

Opportunistic Information-Bottleneck for Goal-Oriented Feature Extraction and Communication

FRANCESCO BINUCCI^{1,2} (Graduate Student Member, IEEE), PAOLO BANELLI¹,
PAOLO DI LORENZO^{2,3} (Senior Member, IEEE), AND SERGIO BARBAROSSA³ (Life Fellow, IEEE)

¹Department of Engineering, University of Perugia, 06128 Perugia, Italy

²Consorzio Nazionale Interuniversitario per le Telecomunicazioni, 43124 Parma, Italy

³Department of Information Engineering, Electronics, and Telecommunications, Sapienza University of Rome, 00184 Rome, Italy

CORRESPONDING AUTHOR: F. BINUCCI (e-mail: francesco.binucci@studenti.unipg.it)

This work was supported in part by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future”, Program “RESTART” under Grant PE00000001, and in part by the 6G-GOALS Project under the 6G SNS-JU Horizon Program under Grant 101139232.

ABSTRACT The Information Bottleneck (IB) method is an information theoretical framework to design a parsimonious and tunable feature-extraction mechanism, such that the extracted features are maximally relevant to a specific learning or inference task. Despite its theoretical value, the IB is based on a functional optimization problem that admits a closed form solution only on specific cases (e.g., Gaussian distributions), making it difficult to be applied in most applications, where it is necessary to resort to complex and approximated variational implementations. To overcome this limitation, we propose an approach to adapt the closed-form solution of the Gaussian IB to a general task. Whichever is the inference task to be performed by a (possibly deep) neural-network, the key idea is to *opportunistically* design a regression sub-task, embedded in the original problem, where we can safely assume a (joint) multivariate normality between the sub-task’s inputs and outputs. In this way we can exploit a fixed and pre-trained neural network to process the input data, using a tunable number of features, to trade data-size and complexity for accuracy. This approach is particularly useful every time a device needs to transmit data (or features) to a server that has to fulfil an inference task, as it provides a principled way to extract the most relevant features for the task to be executed, while looking for the best trade-off between the size of the feature vector to be transmitted, inference accuracy, and complexity. Extensive simulation results testify the effectiveness of the proposed method and encourage to further investigate this research line.

INDEX TERMS Goal-oriented communications, information bottleneck, edge-intelligence.

I. INTRODUCTION

ARTIFICIAL Intelligence (AI) and Machine Learning (ML) are experiencing an unprecedented explosion nowadays, with a proliferation of applications and use-cases, such as Generative Deep Learning [1], Natural Language Processing (NLP) [2], object detection/recognition [3], motion planning [4], and many others.

The sixth generation of mobile networks is expected to be one of the main key-enablers for these applications [5], with an ever-increasing integration of AI/ML blocks within

the communication infrastructure. In many ML application scenarios, there is the need to bring intelligence to peripheral devices (sensors), which however have very limited resources (e.g., IoT devices, Unmanned Aerial Vehicles (UAVs) [6], etc.). The problem is even more critical when the applications to be run are delay-sensitive, or there is a stringent constraint either on the accuracy of the decision to be taken, or on the energy consumption, or both. In these cases, the Edge Intelligence (EI) paradigm [7] offers a solution through computation offloading from the Edge Device (ED) to an

Edge Server (ES), thus enabling learning and inference tasks under strict energy, latency, and reliability constraints.

In an energy saving perspective, it is important to maintain the computational complexity, of the employed learning models, as low as possible. At the same time, taking into account the envisaged exponential traffic growth in the next years [8], [9], we also need to search for new and more sustainable communication paradigms, capable to save bandwidth and power resources, while preserving the effectiveness. To this end, Goal-Oriented Communications (GOCs) [8] represent a hot research topic, suggesting to overcome the classical Shannon paradigm [10], which focuses on recovering all the transmitted bits, irrespective of the purpose those bits are transmitted for. Conversely, GOCs tunes the source encoding rule and adapts the transmission rate focusing directly on the task that motivated the exchange of information, while respecting task-specific performance constraints. This approach makes possible to reach high compression degrees, and to save as much transmission resources as possible, any time an ED decides to offload computations to an ES.

From this viewpoint, the Information Bottleneck (IB) method [11], inspired by rate-distortion theory arguments, is a theoretical framework that can be used to extract from data (and parsimoniously encode) those features that are particularly meaningful for a specific learning or inference task, and, consequently, it is a promising way to formalize and implement GOCs. Specifically, the IB method is designed to retrieve a representation of an input random source that, for any (compact) features size and its associated bit-compressed representation, is as much informative as possible with respect to the target of a specific learning or inference task [11]. Specifically, the IB acts as a supervised Feature Extractor (FE) where, differently from unsupervised methods such as PCA, the compression targets a proper representation for the final decision, rather than the fidelity of the input reconstruction. Unfortunately, since the IB relies on a functional optimization problem based on Mutual Information (MI), it admits a closed-form solution only on specific cases [11]. The most noticeable one is when the input and the output of a regression task are characterized by a joint multi-variate Gaussian distribution [12], [13]. The aim of this paper is to overcome this limitation, proposing a quite general framework to adapt the Gaussian IB principle to tasks that deviate from multi-variate Normal regression, which turns out to be particularly helpful to deploy Edge Intelligence tasks using a GOC philosophy.

Related Works. There are several works in the recent literature where the IB principle is employed to support ML applications [14]. The authors in [15], [16] propose IB-based approaches for medical imaging classification and segmentation. To overcome the difficulty to adapt the IB principle without Gaussian assumptions, their solution relies on a Variational Information Bottleneck (VIB) framework [17] that, considering a variational bound of the IB cost function,

derives a loss function that is employed to train Variational Auto-Encoders (VAEs) [18].

In [19], the authors propose a variation of the original Information Bottleneck problem named Scalable Information Bottleneck, where multiple compressed representations, with increasingly richer features, are considered. The work in [20] proposes a solution for a distributed implementