] type iterative algorithm which enables to compute optimal relevance-complexity encoding mappings by iterating over a set of self-consistent equations and that is most useful when distributions are known or can be estimated with high accuracy, and 2) a variational inference type algorithm in which the the encoders, the decoder, and the prior distributions are parametrized by DNNs and the bound approximated by Monte-Carlo sampling and optimized with stochastic gradient descent. This algorithm makes usage of Kingma et al.'s reparametrization trick [24] and can be seen as a generalization of the variational information bottleneck algorithm in [32] to the distributed case. We also particularize the algorithms for the memoryless vector Gaussian data model.

- We provide numerical results on synthetic and real datasets that support the efficiency of the approaches and algorithms that are developed in this paper.

C. Notation

Throughout, upper case letters denote random variables, e.g., X; lower case letters denote realizations of random variables, e.g., x; and calligraphic letters denote sets, e.g., $\mathcal{X}$. The cardinality of a set is denoted by $|\mathcal{X}|$. For a random variable X with probability mass function (pmf) P_X, we use $P_X(x) = p(x)$, $x \in \mathcal{X}$ for short. Boldface upper case letters denote vectors or matrices, e.g., $\mathbf{X}$, where context should make the distinction clear. We abbreviate a sequence $(X_1, X_2, \ldots, X_n)$ of n random variables by X_n, and we denote the interval $(X_k, X_{k+1}, \cdots, X_j)$ by X_k^j. If the lower index is equal to 1, it will be omitted when there is no ambiguity (e.g., $X^j \triangleq X_1^j$). For random variables $(X_1, X_2, \ldots)$ and a set of integers $\mathcal{K} \subseteq \mathbb{N}$, $X_\mathcal{K}$ denotes the set of variables with indices in $\mathcal{K}$, i.e., $X_\mathcal{K} = \{X_k : k \in \mathcal{K}\}$. If

$\mathcal{K} = \emptyset$, $X_{\mathcal{K}} = \emptyset$. For $k \in \mathcal{K}$ we let $X_{\mathcal{K}/k} = (X_1, \ldots, X_{k-1}, X_{k+1}, \ldots, X_K)$, with $X_0 = X_{K+1} = \emptyset$. Also, for zero-mean random vectors $\mathbf{X}$ and $\mathbf{Y}$, the quantities $\mathbf{\Sigma_x}$, $\mathbf{\Sigma_{x,y}}$ and $\mathbf{\Sigma_{x|y}}$ denote respectively the covariance matrix of the vector $\mathbf{X}$, the covariance matric of vector $(\mathbf{X}, \mathbf{Y})$ and the conditional covariance matrix of $\mathbf{X}$, conditionally on $\mathbf{Y}$, i.e., $\mathbf{\Sigma_x} = \mathrm{E}[\mathbf{X}\mathbf{X}^H]$ $\mathbf{\Sigma_{x,y}} := \mathrm{E}[\mathbf{X}\mathbf{Y}^H]$, and $\mathbf{\Sigma_{x|y}} = \mathbf{\Sigma_x} - \mathbf{\Sigma_{x,y}}\mathbf{\Sigma_y^{-1}}\mathbf{\Sigma_{y,x}}$. Finally, for two probability measures P_X and Q_X on $X \in \mathcal{X}$, the relative entropy or Kullback-Leibler divergence is denoted as $D_{\mathrm{KL}}(P_X \| Q_X)$. That is, if P_X is absolutely continuous with respect to Q_X, $P_X \ll Q_X$, (i.e., for every $x \in \mathcal{X}$ if $P_X(x) > 0$ then $Q_X(x) > 0$), $D_{\mathrm{KL}}(P_X \| Q_X) = \mathbb{E}_{P_X}[\log(P_X(X)/Q_X(X))]$, otherwise $D_{\mathrm{KL}}(P_X \| Q_X) = \infty$.

II. Problem Formulation

Consider the distributed learning model shown in Figure 1. Let an integer $K \geq 2$ be given, and denote $\mathcal{K} = \{1, \ldots, K\}$. Also, let $(X_1, \ldots, X_K, Y)$ be a tuple of random variables which have joint probability mass function $P_{X_{\mathcal{K}}, Y}(x_{\mathcal{K}}, y) := P_{X_1, \ldots, X_K, Y}(x_1, \ldots, x_K, y)$ for $(x_1, \ldots, x_K) \in \mathcal{X}_1 \times \ldots \times \mathcal{X}_K$ and $y \in \mathcal{Y}$, where $\mathcal{X}_k$ for $k \in \mathcal{K}$ designates the alphabet of X_k and $\mathcal{Y}$ designates the alphabet of Y, all assumed to be finite[1]. Throughout, it will be assumed that the following Markov chain holds for all $k \in \mathcal{K}$,

$$X_k \; \multimap \; Y \; \multimap \; X_{\mathcal{K}/k}, \tag{3}$$

i.e., $p(x_{\mathcal{K}}, y) = p(y) \prod_{k=1}^{K} p(x_k | y)$ for $x_k \in \mathcal{X}_K$, $y \in \mathcal{Y}$.

The problem of distributed learning at hand seeks to characterize how accurate one can estimate the target variable Y from the observations $(X_1, \ldots, X_K)$ when those are processed separately, each by a distinct encoder. More specifically, let a training dataset $\{(X_{1,i}, \ldots, X_{K,i}, Y_i)\}_{i=1}^{n}$ consisting of n independent, identically distributed (i.i.d.) random samples drawn from the joint distribution $P_{X_{\mathcal{K}}, Y}$ be given. Encoder $k \in \mathcal{K}$ only observes the sequence X_k^n; and processes it to generate a description $J_k = \phi_k(X_k^n)$ according to some (possibly stochastic) mapping

$$\phi_k : \mathcal{X}_k^n \rightarrow \mathcal{M}_k^{(n)}, \tag{4}$$

where $\mathcal{M}_k^{(n)}$ is an arbitrary set of descriptions. The range of allowable description sets will be specified below. A (possibly stochastic) decoder $\psi(\cdot)$ collects all descriptions $J_{\mathcal{K}} = (J_1, \ldots, J_K)$ and returns an estimate $\hat{Y}^n$ of Y^n as

$$\psi : \mathcal{M}_1^{(n)} \times \ldots \times \mathcal{M}_K^{(n)} \rightarrow \hat{\mathcal{Y}}^n. \tag{5}$$

The accuracy of the estimation $\hat{Y}^n$ is measured in terms of the *relevance*, defined as the information that the descriptions $\phi_1(X_1^n), \ldots, \phi_K(X_K^n)$ *collectively* preserve about Y^n, as measured by Shannon mutual information

$$\Delta^{(n)}(P_{X_{\mathcal{K}}, Y}) := \frac{1}{n} I_{P_{X_{\mathcal{K}}, Y}}(Y^n; \hat{Y}^n), \tag{6}$$

where $\hat{Y}^n = \psi(\phi_1(X_1^n), \ldots, \phi_K(X_K^n))$ and the subscript $P_{X_{\mathcal{K}}, Y}$ indicates that the mutual information is computed under the joint distribution $P_{X_{\mathcal{K}}, Y}$.

---

[1]For simplicity, we assume finite alphabets. The results presented in this paper also extend to random variables with continuous alphabets using standard tools [61]. Results for the vector Gaussian model, which has continuous alphabet are also provided.

If the encoding mappings $\{\phi_k\}_{k=1}^K$ are un-constrained, the maximization of the right hand sight (RHS) of (6) based on the given training dataset usually results in overfitting. A better generalization capability is generally obtained by constraining the *complexity* of the encoders. In what follows, we do so by requiring that the range of the encoding functions at all encoders, i.e., the cardinality of the set of descriptions $\mathcal{M}_k^{(n)}$, are restricted in size. This is the so-called *minimum description length* complexity measure, often used in the learning literature to limit the description lenght of the weights of neural networks [62]. A connection between the use of the minimum description complexity for limiting the description length of the input encoding as described and that of the weights is given in [38]. Specifically, the encoding function $\phi_k(\cdot)$ at Encoder $k \in \mathcal{K}$ needs to satisfy

$$R_k \geq \frac{1}{n} \log |\phi_k(X_k^n)| \quad \text{for all} \quad X_k^n \in \mathcal{X}_k^n. \tag{7}$$

The characterization of the optimal performance for the distributed learning problem that we study in this paper can be cast as that of finding the region of all simultaneously achievable relevance-complexity tuples.

Definition 1. *A tuple $(\Delta, R_1, \ldots, R_K)$ is said to be achievable if there exists a training set length n, encoders ϕ_k for $k = 1, \ldots, K$ and a decoder ψ, such that*

$$\Delta \leq \frac{1}{n} I_{P_{X_{\mathcal{K}},Y}}\left(Y^n; \psi(\phi_1(X_1^n), \ldots, \phi_K(X_K^n))\right) \tag{8}$$

$$R_k \geq \frac{1}{n} \log |\phi_k(X_k^n)| \quad \text{for all} \quad k \in \mathcal{K}. \tag{9}$$

The relevance-complexity region $\mathcal{RI}_{\text{DIB}}$ is given by the closure of all achievable tuples $(\Delta, R_1, \ldots, R_K)$. ☐

Equivalently, one may seek to characterize the mappings $\{\phi_k\}_{k=1}^K$ and ψ that maximize the relevance, as given in the RHS of (6), under the complexity constraint (7). The function

$$\Delta(R_{\mathcal{K}}, P_{X_{\mathcal{K}},Y}) = \max_{\{\phi_k\}_{k=1}^K, \psi} \Delta^{(n)}(P_{X_{\mathcal{K}},Y}) \tag{10}$$

$$\text{such that } nR_k \geq \log |\phi_k(X_k^n)| \quad \text{for all} \quad k \in \mathcal{K}, \tag{11}$$

gives the largest *relevance* for prescribed complexity levels as measured by the rate tuple $R_{\mathcal{K}} = (R_1, \ldots, R_K)$, and will be referred to hereafter as the *relevance-complexity function*.

The main the goal of this paper is to learn the encoders and decoders that achieve any relevance-complexity tuple that lies on the boundary on the region $\mathcal{RI}_{\text{DIB}}$ of optimal relevance-complexity tradeoffs. To that end, another goal of this work is to characterize the region $\mathcal{RI}_{\text{DIB}}$, for both discrete and Gaussian data models. In some cases, for the ease of the exposition, we will be content with the characterization of the relevance-complexity function $\Delta(R_{\mathcal{K}}, P_{X_{\mathcal{K}},Y})$.

III. RELEVANCE-COMPLEXITY TRADEOFFS

A. Relevance-Complexity Region

In this section we first characterize the optimal relevance-complexity region for the distributed learning model of Figure 1. It is now well known that the IB problem in (1) is essentially a remote point-to-point source coding

problem [35] in which distortion is measured under the logarithm loss (log-loss) fidelity criterion [63]. That is, rather than just assigning a deterministic value to each sample of the source, the decoder gives an assessment of the degree of confidence or reliability on each estimate. Specifically, given the output description $j = \phi(x^n)$ of the encoder, the decoder generates a soft-estimate $\hat{y}^n$ of y^n in the form of a probability distribution over $\mathcal{Y}^n$, i.e., $\hat{y}^n = \hat{P}_{Y^n|J}(\cdot)$. The incurred discrepancy between y^n and the estimation $\hat{y}^n$ under log-loss for the observation x^n is then given by

$$\ell_{\log}(y^n, \hat{y}^n) = \frac{1}{n} \log \frac{1}{\hat{P}_{Y^n|J}(y^n|\phi(x^n))}, \tag{12}$$

where $\hat{P}_{Y^n|J}(y^n|\phi(x^n))$ is the value of this distribution for $y^n \in \mathcal{Y}^n$ evaluated for $j = \phi(x^n)$, $x^n \in \mathcal{X}^n$. The optimal tradeoff between complexity and relevance for the IB problem are characterized by the region given by the union of all relevance-complexity pairs $(\Delta, R) \in \mathbb{R}_+^2$ satisfying

$$\Delta \le I(Y; U) \quad \text{and} \quad R \ge I(X; U), \tag{13}$$

for some pmf $P_{U|X}$ with joint distribution of the form $p_Y(y)p_{X|Y}(x|y)p_{U|X}(u|x)$, where U is a latent representation of X. The boundary of this region is equivalent to the one described by the IB principle in (1) if solved for all s.

Likewise, the distributed learning model in Figure 1 is essentially a K-encoder CEO source coding problem under log-loss distortion, in which the decoder is interested in a soft estimate of Y from the descriptions $J_\mathcal{K}$ generated from the observations $X_1, \ldots, X_K$. A shown in Section VII-A, the equivalence of the two problems allows us to characterize the complexity-relevance region $\mathcal{RI}_{\mathrm{DIB}}$ through the rate-distortion region of the K-encoder CEO problem, which has been recently characterized in [59, Theorem 10].

Theorem 1. *The relevance-complexity region $\mathcal{RI}_{\mathrm{DIB}}$ of the distributed learning problem with $P_{X_\mathcal{K},Y}$ for which the Markov chain (3) holds, is given by the union of all tuples $(\Delta, R_1, \ldots, R_K) \in \mathbb{R}_+^{K+1}$ satisfying for all $\mathcal{S} \subseteq \mathcal{K}$,*

$$\Delta \le \sum_{k \in \mathcal{S}} [R_k - I(X_k; U_k|Y, T)] + I(Y; U_{\mathcal{S}^c}|T), \tag{14}$$

for some set of pmfs $\mathbf{P} := \{P_{U_1|X_1,T}, \ldots, P_{U_K|X_K,T}, P_T\}$ with joint distribution of the form

$$p_T(t)p_Y(y) \prod_{k=1}^K p_{X_k|Y}(x_k|y) \prod_{k=1}^K p_{U_k|X_k,T}(u_k|x_k, t). \tag{15}$$

Proof. The proof appears in Section VII-A. $\qquad\square$

Remark 1. *For a given joint data distribution $P_{X_\mathcal{K},Y}$, Theorem 1 extends the single encoder IB principle of Tishby [18] to the distributed learning model with K encoders, which we denote by Distributed Information Bottleneck (DIB) problem. The result characterizes the optimal relevance-complexity tradeoff as a region of achievable tuples $(\Delta, R_1, \ldots, R_K)$ in terms of a distributed representation learning problem involving the optimization over K conditional pmfs $P_{U_k|X_k,T}$ and a pmf P_T. The pmfs $P_{U_k|X_k,T}$ correspond to stochastic encodings of the observation X_k to a latent variable, or representation, U_k which captures the relevant information of Y in observation X_k. Variable T corresponds to a time-sharing among different encoding mappings, see, e.g., [61]. For such encoders,*

the optimal decoder is implicitly given by the conditional pmf of Y from $U_1, \ldots, U_K$, i.e., $P_{Y|U_\mathcal{K}, T}$.

Remark 2. *In the proof of Theorem 1 it is shown that maximizing relevance, as given in (6), is equivalent to minimizing the average log-loss. Also, it shows that the maximum relevance (and minimum log-loss) is achieved by employing the decoder ψ that generates a soft estimation of Y from the descriptions $J_\mathcal{K}$, given by the conditional distribution $\psi(J_\mathcal{K}) = \{P_{Y^n|J_\mathcal{K}}(y^n|J_\mathcal{K})\}_{y^n \in \mathcal{Y}^n}$.*

Remark 3. *By Proposition 4, it follows that maximizing the relevance in the distributed learning model is also equivalent to maximizing the average log-likelihood of Y from the descriptions $U_\mathcal{K}$. Thus, $\mathcal{RI}_{\mathrm{DIB}}$ also characterizes the optimal tradeoff between likelihood and description complexity in a distributed scenario with K encoders in Figure 1. In addition, if the destination is interested in reconstructing the observations, i.e., $Y = (X_1, \ldots, X_K)$, e.g., as done with Variational Autoencoders (VAEs) [24], [25], the region $\mathcal{RI}_{\mathrm{DIB}}$ also characterizes the optimal tradeoff achievable between maximum log-likelihood and complexity for this case.*

Remark 4. *The characterization of the optimal relevance-complexity region for data models satisfying the Markov chain (3) is connected to the CEO problem under logarithmic loss distortion measure, fully characterized in [59, Theorem 10]. For general IB data models, i.e., without the Markov chain (3) the model connects with the distributed source coding problem under logarithmic loss, whose solution is a longstanding open problem in source coding [59]. Note that while in some practical cases (3) might is not be satisfied, e.g., when X_k are several pictures of a particular object from different angles, the proposed methods in this paper are still applicable, without the interpretation of relevance and complexity.*

B. A Variational Bound

In this section, we consider the problem of learning the encoders and decoders that maximize relevance given a complexity constraint, i.e., that perform on the boundary of $\mathcal{RI}_{\mathrm{DIB}}$ for given $(R_1, \ldots, R_K)$. However, directly learning such encoders is challenging. In what follows, we find a parameterization of the relevance-complexity region characterized in Theorem 1 and derive a variational bound which expresses the optimal encoding and decoding mappings as the solution to an optimization of the average logarithmic loss with a novel regularization term. In Section IV and Section V we provide algorithms to solve this optimization problem. For simplicity, we focus on maximizing relevance under sum-complexity constraint, i.e., $R_{\mathrm{sum}} := \sum_{k=1}^{K} R_k$. The region of achievable relevance-complexity tuples under sum-complexity constraint is defined by:

$$\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{sum}} := \Big\{ (\Delta, R_{\mathrm{sum}}) \in \mathbb{R}_+^2 : \exists (R_1, \ldots, R_K) \in \mathbb{R}_+^K \text{ s.t. } (\Delta, R_1, \ldots, R_K) \in \mathcal{RI}_{\mathrm{DIB}} \text{ and } \sum_{k=1}^{K} R_k = R_{\mathrm{sum}} \Big\}.$$

The region $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{sum}}$ can be characterized as given next.

Proposition 1. *The relevance-complexity region $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{sum}}$ is given by the convex-hull of all tuples $(\Delta, R_{\mathrm{sum}}) \in \mathbb{R}_+^2$*

satisfying $\Delta \leq \Delta(R_{\mathrm{sum}}, P_{X_{\mathcal{K}},Y})$ where

$$\Delta(R_{\mathrm{sum}}, P_{X_{\mathcal{K}},Y}) = \max_{\mathbf{P}} \min \left\{ I(Y; U_{\mathcal{K}}), R_{\mathrm{sum}} - \sum_{k=1}^{K} I(X_k; U_k | Y) \right\}, \tag{16}$$

and where the maximization is over the set of pmfs $\mathbf{P} := \{P_{U_1|X_1}, \ldots, P_{U_K|X_K}\}$ such that the joint pmf factorizes as $p_Y(y) \prod_{k=1}^{K} p_{X_k|Y}(x_k|y) \prod_{k=1}^{K} p_{U_k|X_k}(u_k|x_k)$.

Proof. The proof is given in Section VII-B. $\qquad\square$

The following proposition provides a characterization of the pairs $(\Delta, R_{\mathrm{sum}})$ on the boundary of $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{sum}}$ in terms of a parameter $s \geq 0$, as (Δ_s, R_s), which is more suitable for the derivation of the variational bound.

Proposition 2. *For each tuple* $(\Delta, R_{\mathrm{sum}}) \in \mathbb{R}_+^2$ *on the boundary of the relevance-complexity region* $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{sum}}$ *there exist* $s \geq 0$ *such that* $(\Delta, R_{\mathrm{sum}}) = (\Delta_s, R_s)$, *where*

$$\Delta_s = \frac{1}{(1+s)} \left[(1 + sK)H(Y) + sR_s + \max_{\mathbf{P}} \mathcal{L}_s(\mathbf{P}) \right], \tag{17}$$

$$R_s = I(Y; U_{\mathcal{K}}^*) + \sum_{k=1}^{K} [I(X_k; U_k^*) - I(Y; U_k^*)], \tag{18}$$

and $\mathbf{P}^*$ *is the set of pmfs* $\mathbf{P}$ *that maximize the cost function*

$$\mathcal{L}_s(\mathbf{P}) := -H(Y|U_{\mathcal{K}}) - s \sum_{k=1}^{K} [H(Y|U_k) + I(X_k; U_k)]. \tag{19}$$

Proof. The proof appears in Section VII-C. $\qquad\square$

Remark 5. *The optimization of the DIB cost in* (19) *generalizes the centralized Tishby's information bottleneck formulation in* (1) *to the distibuted learning problem with* K-*encoders in Figure 1. Note that for* $K = 1$ *the optimization in* (17) *reduces to the single encoder cost in* (1) *with a multiplier* $s/(1 + s)$.

Using Proposition 2 it is clear that the encoders $\{P_{U_k|X_k}\}_{k \in \mathcal{K}}$ that achieve the relevance-complexity pair (Δ_s, R_s) can be computed by maximizing the regularized cost (19) for the corresponding value of $s \geq 0$. The corresponding optimal decoder $P_{Y|U_{\mathcal{K}}}$ for these encoders can be found as $P_{Y|U_{\mathcal{K}}}$. Different relevance-complexity pairs (Δ_s, R_s) on the boundary of $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{sum}}$ and the corresponding encoders and decoder achieving it can be obtained by solving (19) for different values of $s \geq 0$ and evaluating (17) and (18) for the resulting solution.

The optimization of (19) generally requires to compute marginal distributions involving the descriptions $U_1, \ldots, U_K$, which might be costly to calculate. To overcome this limitation, in the following we derive a tight variational bound on $\mathcal{L}_s(\mathbf{P})$ which lower bounds the DIB cost function with respect to some arbitrary distributions. Let us consider the arbitrary decoder $Q_{Y|U_1,\ldots,U_K}(y|u_1, \ldots, u_K)$ for $y \in \mathcal{Y}$, $u_1 \in \mathcal{U}_1, \ldots, u_K \in \mathcal{U}_K$, the K decoders $Q_{Y|U_k}(y|u_k)$ for $k \in \mathcal{K}$ for $y \in \mathcal{Y}$, $u_k \in \mathcal{U}_k$, and latent variable priors $Q_{U_k}(u_k)$, $k \in \mathcal{K}$, $u_k \in \mathcal{U}_k$. For short, we denote

$$\mathbf{Q} := \{Q_{Y|U_1,\ldots,U_K}, Q_{Y|U_1}, \ldots, Q_{Y|U_K}, Q_{U_1}, \ldots, Q_{U_K}\}.$$

We define the *variational DIB cost function* $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$ as

$$\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}) := \underbrace{\mathbb{E}[\log Q_{Y|U_{\mathcal{K}}}(Y|U_{\mathcal{K}})]}_{\text{av. logarithmic-loss}} + s \underbrace{\sum_{k=1}^{K} \Big(\mathbb{E}[\log Q_{Y|U_k}(Y|U_k)] - D_{\mathrm{KL}}(P_{U_k|X_k} \| Q_{U_k}) \Big)}_{\text{regularizer}}.$$

Next lemma states that $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$ is a lower bound to $\mathcal{L}_s(\mathbf{P})$ for all distributions $\mathbf{Q}$.

Lemma 1. *For fixed pmfs $\mathbf{P}$, we have*

$$\mathcal{L}_s(\mathbf{P}) \geq \mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}), \qquad \text{for all pmfs } \mathbf{Q}. \tag{20}$$

In addition, there exists a unique $\mathbf{Q}$ that achieves the maximum $\max_{\mathbf{Q}} \mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}) = \mathcal{L}_s(\mathbf{P})$, and is given by

$$Q_{U_k}^* = P_{U_k}, \qquad Q_{Y|U_k}^* = P_{Y|U_k}, \quad k = 1, \ldots, K, \tag{21}$$

$$Q_{Y|U_1,\ldots,U_k}^* = P_{Y|U_1,\ldots,U_K}, \tag{22}$$

where P_{U_k}, $P_{Y|U_k}$ and $P_{Y|U_1,\ldots,U_K}$ are computed from $\mathbf{P}$.

Proof. The proof appears in Section VII-D. □

Using the above, the optimization in (17) can be written in terms of the variational DIB cost function as follows:

$$\max_{\mathbf{P}} \mathcal{L}_s(\mathbf{P}) = \max_{\mathbf{P}} \max_{\mathbf{Q}} \mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}). \tag{23}$$

Remark 6. *The variational DIB cost function in (20) consist of the average logarithmic loss (or cross-entropy) of estimating Y form all the latent space representations $U_1, \ldots, U_K$ using the decoder $Q_{Y|U_1,\ldots,U_K}$, and a regularization term that is composed of two types of terms: i) the Kullback-Leiber divergence between encoders $P_{U_k|X_k}$ and the priors Q_{U_k}, which also appears in the single encoder version of the bound, and ii) the logarithmic loss of estimating the target variable Y from each individual latent space representation U_k using a decoder $Q_{Y|U_k}$, which does not appear in the single encoder case. An example of learning architecture which can be trained to minimize this cost function using neural networks is shown in Figure 2 for $K = 2$ observations.*

Remark 7. *The variational DIB cost in (20) is a generalization to distributed learning with K-encoders of the evidence lower bound (ELBO) of the target variable Y given the representations $U_1, \ldots, U_K$ [34], [64]. If $Y = (X_1, \ldots, X_K)$, the bound generalizes the ELBO used for VAEs to the setting of $K \geq 2$ encoders. Also note that (20) provides an operational meaning to the β-VAE cost [25] with $\beta = s/(1+s)$, as a criteria to design estimators on the relevance-complexity plane, for different β values.*

C. Relevance-Complexity for Vector Gaussian Model

In this section we consider the memoryless vector Gaussian data model and show that the encoders and decoders maximizing the relevance-complexity tradeoff correspond to Gaussian distributions. In this model, $(\mathbf{X}_1, \ldots, \mathbf{X}_K, \mathbf{Y})$ are jointly Gaussian and satisfy the Markov chain (3). Without loss in generality, let the target variable be a zero-mean

multivariate Gaussian $\mathbf{Y} \in \mathbb{C}^{n_y}$, with covariance matrix $\boldsymbol{\Sigma_y}$, i.e., $\mathbf{Y} \sim \mathcal{CN}(\mathbf{y}; \mathbf{0}, \boldsymbol{\Sigma_y})$. Encoder k, $k = 1, \ldots, K$, observes a noisy observation $\mathbf{X}_k \in \mathbb{C}^{n_k}$, that is given by

$$\mathbf{X}_k = \mathbf{H}_k \mathbf{Y} + \mathbf{N}_k, \tag{24}$$

where $\mathbf{H}_k \in \mathbb{C}^{n_k \times n_y}$ models the linear model connecting $\mathbf{Y}$ to the observation at encoder k, and $\mathbf{N}_k \in \mathbb{C}^{n_k}$, $k = 1, \ldots, K$, is the noise vector at encoder k, assumed to be Gaussian with zero-mean and covariance matrix $\boldsymbol{\Sigma}_k$, and independent from all other noises and $\mathbf{Y}$. Note that it can be shown that for every jointly Gaussian random vector $(\mathbf{X}_1, \ldots, \mathbf{X}_K, \mathbf{Y})$ satisfying (3), there exist matrices $(\mathbf{H}_1, \ldots, \mathbf{H}_K)$ and noises $\mathbf{N}_1, \ldots, \mathbf{N}_K$ as defined.

The vector Gaussian model satisfies the Markov chain (3); and thus, the result of Theorem 1, which can be extended to continuous sources using standard techniques [61], characterizes the relevance-complexity region of this model. Let us denote by $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{G}}$ the relevance-complexity region of the Gaussian model (24). Next theorem characterizes $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{G}}$, and shows that there is not need for time-sharing, i.e., $T = \emptyset$ and that the optimal stochastic mappings $P_{U_k|X_k}$, $k \in \mathcal{K}$, can be restricted without loss in optimality to be multivariate Gaussian distributions as,

$$\mathbf{U}_k = \mathbf{A}_k \mathbf{X}_k + \mathbf{Z}_k \sim \mathcal{CN}(\mathbf{u}_k; \mathbf{A}_k \mathbf{X}_k, \boldsymbol{\Sigma}_{z,k}), \tag{25}$$

where $\mathbf{A}_k \in \mathbb{C}^{n_k \times n_k}$ projects the observation $\mathbf{X}_k$ and $\mathbf{Z}_k$ is a zero-mean Gaussian noise with covariance $\boldsymbol{\Sigma}_{z,k}$.

Theorem 2. *The relevance-complexity region $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{G}}$ for the vector Gaussian model is given by the union of all tuples $(\Delta, R_1, \ldots, R_L)$ satisfying for all $\mathcal{S} \subseteq \mathcal{K}$*

$$\Delta \leq \sum_{k \in \mathcal{S}} \left(R_k + \log \left| \mathbf{I} - \boldsymbol{\Sigma}_k^{1/2} \boldsymbol{\Omega}_k \boldsymbol{\Sigma}_k^{1/2} \right| \right) + \log \left| \mathbf{I} + \sum_{k \in \mathcal{S}^c} \boldsymbol{\Sigma}_\mathbf{y}^{1/2} \mathbf{H}_k^\dagger \boldsymbol{\Omega}_k \mathbf{H}_k \boldsymbol{\Sigma}_\mathbf{y}^{1/2} \right|,$$

for some $\mathbf{0} \preceq \boldsymbol{\Omega}_k \preceq \boldsymbol{\Sigma}_k^{-1}$. In addition, the relevance-complexity region is achievable with $T = \emptyset$ and pmfs

$$P_{U_k|X_k,T}^*(\mathbf{u}_k|\mathbf{x}_k, t) = \mathcal{CN}(\mathbf{u}_k; \mathbf{x}_k, \boldsymbol{\Sigma}_k^{1/2}(\boldsymbol{\Omega}_k - \mathbf{I})\boldsymbol{\Sigma}_k^{1/2}).$$

Proof. The proof is given in Section VII-E. $\qquad\square$

Remark 8. *Theorem 2 extends the result of [29], [30] and [65] on the relevance-complexity tradeoff characterization of the single-encoder IB problem for jointly Gaussian sources (3) to K encoders.*

Note that, from (25) and the proof of Thereom 2, $\boldsymbol{\Omega}_k$, $k = 1, \ldots, K$ can be understood as parameterizations of the covariance matrixs of the noises $\boldsymbol{\Sigma}_{z,k}$ in (25). By Theorem 2, the optimal encoders $P_{U_k|X_k}$ are multivariate Gaussian distributions as in (25). Thus, the decoder $P_{Y|U_1,\ldots,U_K}$ is also a multivariate Gaussian distribution, since any conditional distribution of jointly Gaussian variables is also a Gaussian distribution. Similarly, the optimal distributions $\mathbf{Q}$ in Lemma 1 for given multivariate Gaussian distributions $\mathbf{P}$ are multivariate Gaussian distributions.

Corollary 1. *If $(\mathbf{X}_1, \ldots, \mathbf{X}_K, \mathbf{Y})$ are jointly Gaussian as in (24), for given encoders $\mathbf{P}$ with multivariate Gaussian distribution as in (25), the optimal distributions $\mathbf{Q}$ in (21) and (22) are multivariate Gaussian distributions.*

For the vector Gaussian model, the optimal encoding and decoding mappings can be computed explicitly. However, in general closed form expressions are unfeasible. In the following sections we provide algorithms to find mappings that approximate these optimal distributions.

IV. BLAHUT-ARIMOTO DIB ALGORITHM

In this section, we derive an iterative method to optimize the variational DIB cost function in (23) when the data model is discrete and the joint distribution $P_{X_\mathcal{K}, Y}$ is either known, or a good estimation of it can be obtained from the training samples. In these cases, the maximizing distributions $\mathbf{P}, \mathbf{Q}$ of the variational DIB cost in (23) can be efficiently found by an alternating optimization procedure over $\mathbf{P}$ and $\mathbf{Q}$ similar to the expectation-maximization (EM) algorithm [26] and the standard Blahut-Arimoto (BA) method [60]. An extension to the vector Gaussian data model, which involves random variable with continuous alphabets, is also provided. The main idea of the algorithm is that at iteration t, the optimal distributions $\mathbf{P}^{(t)}$ that maximize the variational D-IB bound $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}^{(t)})$ for fixed $\mathbf{Q}^{(t)}$ can be optimized in closed form and, next, the maximizing pmfs $\mathbf{Q}^{(t)}$ for given $\mathbf{P}^{(t)}$ can be also found analytically. So, starting from an initialization $\mathbf{P}^{(0)}$ and $\mathbf{Q}^{(0)}$ the algorithms performs the following computations successively, until convergence,

$$\mathbf{P}^{(0)} \to \mathbf{Q}^{(0)} \to \mathbf{P}^{(1)} \to \dots \to \mathbf{P}^{(t)} \to \mathbf{Q}^{(t)} \to \dots \tag{26}$$

We refer to such algorithm as "Blahut-Arimoto Distributed Information Bottleneck Algorithm (BA-DIB)". Algorithm 1 describes the steps taken by BA-DIB to successively maximize $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$ by solving a concave optimization problem over $\mathbf{P}$ and over $\mathbf{Q}$ at each iteration.

Lemma 2. *The function $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$ is concave in $\mathbf{P}$ and in $\mathbf{Q}$.*

Proof. The proof follows by using the log-sum inequality [66] and the convexity of the mapping $x \mapsto x \log x$; and it is omitted for brevity. $\qquad\square$

For fixed $\mathbf{P}^{(t)}$, the optimal $\mathbf{Q}^{(t)}$ maximizing the variational D-IB bound in (20) follows from Lemma 1 as given by (21)-(22). For fixed $\mathbf{Q}^{(t)}$, the optimal $\mathbf{P}^{(t)}$ can be found using the following lemma.

Lemma 3. *For fixed $\mathbf{Q}$, there exists a $\mathbf{P}$ that achieves the maximum $\max_{\mathbf{P}} \mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$, where $P_{U_k|X_k}$ is given by*

$$p^*(u_k|x_k) = q(u_k) \frac{\exp\left(-\psi_s(u_k, x_k)\right)}{\sum_{u_k \in \mathcal{U}_k} q(u_k) \exp(-\psi_s(u_k, x_k))}, \tag{27}$$

for $u_k \in \mathcal{U}_k$ and $x_k \in \mathcal{X}_k$, $k \in \mathcal{K}$, and where we define

$$\psi_s(u_k, x_k) := D_{\mathrm{KL}}(P_{Y|x_k} || Q_{Y|u_k}) + \frac{1}{s} \mathbb{E}_{U_{\mathcal{K}\setminus k}|x_k} [D_{\mathrm{KL}}(P_{Y|U_{\mathcal{K}\setminus k}, x_k} || Q_{Y|U_{\mathcal{K}\setminus k}, u_k}))]. \tag{28}$$

Proof. Due to its concavity, to maximize $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$ with respect to $\mathbf{P}$ for given $\mathbf{Q}$, we add the Lagrange multipliers $\lambda_{x_k} \geq 0$ for each constraint $\sum_{u_k \in \mathcal{U}_k} p(u_k|x_k) = 1$ with $x_k \in \mathcal{X}_k$. For each s, $\lambda_{x_k} \geq 0$ and $p(u_k|x_k)$ can be

---

Algorithm 1 BA-DIB training algorithm for discrete data

1: **Inputs:**
   discrete pmf $P_{X_1,\dots,X_k,Y}$, parameter $s \geq 0$.
2: **output:** optimal $P^*_{U_k|X_k}$, pair (Δ_s, R_s).
3: **initialization**
   Set $t = 0$ and set $\mathbf{P}^{(0)}$ with $p(u_k|x_k) = \frac{1}{|\mathcal{U}_k|}$
   for $u_k \in \mathcal{U}_k$, $x_k \in \mathcal{X}_k$, $k = 1, \dots, K$.
4: **repeat**
5:     Compute $\mathbf{Q}^{(t+1)}$ using (21) and (22).
6:     Compute $\mathbf{P}^{(t+1)}$ using (27).
7:     $t \leftarrow t + 1$.
8: **until** convergence.

---

explicitly found by solving the KKT conditions, e.g.,

$$\frac{\partial}{\partial p(u_k|x_k)} \left[\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}) + \sum_{x_k \in \mathcal{X}_k} \lambda_{x_k} \left(\sum_{u_k \in \mathcal{U}_k} p(u_k|x_k) - 1 \right) \right] = 0.$$

This completes the proof. □

After convergence of the encoders $\mathbf{P}^*$ and decoders $\mathbf{Q}^*$, the target variable Y can be inferred for a new observation using $P^*_{U_K|X_k}(U_k|X_k)$ and the soft estimate provided by $Q^*_{Y|U_1,\dots,U_K}(Y|U_1,\dots,U_K)$.

A. Convergence of the BA-DIB Algorithm

Algorithm 1 essentially falls into the class of the Successive Upper-Bound Minimization (SUM) algorithms [67] in which $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$ acts as a globally tight lower bound on $\mathcal{L}_s(\mathbf{P})$. Algorithm 1 provides a sequence $\mathbf{P}^{(t)}$ for each iteration t, which converges to a stationary point of the problem in (23).

Proposition 3. *Every limit point of the sequence $\mathbf{P}^{(t)}$ generated by Algorithm 1 converges to a stationary point of* (23).

Proof. Let $\mathbf{Q}^*(\mathbf{P}) := \arg\max_{\mathbf{Q}} \mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q})$. From Lemma 1, $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}^*(\mathbf{P}')) \leq \mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}^*(\mathbf{P})) = \mathcal{L}_s(\mathbf{P})$ for $\mathbf{P}' \neq \mathbf{P}$. Since $\mathcal{L}_s(\mathbf{P})$ and $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}^*(\mathbf{P}'))$ satisfy [67, Proposition 1], then $\mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}^*(\mathbf{P}'))$ satisfies A1-A4 in [67]. Convergence to a stationary point of (23) is due to [67, Theorem 1]. □

Remark 9. *The resulting set of self consistent equations* (21)*,* (22) *and* (28) *satisfied by any stationary point of the D-IB problem extend those of the standard point-to-point IB problem* [29] *to the distributed IB problem with $K \geq 2$ encoders. The resulting equations are reminiscent of those for the point-to-point IB problem with an additional divergence term in* (28) *for encoder k averaged over the descriptions at the other $\mathcal{K} \setminus k$ encoders.*

B. BA-DIB Algorithm for the Vector Gaussian Model

Computing the relevance-complexity for the vector Gaussian model from Theorem 2 is a convex optimization problem on $\mathbf{\Omega}_k$, which can be efficiently solved with generic tools. In the following, we extend Algorithm 1 as an alternative method to maximize relevance under sum-complexity constraint for vector Gaussian models.

Algorithm 2 BA-DIB algorithm for the Gaussin Vector D-IB

1: **Inputs:**
   covariance $\boldsymbol{\Sigma}_{\mathbf{y},\mathbf{x}_1,\ldots,\mathbf{x}_k}$, parameter $s \geq 0$.
2: **output:** optimal pairs $(\mathbf{A}_k^*, \boldsymbol{\Sigma}_{\mathbf{z}_k^*})$, $k = 1, \ldots, K$.
3: **initialization**
   Set randomly $\mathbf{A}_k^0$ and $\boldsymbol{\Sigma}_{\mathbf{z}_k^0} \succeq 0$, $k \in \mathcal{K}$.
4: **repeat**
5:    Compute $\boldsymbol{\Sigma}_{\mathbf{x}_k|\mathbf{u}_{\mathcal{K}\setminus k}^t}$ and update for $k \in \mathcal{K}$

$$\boldsymbol{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}} = \mathbf{A}_k^t \boldsymbol{\Sigma}_{\mathbf{x}_k|\mathbf{y}} \mathbf{A}_k^{t,\dagger} + \boldsymbol{\Sigma}_{\mathbf{z}_k^t} \tag{32}$$

$$\boldsymbol{\Sigma}_{\mathbf{u}_k^t|\mathbf{u}_{\mathcal{K}\setminus k}^t} = \mathbf{A}_k^t \boldsymbol{\Sigma}_{\mathbf{x}_k|\mathbf{u}_{\mathcal{K}\setminus k}^t} \mathbf{A}_k^{t,\dagger} + \boldsymbol{\Sigma}_{\mathbf{z}_k^t}, \tag{33}$$

6:    Compute $\boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}}$ as in (30) for $k \in \mathcal{K}$.
7:    Compute $\mathbf{A}_k^{t+1}$ as (31), $k \in \mathcal{K}$.
8:    $t \leftarrow t + 1$.
9: **until** convergence.

For finite alphabet sources the updating rules of $\mathbf{Q}^{(t+1)}$ and $\mathbf{P}^{(t+1)}$ in Algorithm 1 are relatively easy, but they become unfeasible for continuous alphabet sources. We leverage on the optimality of Gaussian encoders, shown in Theorem 2, to restrict the optimization of $\mathbf{P}$ to Gaussian distributions, which are represented by a finite set of parameters, namely its mean and covariance. We show that if $\mathbf{P}^{(t)}$ are Gaussian distributions, then $\mathbf{P}^{(t+1)}$ are also Gaussian distributions, and can be computed with an efficient update algorithm of its representing parameters. In particular, if at time t the k-th distributions $P_{\mathbf{U}_k|\mathbf{X}_k}^{(t)}$ is given by

$$\mathbf{U}_k^t = \mathbf{A}_k^t \mathbf{X}_k + \mathbf{Z}_k^t, \tag{29}$$

where $\mathbf{Z}_k^t \sim \mathcal{CN}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{z}_k^t})$, we show that at $t+1$, for $\mathbf{P}^{(t+1)}$ updated as in (27), the encoder $P_{\mathbf{U}_k|\mathbf{X}_k}^{(t+1)}$ corresponds to $\mathbf{U}_k^{t+1} = \mathbf{A}_k^{t+1} \mathbf{X}_k + \mathbf{Z}_k^{t+1}$, where $\mathbf{Z}_k^{t+1} \sim \mathcal{CN}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}})$ and $\boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}}, \mathbf{A}_k^{t+1}$ are updated as

$$\boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}} = \left(\left(1 + \frac{1}{s} \right) \boldsymbol{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1} - \frac{1}{s} \boldsymbol{\Sigma}_{\mathbf{u}_k^t|\mathbf{u}_{\mathcal{K}\setminus k}^t}^{-1} \right)^{-1}, \tag{30}$$

$$\mathbf{A}_k^{t+1} = \boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}} \left(\left(1 + \frac{1}{s} \right) \boldsymbol{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1} \mathbf{A}_k^t (\mathbf{I} - \boldsymbol{\Sigma}_{\mathbf{x}_k|\mathbf{y}} \boldsymbol{\Sigma}_{\mathbf{x}_k}^{-1}) - \frac{1}{s} \boldsymbol{\Sigma}_{\mathbf{u}_k^t|\mathbf{u}_{\mathcal{K}\setminus k}^t}^{-1} \mathbf{A}_k^t (\mathbf{I} - \boldsymbol{\Sigma}_{\mathbf{x}_k|\mathbf{u}_{\mathcal{K}\setminus k}^t} \boldsymbol{\Sigma}_{\mathbf{x}_k}^{-1}) \right). \tag{31}$$

The detailed update procedure is given in Algorithm 2, and the details of the derivation are relegated to Section VII-F.

Remark 10. *Algorithm 2 generalizes the iterative algorithm for single encoder Gaussian IB in [30] to the Gaussian D-IB with K encoders and sum-complexity constraint. Similarly to the solution in [30], the optimal description at each encoder is given by a noisy linear projection of the observation whose dimensionality is determined by the regularization parameter s and the second order moments between the observed data and the target variable, as well as a term depending on the observed data with*

V. Variational Distributed IB Algorithm

In cases in which the joint distribution of the data is either known perfectly or can be estimated to high accuracy, the maximizing distributions $\mathbf{P}, \mathbf{Q}$ of the variational DIB cost in (23) can be found through BA-DIB algorithm as shown in Section IV. However, in general only a set of training samples $\{(X_{1,i}, \ldots, X_{K,i}, Y_i)\}_{i=1}^n$ are available. In this section, we provide a method to optimize (23) in this situation by parameterizing the encoding and decoding distributions that are to optimize using a family of distributions whose parameters are determined by DNNs. This allows us to formulate (23) in terms of the DNN parameters and optimize it by using the reparameterization trick [64], Monte Carlo sampling, as well as stochastic gradient descent (SGD) type algorithms. The proposed method generalizes the variational framework in [31]–[33], [47], [48] to the scenario with K encoders in Figure 1.

Let $P_{\theta_k}(u_k|x_k)$ denote the family of encoding probability distributions $P_{U_k|X_k}$ over $\mathcal{U}_k$ for each element on $\mathcal{X}_k$, parameterized by the output of a DNN f_{θ_k} with parameters θ_k. A common example is the family of multivariate Gaussian distributions, [32], [64], which are parameterized by the mean $\boldsymbol{\mu}_k^\theta$ and covariance matrix $\boldsymbol{\Sigma}_k^\theta$, i.e., $\boldsymbol{\gamma}_k := (\boldsymbol{\mu}_k^\theta, \boldsymbol{\Sigma}_k^\theta)$. Given an observation X_k, the values of $(\boldsymbol{\mu}_k^\theta, \boldsymbol{\Sigma}_k^\theta)$ are determined by the output of the DNN f_{θ_k} and the corresponding family member is given by $P_{\theta_k}(u_k|x_k) = \mathcal{N}(u_k; \boldsymbol{\mu}_k^\theta, \boldsymbol{\Sigma}_k^\theta)$. For discrete distributions, a common example are concrete variables [68] (or Gumbel-Softmax [69]). More expressive paremeterizations can also be considered, e.g., [70], [71].

Similarly, for decoders $Q_{Y|U_k}$ over $\mathcal{Y}$ for each element on $\mathcal{U}_k$, and the decoding distribution $Q_{Y|U_\mathcal{K}}$ over $\mathcal{Y}$ for each element in $\mathcal{U}_1 \times \cdots \times \mathcal{U}_K$, let $Q_{\psi_k}(y|u_k)$, $Q_{\psi_\mathcal{K}}(y|u_\mathcal{K})$ denote the families of distributions parameterized by the output of the DNNs $f_{\psi_k}, f_{\psi_\mathcal{K}}$, respectively. Finally, for the prior distributions $Q_{U_k}(u_k)$ over $\mathcal{U}_k$ we define the family of distributions $Q_{\varphi_k}(u_k)$, which do not depend on a DNN.

By restricting the optimization of the variational DIB cost in (23) to the encoder, decoder and priors within the families of distributions $P_{\theta_k}(u_k|x_k)$, $Q_{\psi_k}(y|u_k)$, $Q_{\psi_\mathcal{K}}(y|u_\mathcal{K})$ and $Q_{\varphi_k}(u_k)$ we get

$$\max_{\mathbf{P}} \max_{\mathbf{Q}} \mathcal{L}_s^{\mathrm{VB}}(\mathbf{P}, \mathbf{Q}) \geq \max_{\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}} \mathcal{L}_s^{\mathrm{NN}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}), \tag{34}$$

where we use the notation $\boldsymbol{\theta} := [\theta_1, \ldots, \theta_K]$, $\boldsymbol{\phi} := [\phi_1, \ldots, \phi_K, \phi_\mathcal{K}]$ and $\boldsymbol{\varphi} := [\varphi_1, \ldots, \varphi_K]$ to denote the DNN and prior parameters and, the cost in (34) is given by

$$\mathcal{L}_s^{\mathrm{NN}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}) := \mathbb{E}_{P_{Y,X}} \mathbb{E}_{\{P_{\theta_k}(U_k|X_k)\}} \Big[\log Q_{\phi_\mathcal{K}}(Y|U_\mathcal{K}) \\ + s \sum_{k=1}^K \Big(\log Q_{\phi_k}(Y|U_k) - D_{\mathrm{KL}}(P_{\theta_k}(U_k|X_k) \| Q_{\varphi_k}(U_k)) \Big) \Big].$$

Next, we train the DNNs to maximize a Monte Carlo approximation of (34) over $\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}$ using optimization methods such as SGD or ADAM [72] with backpropagation. We use the reparameterization trick [64], to sample from $P_{\theta_k}(U_k|X_k)$. In particular, we consider $P_{\theta_k}(U_k|X_k)$. to belong to a parametric family of distributions that can be sampled by first sampling a random variable Z_k with distribution $P_{Z_k}(z_k)$, $z_k \in \mathcal{Z}_k$ and then transforming the samples using some function $g_{\theta_k} : \mathcal{X}_k \times \mathcal{Z}_k \to \mathcal{U}_k$ parameterized by θ_k, such that $U_k = g_{\theta_k}(x_k, Z_k) \sim P_{\theta_k}(U_k|x_k)$. Various parametric families of distributions fall within this class both for discrete and continuous latent spaces,

---

Algorithm 3 D-VIB training algorithm for the D-IB problem

1: **Inputs:**
   Training Dataset $\mathcal{D}$, parameter $s \geq 0$, DNN parameters $\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}$.
2: **output:** optimized $\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}$ and pair (Δ_s, R_s).
3: **initialization**
   Initialize $\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}$ and set $t = 0$.
4: **repeat**
5:    Randomly select b minibatch samples $X_{\mathcal{K}}^b = (X_1^b, \ldots, X_K^b)$ and the corresponding Y^b form $\mathcal{D}$.
6:    Draw m random samples $Z_k^m \sim P_{Z_k}$, $k = 1, \ldots, K$.
7:    Compute m samples $U_{k,l}^b = g_{\phi_k}(X^b, Z_{k,l}^b)$.
8:    Compute gradients of the empirical DIB cost in (36), $\nabla_{\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}} \sum_{i=1}^b \mathcal{L}_{s,i}^{\mathrm{emp}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi})$ for $(X_{\mathcal{K}}^b, Y^b)$.
9:    Update $(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi})$ using the estimated gradient (e.g. with SGD or ADAM).
10: **until** convergence of $(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi})$.

---

e.g., the Gumbel-Softmax distributions and the Gaussian distributions mentioned above. The reparametrization trick reduces the original optimization to estimating θ_k of the deterministic function g_{θ_k} and allows to compute estimates of the gradient using backpropagation [64]. The variational DIB cost in (34) can be approximated, by sampling m independent samples $\{u_{k,i,j}\}_{j=1}^m \sim P_{\theta_k}(u_k|x_{k,i})$ for each training sample $(x_{1,i}, \ldots, x_{K,i}, y_i)$, $i = 1, \ldots, n$. Sampling is performed by using $u_{k,i,j} = g_{\phi_k}(x_{k,i}, z_{k,j})$ with $\{z_{k,j}\}_{j=1}^m$ i.i.d. sampled from P_{Z_k}. Altogether, we have the empirical DIB cost for the i-th sample in the training dataset:

$$
\mathcal{L}_{s,i}^{\mathrm{emp}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}) := \frac{1}{m} \sum_{j=1}^m \Big[\log Q_{\phi_{\mathcal{K}}}(y_i | u_{1,i,j}, \ldots, u_{K,i,j}) \tag{35}
$$
$$
+ s \sum_{k=1}^K \Big(\log Q_{\phi_k}(y_i|u_{k,i,j}) - D_{\mathrm{KL}}(P_{\theta_k}(U_{k,i}|x_{k,i}) \| Q_{\varphi_k}(U_{k,i})) \Big) \Big].
$$

Finally, we maximize the empirical DIB cost over the DNN parameters $\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}$ over the training data as,

$$
\max_{\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}_{s,i}^{\mathrm{emp}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}). \tag{36}
$$

By the law of large numbers, for large n, m, we have $1/n \sum_{i=1}^n \mathcal{L}_{s,i}^{\mathrm{emp}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi})) \to \mathcal{L}_s^{\mathrm{NN}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi})$ almost surely.

The training of the D-VIB algorithm is detailed in Algorithm 3. After convergence of the DNN parameters to $\boldsymbol{\theta}^*, \boldsymbol{\phi}^*, \boldsymbol{\varphi}^*$, for a new observation, the target variable Y can be inferred by sampling from the encoders $P_{\theta_k^*}(U_k|X_k)$ and the soft estimate provided by decoder $Q_{\phi_{\mathcal{K}}^*}(Y|U_1, \ldots, U_K)$.

A. D-VIB for Regression and Classification

The choice of parametric distributions $P_{\theta_k}(u_k|x_k)$, $Q_{\psi_k}(y|u_k)$, $Q_{\psi_{\mathcal{K}}}(y|u_{\mathcal{K}})$ and $Q_{\varphi_k}(u_k)$ depends on the application. Nevertheless, the parametric families of distributions should be chosen to be expressive enough to approximate the optimal encoders minimizing (19) and the optimal conditional distributions (21) and (22), to minimize the gap between the lower bound given by the variational DIB cost (20) and the original DIB cost function (19).

For example, in classification problems, i.e., $\mathcal{Y}$ is finite, the decoder $Q_{\phi_{\mathcal{K}}}(Y|U_{\mathcal{K}})$ and the decoders $Q_{\phi_k}(Y|U_k)$, used for regularization can be general categorical distributions parameterized by a DNN with a softmax operation
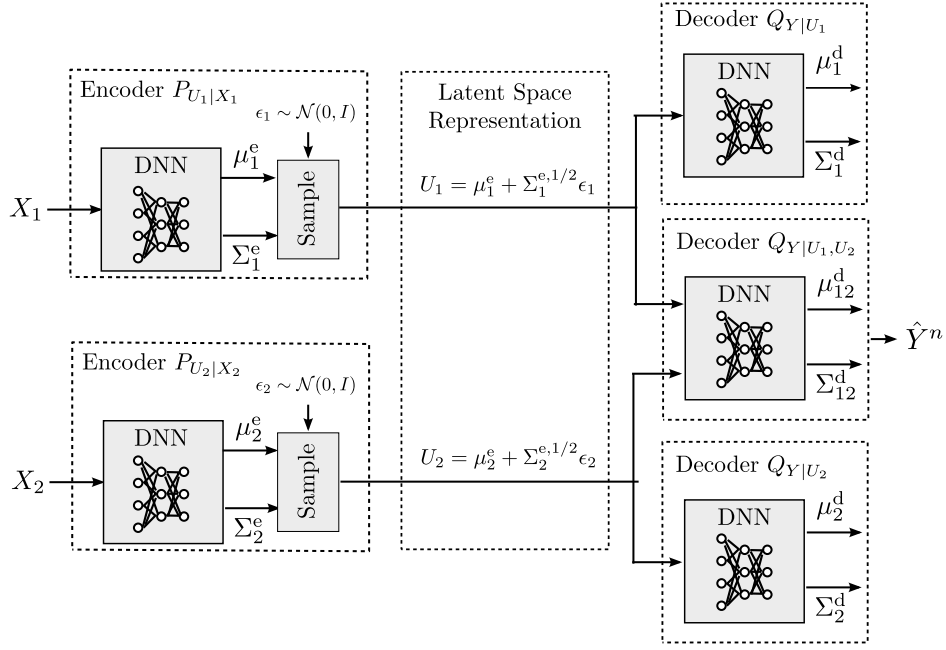
Fig. 2: Distributed learning architecture for the minimization of the variational DIB cost in (23) for $K = 2$ using multivariate Gaussian distributions to parameterize the encoders, decoders and prior distributions. The decoders $Q_{Y|U_1}$ and $Q_{Y|U_2}$ influence in the regularization.

in the last layer, which outputs a vector of dimension $|\mathcal{Y}|$. If the target variable Y is continuous, i.e., in regression problems, the decoders can be chosen to lie within the family of multivariate Gaussian distributions. The priors Q_{φ_k} and the encoders $P_{\theta_k}(U_k|X_k)$, can be chosen as multivariate Gaussian, such that the divergence in the cost function is easy to compute [32], [64], or using more expressive parameterization [70], [71]. See Section VI for some choices for different regression and classification.

B. D-VIB algorithm for Vector Gaussian Model

In the D-VIB in Algorithm 3, the performance depends on the choice of the parametric family of distributions for the encoders, decoder and prior distributions. By Corollary 1, when the underlying data model is multivariate vector Gaussian, the optimal distributions $\mathbf{P}$ and $\mathbf{Q}$ in (23) lie within the family of multivariate Gaussian distributions. Motivated by this observation, we consider the following parameterization for $k \in \mathcal{K}$:

$$P_{\theta_k}(\mathbf{u}_k|\mathbf{x}_k) = \mathcal{N}(\mathbf{u}_k; \boldsymbol{\mu}_k^{\mathrm{e}}, \boldsymbol{\Sigma}_k^{\mathrm{e}}) \tag{37}$$

$$Q_{\phi_{\mathcal{K}}}(\hat{\mathbf{y}}|\mathbf{u}_{\mathcal{K}}) = \mathcal{N}(\hat{\mathbf{y}}; \boldsymbol{\mu}_{\mathcal{K}}^{\mathrm{d}}, \boldsymbol{\Sigma}_{\mathcal{K}}^{\mathrm{d}}) \tag{38}$$

$$Q_{\phi_k}(\hat{\mathbf{y}}|\mathbf{u}_k) = \mathcal{N}(\hat{\mathbf{y}}; \boldsymbol{\mu}_k^{\mathrm{d}}, \boldsymbol{\Sigma}_k^{\mathrm{d}}) \quad k = 1, 2, \tag{39}$$

$$Q_{\varphi_k}(\mathbf{u}_k) = \mathcal{N}(\mathbf{0}, \mathbf{I}). \tag{40}$$

where $\boldsymbol{\mu}_k^{\mathrm{e}}, \boldsymbol{\Sigma}_k^{\mathrm{e}}$ are the output of a DNN f_{θ_k} with input $\mathbf{X}_k$ that encodes the input into a n_{u_k}-dimensional Gaussian distribution, $\boldsymbol{\mu}_{\mathcal{K}}^{\mathrm{d}}, \boldsymbol{\Sigma}_{\mathcal{K}}^{\mathrm{d}}$ are the outputs of a DNN $f_{\phi_{\mathcal{K}}}$ with inputs $\mathbf{U}_1, \ldots, \mathbf{U}_K$, sampled from $P_{\theta_k}(\mathbf{u}_k|\mathbf{x}_k)$ and $\boldsymbol{\mu}_k^{\mathrm{d}}, \boldsymbol{\Sigma}_k^{\mathrm{d}}$ are the output of a DNN f_{ϕ_k} with input $\mathbf{U}_k$, $k = 1, \ldots, K$.
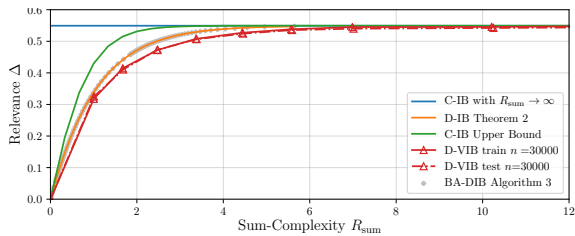
Fig. 3: Relevance vs. sum-complexity tradeoff for vector Gaussian data model with $K = 2$ encoders, $n_y = 1$, $n_1 = n_2 = 3$, and achievable pairs with the BA-DIB and D-VIB algorithms for $n = 30.000$.
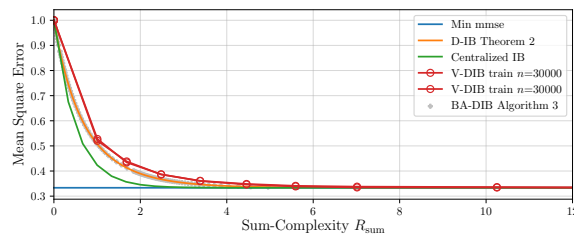
Fig. 4: Mean square error vs. sum-complexity tradeoff for vector Gaussian data model with $K = 2$ encoders, $n_y = 1$, $n_1 = n_2 = 3$, and achievable pairs with the BA-DIB and D-VIB algorithms for $n = 30.000$.

Remark 11. *The above parametrization might result in some performance loss, e.g., observe that $Q_{\varphi_k}(\mathbf{u}_k)$ may not be expressive enough to approximate the optimal distribution in* (21). *Also, note that the distributions that optimize the variational DIB cost* (23) *do not necessarily minimize the empirical DIB cost* (36).

VI. EXPERIMENTAL RESULTS

In this section, we evaluate the relevance-complexity tradeoffs achieved by the proposed algorithms BA-DIB and D-VIB in experiments with synthetic and real data. We also compare the resulting relevance-complexity pairs to the optimal relevance-complexity tradeoff and to an upper bound on the relevance-complexity, which we denote by *Centralized IB upper bound (C-IB)*. The C-IB bound is given by the pairs (Δ, R_{sum}) achievable if $(X_1, \ldots, X_K)$ are encoded jointly at a single encoder with complexity $R_{\text{sum}} = R_1 + \cdots + R_K$ and are given by the centralized IB problem [18]:

$$\Delta_{\text{cIB}}(R_{\text{sum}}) = \max_{P_{U|X_1, \ldots, X_K}} I(Y; U)$$

$$\text{s.t. } R_{\text{sum}} \geq I(X_1, \ldots, X_K; U). \tag{41}$$

Our implementation of the D-VIB in Algorithm 3 uses Tensorflow and its implementation of Adam optimizer [72] over 150 epochs and minibatch size of 64. The learning rate is computed as $0.001 \cdot (0.5)^{\lfloor n_{\text{epoch}}/30 \rfloor}$ at epoch n_{epoch}.

A. Regression for Vector Gaussian Data Model

In this section, we consider a real valued vector Gaussian data model as in Section III-C with $K = 2$ encoders observing a noisy version of an n_y-dimensional Gaussian vector $\mathbf{Y} \sim \mathcal{N}(\mathbf{y}; \mathbf{0}, \mathbf{I})$, as $\mathbf{X}_k = \mathbf{H}_k \mathbf{Y} + \mathbf{N}_k$, where $\mathbf{H}_k \in \mathbb{R}^{n_k \times n_y}$ and the noise is distributed as $\mathbf{N}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $k = 1, 2$.

The optimal complexity-relevance tradeoff for this model is characterized as in Theorem 2 and can be computed by solving the convex problem (2). The C-IB upper bound in (41) is an instance the IB problem for vector Gaussian sources in [30] and it can be computed analytically.

Next, we consider the numerical evaluation of the proposed algorithms BA-DIB and D-VIB for regression of the Gaussian target variable $\mathbf{Y}$ trained using a dataset of n i.i.d. samples $\{(\mathbf{X}_{1,i}, \mathbf{X}_{2,i}, \mathbf{Y}_i)\}_{i=1}^n$ form the described
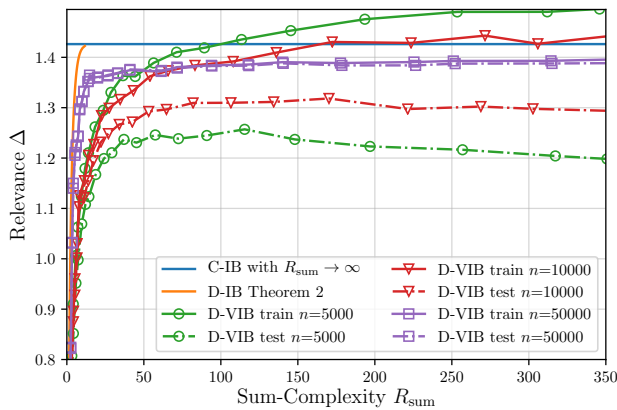
Fig. 5: Relevance vs. sum-complexity tradeoff for vector Gaussian data model with $K = 2$ encoders, $n_y = 2$, $n_1 = n_2 = 3$, and achievable pairs with D-VIB for training dataset $n = \{5.000, 10.000, 50.000\}$.

TABLE I: DNN architecture for Figure 3 and Figure 5

	DNN Layers
Encoder k	dense [512]-ReLu
	dense [512]-ReLu
	dense [512]-ReLu
Lat. space k	dense [256]-ReLu
Decoder 12	dense [256]-ReLu
Decoder k	dense [256]-ReLu

vector Gaussian data model. For BA-DIB, we assume a jointly Gaussian distribution for the data and empirically estimate its joint mean and covariance. Then, we apply Algorithm 2 to compute the relevance-complexity pairs and the corresponding estimators for different values of s. For D-VIB, we do not make any assumption on the data model and we apply Algorithm 3 to train the DNNs determining the encoders and decoders for different values of s. We use the multivariate Gaussian parameterization in (37)-(40) for the DNNs architecture shown in Table I. Specifically, Encoder k, $k = 1, 2$, consists of three dense layers of 512 neurons each followed by rectified linear unit (ReLu) activations. The output of encoder k is processed by a dense layer without nonlinear activation to generate $\boldsymbol{\mu}_k^{\mathrm{e}}$ and $\boldsymbol{\Sigma}_k^{\mathrm{e}}$ of size 512 and 512×512, respectively. Each decoder consists of two dense layers of 512 neurons with ReLu activations. The output of decoder 1, 2 and 12 is processed, each, by a fully connected layer without activation to generate $\boldsymbol{\mu}_k^{\mathrm{d}}$ and $\boldsymbol{\Sigma}_k^{\mathrm{d}}$ and $\boldsymbol{\mu}_{12}^{\mathrm{d}}$ and $\boldsymbol{\Sigma}_{12}^{\mathrm{d}}$, of size 2 and 2×2.

Figure 3 shows the optimal relevance-complexity region of tuples $(\Delta, R_{\mathrm{sum}})$ obtained from Theorem 2 for a vector Gaussian model with $K = 2$ encoders, target variable dimension $n_y = 1$, and observations dimension $n_1 = n_2 = 3$, as well as the C-IB bounds $\Delta_{\mathrm{cIB}}(R_{\mathrm{sum}})$ and $\Delta_{\mathrm{cIB}}(\infty)$. A dataset of 40.000 i.i.d. samples is available which is split into a training set of $n = 30.000$ samples and a test set of 10.000 samples used to evaluate the results from the trained estimators. We show the tuples $(\Delta, R_{\mathrm{sum}})$ resulting from the application of the BA-DIB in Algorithm 2 for different values of $s \in (0, 10]$, and we observe that they lie on the optimal relevance-complexity curve obtained from Theorem 2. The relevance-complexity pairs resulting form the application of the D-VIB algorithm for different values of s in the range $(0, 10]$ calculated as in Proposition 2 are also shown. Figure 4 depicts the mean squared error (MSE) between the original vector $\mathbf{Y}$ and the estimation given by the application of the BA-DIB and the D-VIB algorithm as in Figure 3 .

Figure 5 shows the effect of the size of the training set on generalization of the DNNs trained with the D-VIB in Algorithm 3 for a vector Gaussian model with $K = 2$ encoders, target variable dimension $n_y = 2$, and observations dimension $n_1 = n_2 = 3$, and the same DNN architecture in Table I. The achievable relevance-complexity pairs are shown after training the DNNs with training set lengths of $n = \{5.000, 10.000, 50.000\}$ as well as the corresponding achievable pairs for the test data, and the optimal pairs from Theorem 2. It is observed that for the training data,
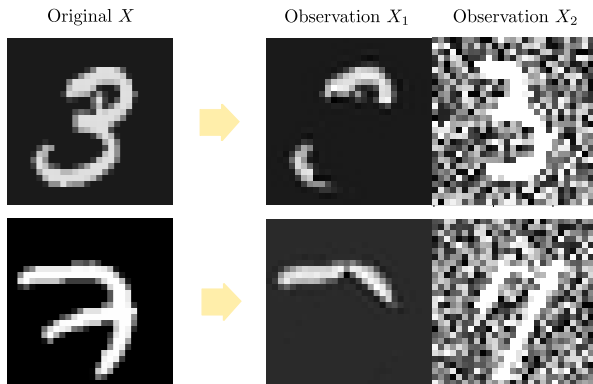
Fig. 6: Two-view handwritten MNIST dataset.

TABLE II: DNN architecture for Figure 7 and Figure8

	DNN Layers
Encoder k	conv. ker. [5,5,32]-ReLu
	maxpool [2,2,2]
	conv. ker. [5,5,64]-ReLu
	maxpool [2,2,2]
	dense [1024]-ReLu
	dropout 0.4
	dense [256]-ReLu
Latent space k	dense [256]-ReLu
Decoder 12	dense [256]-ReLu
Decoder k	dense [256]-ReLu

the larger the complexity, the higher the achievable relevance. Note that for small and intermediate datasets, e.g., with $n = \{5.000, 10.000\}$, there is overfitting of the DNN to the dataset, and the resulting relevance is higher than that allowed by the optimal tradeoff. However, overfitting results in performance loss for test data, which is more pronounced for larger complexity values. Note that the complexity constraint acts as a regularizer which improves generalization of the estimators and that the largest relevance during test is achieved for small complexity values. As the dataset size increases, the difference between resulting relevance-complexity pairs for the test and train datasets is reduced and the better the DNN generalizes. Note the gap to the optimal tradeoff, which could be reduced with more expressive parameterizations than that in (37)-(40).

B. Classification on the multi-view MNIST Dataset

xIn this section, we consider the evaluation of the proposed algorithms for classification on a two-view version of the the MNIST dataset, consisting of 70.000 labeled images of handwritten digits between $\{0, \ldots, 9\}$ observed under two different views, from which the corresponding label $Y \in \{0, \ldots, 9\}$ has to be inferred. Each MNIST image consists of 28×28 grayscale pixels. In this experiment the two views are generated as follows. View 1 is generated by occluding the image with square of 25×25 pixels, randomly rotated for each image with angles uniformly selected in the range $[-\pi/4, \pi/4]$. To generate view 2, for each image in view 1, we randomly select an image of the same digit $(0-9)$ from the original MNIST dataset and add independent random noise uniformly sampled from $[0, 3]$ to each pixel, and the pixel values are truncated to $[0, 1]$. Encoder 1 observes view 1 and Encoder 2 observes view 2, and treat each view as normalized vectors $\mathbf{X}_k \in [0, 1]^{784}$, $k = 1, 2$. We randomly split the 70.000 two-view samples $\{\mathbf{X}_{1,i}, \mathbf{X}_{2,i}, Y_i\}$ into training and test sets of length n and $70.000 - n$, respectively. A subset of the two-view MNIST dataset is shown in Figure 6. To asses the difficulty of estimating the digit from each of the views independently, we consider a standard convolutional neural network (CNN) architecture with dropout which achieves an accuracy of 99.8% in the noiseless MNIST dataset. Then, we independently retrain the same CNN architecture for each encoder input, achieving an accuracy of 92.3% for view 1 and 79.68% for view 2, as shown in Table III. That is, view 1 is less noisy.

We consider the application of the D-VIB algorithm to this model with the CNN architecture in Table II, in which Encoder k, $k = 1, 2$ is parameterized by a $n_{u_k} = 256$ dimensional multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_k^{\mathrm{e}}, \boldsymbol{\Sigma}_k^{\mathrm{e}})$
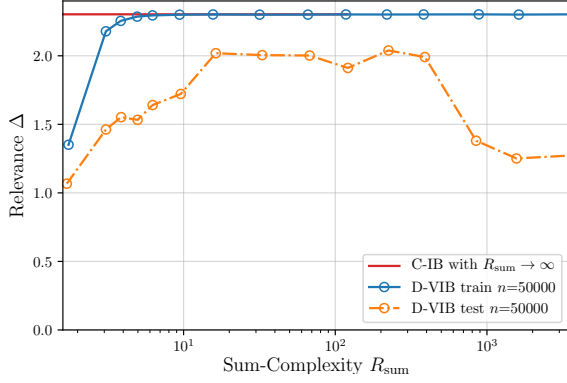
Fig. 7: Relevance vs. sum-complexity tradeoff for the two-view MNIST dataset with $K = 2$ encoders, with the D-VIB algorithm for training dataset $n = 50.000$ and $s \in [10^{-10}, 1]$.
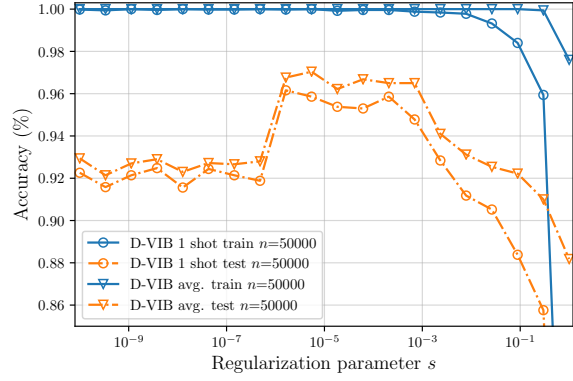

Fig. 8: Train and test accuracy for the two-view MNIST dataset with $K = 2$ encoders, from the estimators resulting with the D-VIB algorithm and the D-VIB avg. for $n = 50.000$ and $s \in [10^{-10}, 1]$.
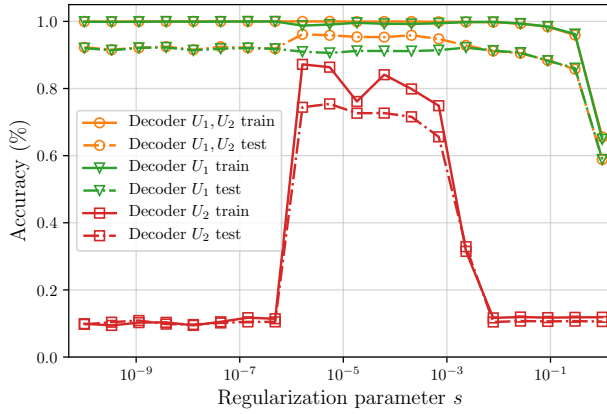

Fig. 9: Train and test accuracy for the two-view MNIST dataset with $K = 2$ encoders, from the estimators using both (U_1, U_2), $Q_{Y|U_1,U_2}$, and only U_1 or U_2, i.e., $Q_{Y|U_1}$ and $Q_{Y|U_2}$, from the application of the D-VIB algorithm for $n = 50.000$ and $s \in [10^{-10}, 1]$.

TABLE III: Accuracy for different algorithms with CNN architectures.

	Accuracy (%)	
	1 shot	avg.
D-VIB	**96.16**	**97.24**
D-VIB-noReg	96.04	96.72
C-VIB	96.01	96.68
Deterministic	93.18	93.18
Independent	92.1 / 79.68	93.1 / 82.01

determined by the output of a DNN f_{θ_k} consisting of the concatenation of convolutional, dense and maxpool layers with ReLu activations and dropout. The output of the last layer is followed by a dense layer without activation that generate $\boldsymbol{\mu}_k^{\mathrm{e}}$ and $\boldsymbol{\Sigma}_k^{\mathrm{e}}$. The prior is chosen as $Q_{\varphi_k}(\mathbf{u}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. Each decoder takes the samples from $P_{\theta_k}(U_k|X_k)$ and processes its inputs with a dense layer DNN ($f_{\phi_{\mathcal{K}}}$ and f_{ϕ_k}) each with 256 neurons and ReLu activation, which outputs a vector $\hat{\mathbf{y}}_i$ of size $|\mathcal{Y}| = 10$ normalized with a softmax, corresponding to a distribution over the one-hot encoding of the digit labels $\{0, \ldots, 9\}$ from the K observations, i.e., we have:

$$Q_{\phi_k}(\hat{\mathbf{y}}_k|\mathbf{u}_k) = \mathrm{Softmax}(f_{\phi_k}(U_k)), \quad k = 1, 2, \text{ and} \tag{42}$$

$$Q_{\phi_{\mathcal{K}}}(\hat{\mathbf{y}}|\mathbf{u}_{\mathcal{K}}) = \mathrm{Softmax}(f_{\phi_{\mathcal{K}}}(U_1, U_2))), \tag{43}$$

where $\mathrm{Softmax}(\mathbf{p})$ for $\mathbf{p} \in \mathbb{R}^d$ is a vector with i-th entry $[\mathrm{Softmax}(\mathbf{p})]_i = \exp(p_i)/\sum_{j=1}^d \exp(p_j)$, $i = 1, \ldots, d$.

For this parameterization, the log-loss terms in the empirical DIB cost (35) reduce to the cross-entropy and the KL divergence terms, can be computed as in (92).

Figure 7 shows the resulting relevance-complexity pairs from the application of the D-VIB Algorithm 3 on

the two-view MNIST dataset for training dataset of size $n = 50.000$ and 15 different s regularization values in the range $[10^{-10}, 1]$ to train different estimators on the relevance-complexity plane. The C-IB limit is shown for $R_{\text{sum}} \to \infty$ assuming that zero classification error is possible form the data, given by $\Delta_{\text{cIB}}(R_{\text{sum}}) = \log 10$. It can be observed that during the training phase, the higher the sum-complexity, the higher the achieved relevance, and that the resulting relevant-complexity pairs perform very close to the theoretical limit. On the contrary, during the test phase, while low sum-complexity results in low achievable relevance, and it increases for intermediate sum-complexity values, the achievable relevance decreases for large values of sum-complexity. Thus, the effect of the regularization due to the complexity constraint results in higher generalization.

After training, the trained CNNs for different values of s can be used for classification. In particular we estimate the labels from the maximum of the estimated conditional distribution $\hat{\mathbf{y}}_i$ for the given observation. We study the case in which a single Monte Carlo sample of U_k is used for prediction, denoted by D-VIB 1 shot, and the case, which we denote by D-VIB avg., in which U_k are sampled M times to generate M estimations $\hat{\mathbf{y}}_{i,m}$, which are averaged to estimate the conditional probability $\hat{\mathbf{y}}_{\text{av}} = \frac{1}{M} \sum_{m=1}^{M} \hat{\mathbf{y}}_{i,m}$, from which the label can be inferred.

Figure 8 shows the resulting accuracy achievable with D-VIB 1 shot and D-VIB avg. using each estimator obtained in Figure 7 with respect to the regularization parameter s. It can be observed that higher accuracy for both methods is obtained at the intermediate regularization parameter values, $s \simeq 10^{-6}$, for which relevance is maximized in Figure 7. Indeed, for any estimator $Q_{Y|U_{\mathcal{K}}}$, the average logarithmic-loss provides an upper bound on the classification error, since by the application of Jensen's inequality, we have

$$P_{\text{error}}(Q_{Y|U_{\mathcal{K}}}) := 1 - \mathrm{E}_{P_{X_{\mathcal{K}},Y}}[Q_{Y|U}] \tag{44}$$

$$\leq 1 - \exp\left(-\mathrm{E}_{P_{X_{\mathcal{K}},Y}}[-\log Q_{Y|U_{\mathcal{K}}}]\right), \tag{45}$$

which justifies the logarithmic loss as a useful surrogate of the probability of error. Note that the largest gains from the averaging with D-VIB avg. occur for the values of s for which the relevance is maximized for the test data.

During the optimization, indirectly, the D-VIB algorithm allocates the complexity for each description U_k of the observation X_k. Figure 9 shows the resulting accuracy achieved with D-VIB 1 shot to estimate Y with the main estimator $Q_{Y|U_1,U_2}$ obtained in Figure 2, as well as the accuracy, achieved by using the regularizing decoders $Q_{Y|U_k}$, $k = 1, 2$ to estimate Y, with respect to the regularization parameter s. In general, the description U_1 from view X_1 (which is less noisy), carries most of the information, while for $10^{-6} \lesssim s \lesssim 10^{-3}$, both descriptions U_1 and U_2 capture relevant information from the views. In this regime, the combination of both views result in an increase of the overall performance for $Q_{Y|U_1,U_2}$.

In order to asses the advantages of the D-VIB trained with Algorithm 3, and the relevance of the logarithmic loss terms in the regularization in the variational DIB cost in (20), we compare the best accuracy achievable by D-VIB in Figure 8, i.e., selecting the estimator corresponding to s for which the accuracy is maximized, with three more estimators for the setup in Figure 1 using CNNs. First, we consider the centralized VIB algorithm, e.g., as in [32], denoted by C-VIB, in which all observations are processed jointly. C-VIB can be trained with the DIB cost in (35) with $K = 1$ and letting $X = (X_1, X_2)$. We also consider the training of an estimator without the regularizing
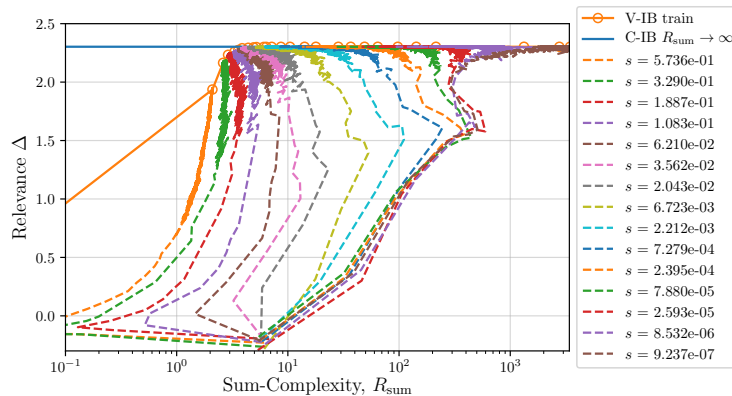
Fig. 10: Evolution of the relevance-complexity pairs during training with D-VIB for different values of the regularization parameter s over the relevance-complexity plane.

decoders $Q_{Y|U_k}$ $k = 1, 2$ and by maximizing the empirical DIB cost in (35) with only the divergence terms in the regularizer using Algorithm 3. The latter corresponds to the naive direct extension of the ELBO bound and the C-VIB to the distributed case, and we denote this scheme by D-VIB-noReg. In particular, for C-VIB, we assume that a single encoder maps the two views $\mathbf{X}_1, \mathbf{X}_2$ to a Gaussian latent of dimension 256 using a CNN network of similar architecture as the encoders in Table II, and uses a singe decoder with 256 activations. Finally, we consider a standard deterministic CNN with dropout regularization, using an architecture as that of the D-VIB-noReg, i.e., without regularizing decoders, and with a deterministic mapping for the latent space layer. For estimation we also consider the single shot and averaging of multiple estimations of the estimated conditional distribution as in D-VIB avg. Table III shows the advantages of D-VIB over other approaches. The advantage of D-VIB over C-VIB is explained by the advantage of training the latent space embedding for each observation separately, which allows to adjust better the CNN encoding-decoding parameters to the statistics of each observation, justifying the use of D-VIB for multi-view learning even if the data is available in a centralized manner.

C. Evolution of the Relevance-Complexity pairs

In this section, we consider the evolution of the achievable relevance-complexity pairs over the relevance-complexity plane during the training phase when using the D-VIB as given in Algorithm 3. We are interested in the trajectory followed by the pairs for different regularization parameter value s as the DNN parameters are updated from a random initialization until their convergence. In particular, for a given value s, let $(\boldsymbol{\theta}^t, \boldsymbol{\phi}^t, \boldsymbol{\varphi}^t)$ denote the DNNs parameters determining encoders $P_{\theta_k^t}(U_k|X_k)$, decoders $Q_{\phi_{\mathcal{K}}^t}(Y|U_{\mathcal{K}})$ and $Q_{\phi_k^t}(Y|U_k)$, and priors $Q_{\varphi_k^t}$ at iteration t of Algorithm 3. The achievable relevance pair (Δ_s^t, R_s^t) on the (Δ, R_{sum}) plane at iteration t with such estimator can be calculated by evaluating (17)-(18) for the current DNN parameters.

For simplicity, we consider the standard MNIST database with 70.000 labeled images of handwritten digits between $\{0, \ldots, 9\}$, observed under a single view consisting of a single view of 28×28 grayscale pixels per sample, i.e., $K = 1$. We consider a parameterization as in Section VI-B with an encoder with two layers of 1024 neurons each with ReLu activations, whose output is mapped to generate the mean and covariance of a multivariate Gaussian embedding of size 256. The decoder consists of a single layer of 256 activations which generate a

conditional probability of size 10. The DNN parameters are initialized with Glorot's initialization [73] and updated with Algorithm 3 until convergence for different values of the regularization parameter s.

Figure 10 shows the evolution of the pairs (Δ_s^t, R_s^t) for different values of the regularization parameter s during training starting from the achievable pairs with randomly initialized weights around $(\Delta_s^{t_0}, R_s^{t_0}) \simeq (-0.2, 0.5)$ until their converge to the corresponding achievable point on the V-IB training curve. Three different type of trajectories can be observed depending on the value of s: i) For $s \in [5.736e - 01, 2.043e - 02)$ (curves on left), the complexity decreases without significant improvement in the relevance. Then relevance improves at the cost of higher complexity and starts reducing complexity just before convergence on the point on the V-IB train curve; ii) for $s \in [2.043e - 02, 2.594e - 05)$ the algorithm directly starts improving relevance and starts reducing complexity before convergence; iii) For $s \leq 2.594e - 05$ (curves on right) relevance improves by increasing complexity, then complexity is reduced while improving relevance, and finally they continue improving relevance by increasing complexity until convergence.

The evolution of the relevance-complexity pairs at different layers of a DNN has been studied in [37] and [74], where relevance is measured as the mutual information between Y and the output at layer l, and complexity is given by the mutual information between the observed data and the output. Under tanh activations, only curves of type ii) are observed in [37] while a different type of trajectory has been observed under ReLu activations in [74]. Figure 10 shows that very different types of trajectories can result for the same DNN architecture, i.e., weights and activations, depending on the regularization parameter. Figure 10 shows a phase transition between the three types of trajectories on the plane. The study of the trajectories followed by the algorithm for different activations is left as future work.

VII. Proofs

A. Proof of Theorem 1

In this proof, we show the connection between the distributed learning problem under study and the K-encoder CEO problem studied in [59]. For the K-encoder CEO problem, let us consider K encoding functions $\phi_k : \mathcal{X}_k \to \mathcal{M}_k^{(n)}$ satisfying the complexity constraint (7) and a decoding function $\tilde{\psi} : \mathcal{M}_1^{(n)} \times \ldots \times \mathcal{M}_K^{(n)} \to \hat{\mathcal{Y}}^n$, that produces a probabilistic estimate of Y from the outputs of the encoders, i.e., $\hat{\mathcal{Y}}^n$ is the set of distributions on $\mathcal{Y}$. The quality of the estimation is measured in terms of the average log-loss.

Definition 2. *A tuple $(D, R_1, \ldots, R_K)$ is said to be achievable in the K-encoder CEO problem for $P_{X_{\mathcal{K}}, Y}$ for which the Markov chain (3) holds, if there exists a length n, encoders ϕ_k for $k \in \mathcal{K}$ and a decoder $\tilde{\psi}$, such that*

$$D \geq \mathrm{E}\left[\frac{1}{n} \log \frac{1}{\hat{P}_{Y^n | J_{\mathcal{K}}}(Y^n | \phi_1(X_1^n), \ldots, \phi_K(X_K^n))}\right], \tag{46}$$

$$R_k \geq \frac{1}{n} \log |\phi_k(X_k^n)| \quad \text{for all} \ \ k \in \mathcal{K}. \tag{47}$$

The rate-distortion region $\mathcal{RD}_{\mathrm{CEO}}$ is given by the closure of all achievable tuples $(D, R_1, \ldots, R_K)$. □

The following lemma shows that the minimum average logarithmic loss is the conditional entropy of Y given the descriptions. The result is essentially equivalent to [59, Lemma 1] and it is provided for completeness.

Lemma 4. *Let us consider $P_{X_\mathcal{K}, Y}$ and the encoders $J_k = \phi_k(X_k^n)$, $k \in \mathcal{K}$ and the decoder $\hat{Y}^n = \tilde{\psi}(J_\mathcal{K})$. Then,*

$$\mathrm{E}[\ell_{\log}(Y^n, \hat{Y}^n)] \geq H(Y^n | J_\mathcal{K}), \tag{48}$$

with equality if and only if $\tilde{\psi}(J_\mathcal{K}) = \{P_{Y^n | J_\mathcal{K}}(y^n | J_\mathcal{K})\}_{y^n \in \mathcal{Y}^n}$.

Proof. Let $Z := (J_1, \ldots, J_K)$ be the argument of $\tilde{\psi}$ and $\hat{P}(y^n | z)$ be a distribution on $\mathcal{Y}^n$. We have for $Z = z$:

$$\mathrm{E}[\ell_{\log}(Y^n, \hat{Y}^n) | Z = z] = \sum_{y^n \in \mathcal{Y}^n} P(y^n | z) \log\left(\frac{1}{\hat{P}(y^n | z)}\right) \tag{49}$$

$$= \sum_{y^n \in \mathcal{Y}^n} P(y^n | z) \log\left(\frac{P(y^n | z)}{\hat{P}(y^n | z)}\right) + H(Y^n | Z = z) \tag{50}$$

$$= D_{\mathrm{KL}}(P(y^n | z) \| \hat{P}(y^n | z)) + H(Y^n | Z = z) \tag{51}$$

$$\geq H(Y^n | Z = z), \tag{52}$$

where (52) is due to the non-negativity of the KL divergence and the equality holds if and only if for $\hat{P}(y^n | z) = P(y^n | z)$ where $P(y^n | z) = \Pr\{Y^n = y^n | Z = z\}$ for all z and $y^n \in \mathcal{Y}^n$. Averaging over Z completes the proof. $\square$

Essentially, Lemma 4 states that minimizing the average log-loss is equivalent to maximizing relevance as given in (8). Formally, the connection between the distributed learning problem under study and the K-encoder CEO problem studied in [59] can be formulated as stated next.

Proposition 4. *A tuple $(\Delta, R_1, \ldots, R_K) \in \mathcal{RI}_{\mathrm{DIB}}$ if and only if $(H(Y) - \Delta, R_1, \ldots, R_K) \in \mathcal{RD}_{\mathrm{CEO}}$.*

Proof. Let the tuple $(\Delta, R_1, \ldots, R_K) \in \mathcal{RI}_{\mathrm{DIB}}$ be achievable for some encoders ϕ_k, i.e., (8) and (9) hold. It follows by Lemma 4 that by letting the decoding function $\tilde{\psi}(J_\mathcal{K}) = \{P_{Y^n | J_\mathcal{K}}(y^n | J_\mathcal{K})\}$, we have $\mathrm{E}[\ell_{\log}(Y^n, \hat{Y}^n) | J_\mathcal{K}] = H(Y^n | J_\mathcal{K})$, and hence $(H(Y) - \Delta, R_1, \ldots, R_K) \in \mathcal{RD}_{\mathrm{CEO}}$.

Conversely, assume the tuple $(D, R_1, \ldots, R_K) \in \mathcal{RD}_{\mathrm{CEO}}$ is achievable. It follows by Lemma 4 that $H(Y) - D \leq H(Y^n) - H(Y^n | J_\mathcal{K}) = I(Y^n; J_\mathcal{K})$, which implies $(\Delta, R_1, \ldots, R_K) \in \mathcal{RI}_{\mathrm{DIB}}$ with $\Delta = H(Y) - D$. $\square$

The characterization of rate-distortion region $\mathcal{R}_{\mathrm{CEO}}$ has been established recently in [59, Theorem 10]. The proof of the theorem is completed by noting that Proposition 4 implies that the result in [59, Theorem 10] can be applied to characterize the region $\mathcal{RI}_{\mathrm{DIB}}$ as given in Theorem 1. $\square$

B. Proof of Proposition 1

For simplicity of exposition, the proof is given for the case $K = 2$ encoders. The proof for $K > 2$ follows similarly. By the definition of $\mathcal{RI}_{\mathrm{DIB}}^{\mathrm{sum}}$, the tuple $(\Delta, R_{\mathrm{sum}}) \in \mathbb{R}_+^2$ is achievable for some random variables

Y, X_1, X_2, U_1, U_2 with joint pmf satisfying (15), if it holds that

$$\Delta \leq I(Y; U_1, U_2) \tag{53}$$

$$\Delta \leq R_1 - I(X_1; U_1|Y) + I(Y; U_2) \tag{54}$$

$$\Delta \leq R_2 - I(X_2; U_2|Y) + I(Y; U_1) \tag{55}$$

$$\Delta \leq R_1 + R_2 - I(X_1; U_1|Y) - I(X_2; U_2|Y) \tag{56}$$

$$R_1 + R_2 \leq R_{\text{sum}}. \tag{57}$$

The application of the Fourier-Motzkin elimination to project out R_1 and R_2 reduces the system on inequalities (53)-(57) to the following system of inequalities

$$\Delta \leq I(Y; U_1, U_2) \tag{58}$$

$$\Delta \leq R_{\text{sum}} - I(X_1; U_1|Y) - I(X_2; U_2|Y) \tag{59}$$

$$2\Delta \leq R_{\text{sum}} - I(X_1; U_1|Y) - I(X_2; U_2|Y) \tag{60}$$

$$+ I(Y; U_1) + I(Y; U_2) \tag{61}$$

It follows due to the Markov chain $U_1 \multimap X_1 \multimap Y \multimap X_2 \multimap U_2$ that we have $I(Y; U_1, U_2) \leq I(Y; U_1) + I(Y; U_2)$. Therefore, inequality (61) is redundant as it is implied by (58) and (59). This completes the proof. $\qquad\square$

C. Proof of Proposition 2

Suppose that $\mathbf{P}^*$ yields the maximum in (17). Then,

$$(1+s)\Delta_s = (1+sK)H(Y) + sR_s + \mathcal{L}_s(\mathbf{P}^*) \tag{62}$$

$$= (1+sK)H(Y) + sR_s \tag{63}$$

$$+ \left(-H(Y|U_{\mathcal{K}}^*) - s\sum_{k=1}^{K}[H(Y|U_k^*) + I(X_k; U_k^*)] \right) \tag{64}$$

$$= (1+sK)H(Y) + sR_s \tag{65}$$

$$+ (-H(Y|U_{\mathcal{K}}^*) - s(R_s - I(Y; U_{\mathcal{K}}^*) + KH(Y))) \tag{66}$$

$$= (1+s)I(Y; U_{\mathcal{K}}^*) \tag{67}$$

$$\leq (1+s)\Delta(R_s, P_{X_{\mathcal{K}}, Y}), \tag{68}$$

where (63) is due to the definition of $\mathcal{L}_s(\mathbf{P})$ in (19); (65) follows since we have $\sum_{k=1}^{K}[I(X_k; U_k^*) + H(Y|U_k^*)] = R_s - I(Y; U_{\mathcal{K}}^*) + KH(Y)$ from the definition of R_s in (18); and (68) follows from the definition in (16).

Conversely, if $\mathbf{P}^*$ is the solution to the maximization in the function $\Delta(R_{\text{sum}}, P_{X_{\mathcal{K}}, Y})$ in (16) such that $\Delta(R_{\text{sum}}, P_{X_{\mathcal{K}}, Y}) = \Delta_s$, then $\Delta_s \leq I(Y; U_{\mathcal{K}}^*)$ and $\Delta_s \leq R - \sum_{k=1}^{K} I(X_k; U_k^*|Y)$ and we have, for any $s \geq 0$, that

$$\Delta(R_{\text{sum}}, P_{X_{\mathcal{K}}, Y}) = \Delta_s$$

$$\leq \Delta_s - (\Delta_s - I(Y;U_{\mathcal{K}}^*)) - s\left(\Delta_s - R_{\text{sum}} + \sum_{k=1}^{K} I(X_k;U_k^*|Y)\right)$$

$$= I(Y;U_{\mathcal{K}}^*) - s\Delta_s + sR_{\text{sum}} - s\sum_{k=1}^{K} I(X_k;U_k^*|Y)$$

$$= H(Y) - s\Delta_s + sR_{\text{sum}} - H(Y|U_{\mathcal{K}}^*) - s\sum_{k=1}^{K}[I(X_k;U_k^*) + H(Y|U_k^*)] + sKH(Y) \quad (69)$$

$$\leq H(Y) - s\Delta_s + sR_{\text{sum}} + \mathcal{L}_s^* + sKH(Y) \quad (70)$$

$$= H(Y) - s\Delta_s + sR_{\text{sum}} + sKH(Y) - ((1+sK)H(Y) + sR_s - (1+s)\Delta_s) \quad (71)$$

$$= \Delta_s + s(R_{\text{sum}} - R_s), \quad (72)$$

where in (69) we have $\sum_{k=1}^{K} I(X_k;U_k|Y) = -KH(Y) + \sum_{k=1}^{K} I(X_k;U_k) + H(Y|U_k)$ due to the Markov chain $U_k - X_k - Y - (X_{\mathcal{K}\setminus k}, U_{\mathcal{K}\setminus k})$; (70) follows since $\mathcal{L}_s^*$ is the maximum over all possible distributions $\mathbf{P}$ (not necessarily $\mathbf{P}^*$ maximizing $\Delta(R_{\text{sum}}, P_{X_{\mathcal{K}},Y})$); and (71) is due to (17).

Finally, (72) is valid for any $R_{\text{sum}} \geq 0$ and $s \geq 0$. Given s, and hence (Δ_s, R_s), letting $R = R_s$ yields $\Delta(R_s, P_{X_{\mathcal{K}},Y}) \leq \Delta_s$. Together with (68), this completes the proof. $\square$

D. Proof of Lemma 1

The proof follows by deriving the following bounds. For any pmf $Q_{Y|Z}(y|z)$, $y \in \mathcal{Y}$ and $z \in \mathcal{Z}$, e.g., $\mathcal{Z} = \mathcal{U}_{\mathcal{K}}$ or $\mathcal{Z} = \mathcal{U}_k$, proceeding similarly to (52) and averaging over Z, we have

$$H(Y|Z) = \mathbb{E}[-\log Q_{Y|Z}(Y|Z)] - D_{\text{KL}}(P_{Y|Z}\|Q_{Y|Z}).$$

Similarly, we have

$$I(X_k;U_k) = H(U_k) - H(U_k|X_k) \quad (73)$$

$$= D_{\text{KL}}(P_{U_k|X_k}\|Q_{U_k}) - D_{\text{KL}}(P_{U_k}\|Q_{U_k}) \quad (74)$$

Applying these identities to (19), we have

$$\mathcal{L}_s(\mathbf{P}) = \mathcal{L}_s^{\text{VB}}(\mathbf{P},\mathbf{Q}) + D_{\text{KL}}(P_{Y|U_{\mathcal{K}}}\|Q_{Y|U_{\mathcal{K}}}) + s\sum_{k=1}^{K}(D_{\text{KL}}(P_{Y|U_k}\|Q_{Y|U_k}) + D_{\text{KL}}(P_{U_k}\|Q_{U_k}))$$

$$\geq \mathcal{L}_s^{\text{VB}}(\mathbf{P},\mathbf{Q}), \quad (75)$$

where (75) follows since the KL-divergence is always positive and equality is met iff $\mathbf{Q}^*$ is given as (21) and (22). $\square$

E. Proof of Theorem 2

The proof of Theorem 2 relies on deriving an outer bound on the relevance-complexity region described by (14), and showing that it is achievable with Gaussian pmfs and without time-sharing. In doing so, we use the technique of [75, Theorem 8] which relies on the de Bruijn identity and the properties of Fisher information and MMSE.

Lemma 5. *[75], [76] Let $(\mathbf{X}, \mathbf{Y})$ be a pair of random vectors with pmf $p(\mathbf{x}, \mathbf{y})$. We have*

$$\log|(\pi e)\mathbf{J}^{-1}(\mathbf{X}|\mathbf{Y})| \leq h(\mathbf{X}|\mathbf{Y}) \leq \log|(\pi e)\mathrm{mmse}(\mathbf{X}|\mathbf{Y})|,$$

where the conditional Fischer information matrix is defined as $\mathbf{J}(\mathbf{X}|\mathbf{Y}) := \mathrm{E}[\nabla \log p(\mathbf{X}|\mathbf{Y}) \nabla \log p(\mathbf{X}|\mathbf{Y})^\dagger]$, *and the minimum mean square error (MMSE) matrix is* $\mathrm{mmse}(\mathbf{X}|\mathbf{Y}) := \mathrm{E}[(\mathbf{X} - \mathrm{E}[\mathbf{X}|\mathbf{Y}])(\mathbf{X} - \mathrm{E}[\mathbf{X}|\mathbf{Y}])^\dagger]$.

First, we outer bound the relevance-complexity region in Theorem 1 for $(\mathbf{Y}, \mathbf{X}_\mathcal{K})$ as in (24). For $t \in \mathcal{T}$ and fixed $\prod_{k=1}^{K} p(\mathbf{u}_k|\mathbf{x}_k, t)$, choose $\mathbf{\Omega}_{k,t}$, $k = 1, \ldots, K$ satisfying $\mathbf{0} \preceq \mathbf{\Omega}_{k,t} \preceq \mathbf{\Sigma}_k^{-1}$ such that

$$\mathrm{mmse}(\mathbf{Y}_k|\mathbf{X}, \mathbf{U}_{k,t}, t) = \mathbf{\Sigma}_k - \mathbf{\Sigma}_k \mathbf{\Omega}_{k,t} \mathbf{\Sigma}_k. \tag{76}$$

Such $\mathbf{\Omega}_{k,t}$ always exists since $\mathbf{0} \preceq \mathrm{mmse}(\mathbf{X}_k|\mathbf{Y}, \mathbf{U}_{k,t}, t) \preceq \mathbf{\Sigma}_k^{-1}$, for all $t \in \mathcal{T}$, and $k \in \mathcal{K}$. We have from (14),

$$I(\mathbf{X}_k; \mathbf{U}_k|\mathbf{Y}, t) \geq \log|\mathbf{\Sigma}_k| - \log|\mathrm{mmse}(\mathbf{X}_k|\mathbf{Y}, \mathbf{U}_{k,t}, t)|$$

$$= -\log|\mathbf{I} - \mathbf{\Sigma}_k^{1/2} \mathbf{\Omega}_{k,t} \mathbf{\Sigma}_k^{1/2}|, \tag{77}$$

where the inequality is due to Lemma 5, and (77) is due to (76).

On the other hand, we have

$$I(\mathbf{Y}; \mathbf{U}_{\mathcal{S}^c, t}|t) \leq \log|\mathbf{\Sigma_y}| - \log|\mathbf{J}^{-1}(\mathbf{Y}|\mathbf{U}_{\mathcal{S}^c, t}, t)| \tag{78}$$

$$= \log\left|\sum_{k \in \mathcal{S}^c} \mathbf{\Sigma_y}^{1/2} \mathbf{H}_k^\dagger \mathbf{\Omega}_{k,t} \mathbf{H}_k \mathbf{\Sigma_y}^{1/2} + \mathbf{I}\right|, \tag{79}$$

where (78) is due to Lemma 5; and (79) is due to to the following equality connecting the MMSE matrix (76) and the Fisher information as in [43], [77], [78], proven below:

$$\mathbf{J}(\mathbf{Y}|\mathbf{U}_{\mathcal{S}^c, t}, t) = \sum_{k \in \mathcal{S}^c} \mathbf{H}_k^\dagger \mathbf{\Omega}_{k,t} \mathbf{H}_k + \mathbf{\Sigma_y}^{-1}. \tag{80}$$

In order to show (80), we use de Brujin identity to relate the Fisher information with the MMSE as given in the following lemma from [75].

Lemma 6. *Let $(\mathbf{V}_1, \mathbf{V}_2)$ be a random vector with finite second moments and $\mathbf{N} \sim \mathcal{CN}(\mathbf{0}, \mathbf{\Sigma}_N)$ independent of $(\mathbf{V}_1, \mathbf{V}_2)$. Then*

$$\mathrm{mmse}(\mathbf{V}_2|\mathbf{V}_1, \mathbf{V}_2 + \mathbf{N}) = \mathbf{\Sigma}_N - \mathbf{\Sigma}_N \mathbf{J}(\mathbf{V}_2 + \mathbf{N}|\mathbf{V}_1) \mathbf{\Sigma}_N.$$

From the MMSE of Gaussian random vectors [61],

$$\mathbf{Y} = \mathrm{E}[\mathbf{Y}|\mathbf{X}_{\mathcal{S}^c}] + \mathbf{Z}_{\mathcal{S}^c} = \sum_{k \in \mathcal{S}^c} \mathbf{G}_k \mathbf{X}_k + \mathbf{Z}_{\mathcal{S}^c}, \tag{81}$$

where $\mathbf{G}_k = \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}} \mathbf{H}_k^\dagger \mathbf{\Sigma}_k^{-1}$ and $\mathbf{Z}_{\mathcal{S}^c} \sim \mathcal{CN}(\mathbf{0}, \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}})$, and

$$\mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}}^{-1} = \mathbf{\Sigma}_{\mathbf{y}}^{-1} + \sum_{k \in \mathcal{S}^c} \mathbf{H}_k^\dagger \mathbf{\Sigma}_k^{-1} \mathbf{H}_k. \tag{82}$$

Note that $\mathbf{Z}_{\mathcal{S}^c}$ is independent of $\mathbf{Y}_{\mathcal{S}^c}$ due to the orthogonality principle of the MMSE and its Gaussian distribution. Hence, it is also independent of $\mathbf{U}_{\mathcal{S}^c,q}$. We have

$$\mathrm{mmse}\left(\sum_{k \in \mathcal{S}^c} \mathbf{G}_k \mathbf{X}_k \Big| \mathbf{Y}, \mathbf{U}_{\mathcal{S}^c,t}, t\right) = \sum_{k \in \mathcal{S}^c} \mathbf{G}_k \mathrm{mmse}\left(\mathbf{X}_k | \mathbf{Y}, \mathbf{U}_{\mathcal{S}^c,t}, t\right) \mathbf{G}_k^\dagger \tag{83}$$

$$= \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}} \sum_{k \in \mathcal{S}^c} \mathbf{H}_k^\dagger \left(\mathbf{\Sigma}_k^{-1} - \mathbf{\Omega}_k\right) \mathbf{H}_k \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}}, \tag{84}$$

where (83) follows since the cross terms are zero due to the Markov chain $(\mathbf{U}_{k,t}, \mathbf{X}_k) \,\multimap\, \mathbf{Y} \,\multimap\, (\mathbf{U}_{\mathcal{K}/k,t}, \mathbf{X}_{\mathcal{K}/k})$, see [77, Appendix V]; and (84) follows due to (76) and $\mathbf{G}_k$. Finally,

$$\mathbf{J}(\mathbf{Y}|\mathbf{U}_{\mathcal{S}^c,t}, t) = \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}}^{-1} - \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}}^{-1} \mathrm{mmse}\left(\sum_{k \in \mathcal{S}^c} \mathbf{G}_k \mathbf{X}_k \Big| \mathbf{Y}, \mathbf{U}_{\mathcal{S}^c,t}, t\right) \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}}^{-1} \tag{85}$$

$$= \mathbf{\Sigma}_{\mathbf{y}|\mathbf{x}_{\mathcal{S}^c}}^{-1} - \sum_{k \in \mathcal{S}^c} \mathbf{H}_k^\dagger \left(\mathbf{\Sigma}_k^{-1} - \mathbf{\Omega}_{k,t}\right) \mathbf{H}_k \tag{86}$$

$$= \mathbf{\Sigma}_{\mathbf{y}}^{-1} + \sum_{k \in \mathcal{S}^c} \mathbf{H}_k^\dagger \mathbf{\Omega}_{k,t} \mathbf{H}_k, \tag{87}$$

where (85) is due to Lemma 6; (86) is due to (84); and (87) follows due to (82).

Then, averaging over the time sharing T and letting $\bar{\mathbf{\Omega}}_k := \sum_{t \in \mathcal{T}} p(t) \mathbf{\Omega}_{k,t}$. Then, we have from (77)

$$I(\mathbf{X}_k; \mathbf{U}_k | \mathbf{Y}, T) \geq -\sum_{t \in \mathcal{T}} p(t) \log |\mathbf{I} - \mathbf{\Sigma}_k^{1/2} \mathbf{\Omega}_{k,t} \mathbf{\Sigma}_k^{1/2}|$$

$$\geq -\log |\mathbf{I} - \mathbf{\Sigma}_k^{1/2} \bar{\mathbf{\Omega}}_k \mathbf{\Sigma}_k^{1/2}|, \tag{88}$$

where (88) follows from the concavity of the log-det function and Jensen's inequality.

Similarly, from (79) and Jensen's Inequality we have

$$I(\mathbf{Y}; \mathbf{U}_{\mathcal{S}^c} | T) \leq \log \left| \sum_{k \in \mathcal{S}^c} \mathbf{\Sigma}_{\mathbf{y}}^{1/2} \mathbf{H}_k^\dagger \bar{\mathbf{\Omega}}_k \mathbf{H}_k \mathbf{\Sigma}_{\mathbf{y}}^{1/2} + \mathbf{I} \right|. \tag{89}$$

The outer bound on $\mathcal{RI}_{\mathrm{DIB}}$ is obtained by applying (88) and (89) in (14), noting that $\mathbf{\Omega}_k = \sum_{t \in \mathcal{T}} p(t) \mathbf{\Omega}_{k,t} \preceq \mathbf{\Sigma}_k^{-1}$ since $\mathbf{0} \preceq \mathbf{\Omega}_{k,t} \preceq \mathbf{\Sigma}_k^{-1}$, and taking the union over $\mathbf{\Omega}_k$ satisfying $\mathbf{0} \preceq \mathbf{\Omega}_k \preceq \mathbf{\Sigma}_k^{-1}$.

Finally, the proof is completed by noting that the outer bound is achieved with $T = \emptyset$ and multivariate Gaussian distributions $p^*(\mathbf{u}_k | \mathbf{x}_k, t) = \mathcal{CN}(\mathbf{x}_k, \mathbf{\Sigma}_k^{1/2}(\mathbf{\Omega}_k - \mathbf{I})\mathbf{\Sigma}_k^{1/2})$. $\qquad\square$

F. Derivation of Algorithm 2

In this section, we derive the update rules in Algorithm 2 and show that the Gaussian distribution is invariant to the update rules in Algorithm 1, in line with Theorem 2.

First, we recall that if $(\mathbf{X}_1, \mathbf{X}_2)$ are jointly Gaussian, then

$$P_{\mathbf{X}_2|\mathbf{X}_1=\mathbf{x}_1} = \mathcal{CN}(\boldsymbol{\mu}_{\mathbf{x}_2|\mathbf{x}_1}, \boldsymbol{\Sigma}_{\mathbf{x}_2|\mathbf{x}_1}), \tag{90}$$

where $\boldsymbol{\mu}_{\mathbf{x}_2|\mathbf{x}_1} := \mathbf{K}_{\mathbf{x}_2|\mathbf{x}_1}\mathbf{x}_1$, with $\mathbf{K}_{\mathbf{x}_2|\mathbf{x}_1} := \boldsymbol{\Sigma}_{\mathbf{x}_2,\mathbf{x}_1}\boldsymbol{\Sigma}_{\mathbf{x}_1}^{-1}$.

Then, for $\mathbf{Q}^{(t+1)}$ computed as in (21) and (22) from $\mathbf{P}^{(t)}$, which is a set of Gaussian distributions, we have

$$Q_{\mathbf{Y}|\mathbf{u}_k}^{(t+1)} = \mathcal{CN}(\boldsymbol{\mu}_{\mathbf{y}|\mathbf{u}_k^t}, \boldsymbol{\Sigma}_{\mathbf{y}|\mathbf{u}_k^t}),$$

$$Q_{\mathbf{Y}|\mathbf{u}_{\mathcal{K}}}^{(t+1)} = \mathcal{CN}(\boldsymbol{\mu}_{\mathbf{y}|\mathbf{u}_{\mathcal{K}}^t}, \boldsymbol{\Sigma}_{\mathbf{y}|\mathbf{u}_{\mathcal{K}}^t}).$$

Next, we look at the update $\mathbf{P}^{(t+1)}$ as in (27) from given $\mathbf{Q}^{(t+1)}$. First, we have that $p(\mathbf{u}_k^t)$ is the marginal of $\mathbf{U}_k^t$, given by $\mathbf{U}_k^t \sim \mathcal{CN}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{u}_k^t})$ where $\boldsymbol{\Sigma}_{\mathbf{u}_k^t} = \mathbf{A}_k^t \boldsymbol{\Sigma}_{\mathbf{x}_k}\mathbf{A}_k^{t,H} + \boldsymbol{\Sigma}_{\mathbf{z}_k^t}$.

Then, to compute $\psi_s(\mathbf{u}_k^t, \mathbf{x}_k)$, first, we note that

$$E_{U_{\mathcal{K}\backslash k}|x_k}[D_{\mathrm{KL}}(P_{Y|U_{\mathcal{K}\backslash k},x_k}||Q_{Y|U_{\mathcal{K}\backslash k},u_k})] = D_{\mathrm{KL}}(P_{Y,U_{\mathcal{K}\backslash k}|x_k}||Q_{Y,U_{\mathcal{K}\backslash k}|u_k}) - D_{\mathrm{KL}}(P_{U_{\mathcal{K}\backslash k}|x_k}||Q_{U_{\mathcal{K}\backslash k}|u_k}), \tag{91}$$

and that for two generic multivariate Gaussian distributions $P_1 \sim \mathcal{CN}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $P_2 \sim \mathcal{CN}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ in $\mathbb{C}^N$, the KL divergence is computed as

$$D_{\mathrm{KL}}(P_1||P_2) = \frac{1}{2}\left((\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T \boldsymbol{\Sigma}_2^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + \log|\boldsymbol{\Sigma}_2\boldsymbol{\Sigma}_1^{-1}| - d + \mathrm{tr}\{\boldsymbol{\Sigma}_2^{-1}\boldsymbol{\Sigma}_1\}\right). \tag{92}$$

Applying (91) and (92) in (28) and noting that all involved distributions are Gaussian, it follows that $\psi_s(\mathbf{u}_k^t, \mathbf{x}_k)$ is a quadratic form. Then, since $p(\mathbf{u}_k^t)$ is Gaussian, the product $\log(p(\mathbf{u}_k^t)\exp(-\psi_s(\mathbf{u}_k^t, \mathbf{x}_k)))$ is also a quadratic form, and identifying constant, first and second order terms, we have

$$\log p^{(t+1)}(\mathbf{u}_k|\mathbf{x}_k) = Z(\mathbf{x}_k) + (\mathbf{u}_k - \boldsymbol{\mu}_{\mathbf{u}_k^{t+1}|\mathbf{x}_k})^H \boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}}^{-1}(\mathbf{u}_k - \boldsymbol{\mu}_{\mathbf{u}_k^{t+1}|\mathbf{x}_k}), \tag{93}$$

where $Z(\mathbf{x}_k)$ is a normalization term independent of $\mathbf{u}_k$, and

$$\boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}}^{-1} = \boldsymbol{\Sigma}_{\mathbf{u}_k^t}^{-1} + \mathbf{K}_{\mathbf{y}|\mathbf{u}_k^t}^H \boldsymbol{\Sigma}_{\mathbf{y}|\mathbf{u}_k^t}^{-1}\mathbf{K}_{\mathbf{y}|\mathbf{u}_k^t} + \frac{1}{s}\mathbf{K}_{\mathbf{yu}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^H \boldsymbol{\Sigma}_{\mathbf{yu}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^{-1}\mathbf{K}_{\mathbf{yu}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t} - \frac{1}{s}\mathbf{K}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^H \boldsymbol{\Sigma}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^{-1}\mathbf{K}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}, \tag{94}$$

$$\boldsymbol{\mu}_{\mathbf{u}_k^{t+1}|\mathbf{x}_k} = \boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}}\left(\mathbf{K}_{\mathbf{y}|\mathbf{u}_k^t}^H \boldsymbol{\Sigma}_{\mathbf{y}|\mathbf{u}_k^t}^{-1}\boldsymbol{\mu}_{\mathbf{y}|\mathbf{x}_k} + \frac{1}{s}\mathbf{K}_{\mathbf{y},\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^H \boldsymbol{\Sigma}_{\mathbf{y},\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^{-1}\boldsymbol{\mu}_{\mathbf{y},\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{x}_k} - \frac{1}{s}\mathbf{K}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}\boldsymbol{\Sigma}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^{-1}\boldsymbol{\mu}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{x}_k}\right). \tag{95}$$

This shows that $p^{(t+1)}(\mathbf{u}_k|\mathbf{x}_k)$ is a multivariate Gaussian distribution and that $\mathbf{U}_k^{t+1}|\{\mathbf{X}_k = \mathbf{x}_k\}$ is also a multivariate Gaussian distributed as $\mathcal{CN}(\boldsymbol{\mu}_{\mathbf{u}_k^{t+1}|\mathbf{x}_k}, \boldsymbol{\Sigma}_{\mathbf{z}_k^{t+1}})$.

Next, we simplify (94) and (95) to obtain the update rules (30) and (31). From the matrix inversion lemma, similarly to [30], for $(\mathbf{X}_1, \mathbf{X}_2)$ jointly Gaussian we have

$$\boldsymbol{\Sigma}_{\mathbf{x}_2|\mathbf{x}_1}^{-1} = \boldsymbol{\Sigma}_{\mathbf{x}_2}^{-1} + \mathbf{K}_{\mathbf{x}_1|\mathbf{x}_2}^H \boldsymbol{\Sigma}_{\mathbf{x}_1|\mathbf{x}_2}^{-1}\mathbf{K}_{\mathbf{x}_1|\mathbf{x}_2}. \tag{96}$$

Applying (96), in (94) we have

$$\mathbf{\Sigma}_{\mathbf{z}_k^{t+1}}^{-1} = \mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1} + \frac{1}{s}\mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t}^{-1} - \frac{1}{s}\mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{u}_{\mathcal{K}\backslash k}^t}^{-1}, \tag{97}$$

$$= \left(1 + \frac{1}{s}\right)\mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1} - \frac{1}{s}\mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{u}_{\mathcal{K}\backslash k}^t}^{-1}, \tag{98}$$

where (98) is due to the Markov chain $\mathbf{U}_k \; \multimap \; \mathbf{Y} \; \multimap \; \mathbf{U}_{\mathcal{K}\backslash k}$.

Then, also from the matrix inversion lemma, we have for jointly Gaussian $(\mathbf{X}_1, \mathbf{X}_2)$,

$$\mathbf{\Sigma}_{\mathbf{x}_2|\mathbf{x}_1}^{-1}\mathbf{\Sigma}_{\mathbf{x}_2,\mathbf{x}_1}\mathbf{\Sigma}_{\mathbf{x}_1}^{-1} = \mathbf{\Sigma}_{\mathbf{x}_2}^{-1}\mathbf{\Sigma}_{\mathbf{x}_2,\mathbf{x}_1}\mathbf{\Sigma}_{\mathbf{x}_1|\mathbf{x}_2}^{-1}. \tag{99}$$

Applying (99) to (95), for the first term in (95), we have

$$\mathbf{K}_{\mathbf{y}|\mathbf{u}_k^t}^H \mathbf{\Sigma}_{\mathbf{y}|\mathbf{u}_k^t}^{-1} \boldsymbol{\mu}_{\mathbf{y}|\mathbf{x}_k} = \mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1}\mathbf{\Sigma}_{\mathbf{y},\mathbf{u}_k^t}\mathbf{\Sigma}_{\mathbf{y}}^{-1}\boldsymbol{\mu}_{\mathbf{y}|\mathbf{x}_k} \tag{100}$$

$$= \mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1}\mathbf{A}_k^t\mathbf{\Sigma}_{\mathbf{x}_k,\mathbf{y}}\mathbf{\Sigma}_{\mathbf{y}}^{-1}\mathbf{\Sigma}_{\mathbf{y},\mathbf{x}_k}\mathbf{\Sigma}_{\mathbf{x}_k}^{-1}\mathbf{x}_k$$

$$= \mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1}\mathbf{A}_k^t(\mathbf{I} - \mathbf{\Sigma}_{\mathbf{x}_k|\mathbf{y}}\mathbf{\Sigma}_{\mathbf{x}_k}^{-1})\mathbf{x}_k, \tag{101}$$

where $\mathbf{\Sigma}_{\mathbf{y},\mathbf{u}_k^t} = \mathbf{A}_k^t\mathbf{\Sigma}_{\mathbf{x}_k,\mathbf{y}}$; and (101) is due to the definition of $\mathbf{\Sigma}_{\mathbf{x}_k|\mathbf{y}}$. Similarly, for the second term in (95), we have

$$\mathbf{K}_{\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t} \mathbf{\Sigma}_{\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^{-1} \boldsymbol{\mu}_{\mathbf{y},\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{x}_k} = \mathbf{\Sigma}_{\mathbf{u}_k^t}^{-1}\mathbf{\Sigma}_{\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t,\mathbf{u}_k}\mathbf{\Sigma}_{\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^{-1}\boldsymbol{\mu}_{\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{x}_k} \tag{102}$$

$$= \mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t}^{-1}\mathbf{A}_k^t(\mathbf{I} - \mathbf{\Sigma}_{\mathbf{x}_k|\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t}\mathbf{\Sigma}_{\mathbf{x}_k}^{-1})\mathbf{x}_k, \tag{103}$$

$$= \mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}^{-1}\mathbf{A}_k^t(\mathbf{I} - \mathbf{\Sigma}_{\mathbf{x}_k|\mathbf{y}}\mathbf{\Sigma}_{\mathbf{x}_k}^{-1})\mathbf{x}_k, \tag{104}$$

where we use $\mathbf{\Sigma}_{\mathbf{u}_k^t,\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t} = \mathbf{A}_k^t\mathbf{\Sigma}_{\mathbf{x}_k,\mathbf{y}\mathbf{u}_{\mathcal{K}\backslash k}^t}$; and (104) is due to the Markov chain $\mathbf{U}_k \; \multimap \; \mathbf{Y} \; \multimap \; \mathbf{U}_{\mathcal{K}\backslash k}$.

For the third term in (95),

$$\mathbf{K}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t} \mathbf{\Sigma}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{u}_k^t}^{-1} \boldsymbol{\mu}_{\mathbf{u}_{\mathcal{K}\backslash k}^t|\mathbf{x}_k} = \mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{u}_{\mathcal{K}\backslash k}^t}^{-1}\mathbf{A}_k^t(\mathbf{I} - \mathbf{\Sigma}_{\mathbf{x}_k|\mathbf{u}_{\mathcal{K}\backslash k}^t}\mathbf{\Sigma}_{\mathbf{x}_k}^{-1})\mathbf{x}_k. \tag{105}$$

Equation (31) follows by noting that $\boldsymbol{\mu}_{\mathbf{u}_k^{t+1}|\mathbf{x}_k} = \mathbf{A}_k^{t+1}\mathbf{x}_k$, and that from (95) $\mathbf{A}_k^{t+1}$ can be identified as in (31).

Finally, note that due to (29), $\mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{y}}$ and $\mathbf{\Sigma}_{\mathbf{u}_k^t|\mathbf{u}_{\mathcal{K}\backslash k}}$ are given as in (32) and (33), where $\mathbf{\Sigma}_{\mathbf{x}_k|\mathbf{y}} = \mathbf{\Sigma}_k$ and $\mathbf{\Sigma}_{\mathbf{x}_k|\mathbf{u}_{\mathcal{K}\backslash k}^t}$ follows from its definition. This completes the proof. $\qquad\square$

REFERENCES

[1] I. E. Aguerri and A. Zaidi, "Distributed information bottleneck method for discrete and Gaussian sources," in *Proc. IEEE International Zurich Seminar on Communications*, Zurich, Switzerland, Feb. 2018, pp. 154–157.

[2] D. Erdogmus, "Information theoretic learning: Renyi's entropy and its applications to adaptive system training," Ph.D. dissertation, University of Florida Gainesville, Florida, 2002.

[3] J. C. Principe, N. R. Euliano, and W. C. Lefebvre, *Neural and adaptive systems: fundamentals through simulations*. Wiley New York, 2000, vol. 672.

[4] J. Jiao, T. A. Courtade, K. Venkat, and T. Weissman, "Justification of logarithmic loss via the benefit of side information," *IEEE Trans. on Information Theory*, vol. 61, no. 10, pp. 5357–5365, 2015.

[5] A. Painsky and G. W. Wornell, "On the universality of the logistic loss function," *arXiv preprint arXiv:1805.03804*, 2018.

[6] R. Linsker, "Self-organization in a perceptual network," *Computer*, vol. 21, no. 3, pp. 105–117, 1988.

[7] J. R. Quinlan, *C4. 5: Programs for Machine Learning*. Elsevier, 2014.

[8] C. Chow and C. Liu, "Approximating discrete probability distributions with dependence trees," *IEEE Trans. on Information Theory*, vol. 14, no. 3, pp. 462–467, 1968.

[9] C. Olsen, P. E. Meyer, and G. Bontempi, "On the impact of entropy estimation on transcriptional regulatory network inference based on mutual information," *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2009, no. 1, p. 308959, 2008.

[10] J. P. Pluim, J. A. Maintz, and M. A. Viergever, "Mutual-information-based registration of medical images: a survey," *IEEE Trans. on Medical Imaging*, vol. 22, no. 8, pp. 986–1004, 2003.

[11] P. Viola and W. M. Wells III, "Alignment by maximization of mutual information," *International journal of computer vision*, vol. 24, no. 2, pp. 137–154, 1997.

[12] N. Cesa-Bianchi and G. Lugosi, *Prediction, learning and games*. New York, USA: Cambridge, Univ. Press, 2006.

[13] E. L. Lehmann and G. Casella, *Theory of point estimation*. Springer Science & Business Media, 2006.

[14] O. Bousquet and A. Elisseeff, "Stability and generalization," *Journal of Machine Learning Research*, vol. 2, no. Mar, pp. 499–526, 2002.

[15] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan, "Learnability, stability and uniform convergence," *Journal of Machine Learning Research*, vol. 11, no. Oct, pp. 2635–2670, 2010.

[16] A. Xu and M. Raginsky, "Information-theoretic analysis of generalization capability of learning algorithms," in *Advances in Neural Information Processing Systems*, 2017, pp. 2521–2530.

[17] D. Russo and J. Zou, "How much does your data exploration overfit? controlling bias via information usage," *arXiv preprint arXiv:1511.05219*, 2015.

[18] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annual Allerton Conf. on Comm., Control, and Computing*, 1999, pp. 368–377.

[19] J. C. Principe, *Information theoretic learning: Renyi's entropy and kernel perspectives*. Springer Science & Business Media, 2010.

[20] S. Yu and J. C. Principe, "Understanding autoencoders with information theoretic concepts," *arXiv preprint arXiv:1804.00057*, 2018.

[21] S. Yu, R. Jenssen, and J. C. Principe, "Understanding convolutional neural network training with information theory," *arXiv preprint arXiv:1804.06537*, 2018.

[22] Y. Kong and G. Schoenebeck, "Water from two rocks: Maximizing the mutual information," *arXiv preprint arXiv:1802.08887*, 2018.

[23] Y. Ugur, I. E. Aguerri, and A. Zaidi, "A generalization of Blahut-Arimoto algorithm to computing rate-distortion regions of multiterminal source coding under logarithmic loss," in *Proc. of IEEE Info. Theory Workshop, ITW*, Kaohsiung, Taiwan, Nov. 2017.

[24] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *ArXiv e-prints*, Dec. 2013.

[25] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "β-vae: Learning basic visual concepts with a constrained variational framework," 2016.

[26] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society, Series B*, vol. 39, no. 1, pp. 1–38, 1977.

[27] N. Slonim and N. Tishby, "Document clustering using word clusters via the information bottleneck method," in *Proc. of 23rd Ann. Int'l ACM SIGIR Conf. on Res. and Dev. in Info. Retrieval*, 2000.

[28] N. Slonim, G. S. Atwal, G. Tkačik, and W. Bialek, "Information-based clustering," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 51, pp. 18 297–18 302, 2005.

[29] A. Globerson and N. Tishby, "On the optimality of the Gaussian information bottleneck curve," *Hebrew University Tech. Report*, 2004.

[30] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, "Information bottleneck for Gaussian variables." *Journal of Machine Learning Research*, vol. 6, pp. 165–188, Feb. 2005.

[31] M. Chalk, O. Marre, and G. Tkacik, "Relevant sparse codes with variational information bottleneck," 05 2016.

[32] A. Alemi, I. Fischer, J. Dillon, and K. Murphy, "Deep variational information bottleneck," in *ICLR*, 2017. [Online]. Available: https://arxiv.org/abs/1612.00410

[33] A. Achille and S. Soatto, "Information dropout: Learning optimal representations through noisy computation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. PP, no. 99, pp. 1–1, 2018.

[34] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *Proceedings of the 31st International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, E. P. Xing and T. Jebara, Eds., vol. 32, no. 2. Bejing, China: PMLR, 22–24 Jun 2014, pp. 1278–1286.

[35] R. Dobrushin and B. Tsybakov, "Information transmission with additional noise," *IRE Trans. on Information Theory*, vol. 8, no. 5, pp. 293–304, Sep. 1962.

[36] H. S. Witsenhausen and A. D. Wyner, "A conditional entropy bound for a pair of discrete random variables," *IEEE Trans. on Information Theory*, vol. 21, no. 5, pp. 493–501, September 1975.

[37] R. Shwartz-Ziv and N. Tishby, "Opening the black box of deep neural networks via information." [Online]. Available: http://arxiv.org/abs/1703.00810

[38] A. Achille and S. Soatto, "Emergence of Invariance and Disentangling in Deep Representations," *ArXiv e-prints*, Jun. 2017.

[39] D. A. McAllester, "A PAC-bayesian tutorial with a dropout bound," *CoRR*, vol. abs/1307.2118, 2013. [Online]. Available: http://arxiv.org/abs/1307.2118

[40] K. Kittichokechai and G. Caire, "Privacy-constrained remote source coding," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2016, pp. 1078–1082.

[41] C. Tian and J. Chen, "Successive refinement for hypothesis testing and lossless one-helper problem," *IEEE Trans. Info. Theory*, vol. 54, no. 10, pp. 4666–4681, Oct. 2008.

[42] S. Sreekumar, D. Gündüz, and A. Cohen, "Distributed Hypothesis Testing Under Privacy Constraints," *ArXiv e-prints*, Jul. 2018.

[43] I. E. Aguerri, A. Zaidi, G. Caire, and S. Shamai, "On the capacity of cloud radio access networks with oblivious relaying," *CoRR*, vol. abs/1710.09275, 2017. [Online]. Available: http://arxiv.org/abs/1710.09275

[44] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.

[45] L. Paninski, "Estimation of entropy and mutual information," *Neural computation*, vol. 15, no. 6, pp. 1191–1253, 2003.

[46] J. Jiao, K. Venkat, Y. Han, and T. Weissman, "Minimax estimation of functionals of discrete distributions," *IEEE Trans. on Information Theory*, vol. 61, no. 5, pp. 2835–2885, 2015.

[47] X. B. Peng, A. Kanazawa, S. Toyer, P. Abbeel, and S. Levine, "Variational discriminator bottleneck: Improving imitation learning, inverse rl, and gans by constraining information flow," 2018.

[48] B. Dai, C. Zhu, and D. Wipf, "Compressing neural networks using the variational information bottleneck," 2018.

[49] C. Xu, D. Tao, and C. Xu, "A survey on multi-view learning," *arXiv preprint arXiv:1304.5634*, 2013.

[50] W. Wang, R. Arora, K. Livescu, and J. Bilmes, "On deep multi-view representation learning: Objectives and optimization," 2016.

[51] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," in *Proc. of the eleventh annual conf. on computational learning theory, ACM*, 1998, pp. 92–100.

[52] P. Dhillon, D. P. Foster, and L. H. Ungar, "Multi-view learning of word embeddings via CCA," in *Advances in neural information processing systems*, 2011, pp. 199–207.

[53] A. Kumar and H. Daumé, "A co-training approach for multi-view spectral clustering," in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 393–400.

[54] M. Gönen and E. Alpaydın, "Multiple kernel learning algorithms," *Journal of Machine Learning Research*, vol. 12, no. Jul, 2011.

[55] Y. Jia, M. Salzmann, and T. Darrell, "Factorized latent spaces with structured sparsity," in *Advances in Neural Information Processing Systems*, 2010, pp. 982–990.

[56] V. Vapnik, *The nature of statistical learning theory*. Springer science & business media, 2013.

[57] A. A. Alemi, B. Poole, I. Fischer, J. V. Dillon, R. A. Saurous, and K. Murphy, "Fixing a broken elbo," 2017.

[58] A. A. Alemi and I. Fischer, "Therml: Thermodynamics of machine learning," 2018.

[59] T. A. Courtade and T. Weissman, "Multiterminal source coding under logarithmic loss," *IEEE Trans. on Information Theory*, vol. 60, no. 1, pp. 740–761, Jan. 2014.

[60] R. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. on Information Theory*, vol. 18, no. 4, pp. 460–473, Jul 1972.

[61] A. E. Gamal and Y.-H. Kim, *Network Information Theory*. Cambridge University Press, 2011.

[62] G. E. Hinton and D. van Camp, "Keeping the neural networks simple by minimizing the description length of the weights," in *Proceedings of the Sixth Annual Conference on Computational Learning Theory*, ser. COLT '93. New York, NY, USA: ACM, 1993, pp. 5–13. [Online]. Available: http://doi.acm.org/10.1145/168304.168306

[63] P. Harremoes and N. Tishby, "The information bottleneck revisited or how to choose a good distortion measure," in *Proc. IEEE Int. Symp. Information Theory*, Jun. 2007, pp. 566–570.

[64] D. P. Kingma and M. Welling, "Auto-encoding variational bayes." *CoRR*, vol. abs/1312.6114, 2013. [Online]. Available: http://dblp.uni-trier.de/db/journals/corr/corr1312.html#KingmaW13

[65] A. Winkelbauer, S. Farthofer, and G. Matz, "The rate-information trade-off for gaussian vector channels," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*. IEEE, 2014, pp. 2849–2853.

[66] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, 1991.

[67] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM J. on Opt.*, vol. 23, pp. 1126–1153, Jun. 2013.

[68] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," 2017. [Online]. Available: https://arxiv.org/abs/1611.01144

[69] C. J. Maddison, A. Mnih, and Y. W. Teh, "The concrete distribution: A continuous relaxation of discrete random variables," *CoRR*, vol. abs/1611.00712, 2016. [Online]. Available: http://arxiv.org/abs/1611.00712

[70] D. P. Kingma, T. Salimans, and M. Welling, "Improving variational inference with inverse autoregressive flow," *CoRR*, vol. abs/1606.04934, 2016. [Online]. Available: http://arxiv.org/abs/1606.04934

[71] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked Autoregressive Flow for Density Estimation," *ArXiv e-prints*, May 2017.

[72] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization." *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: http://dblp.uni-trier.de/db/journals/corr/corr1412.html#KingmaB14

[73] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS10). Society for Artificial Intelligence and Statistics*, 2010.

[74] A. M. Saxe, Y. Bansal, J. Dapello, M. Advani, A. Kolchinsky, B. D. Tracey, and D. D. Cox, "On the information bottleneck theory of deep learning," in *International Conference on Learning Representations*, 2018.

[75] E. Ekrem and S. Ulukus, "An outer bound for the vector Gaussian CEO problem," *IEEE Trans. on Inf. Theory*, vol. 60, no. 11, pp. 6870–6887, Nov 2014.

[76] A. Dembo, T. M. Cover, and J. A. Thomas, "Information theoretic inequalities," *IEEE Trans. on Inf. Theory*, vol. 37, no. 6, pp. 1501–1518, Nov 1991.

[77] E. Ekrem and S. Ulukus, "An outer bound for the vector Gaussian CEO problem," *IEEE Trans. on Inf. Th.*, vol. 60, no. 11, pp. 6870–6887, Nov. 2014.

[78] Y. Zhou, Y. Xu, W. Yu, and J. Chen, "On the optimal fronthaul compression and decoding strategies for uplink cloud radio access networks," *IEEE Trans. on Information Theory*, vol. 62, no. 12, pp. 7402–7418, Dec. 2016.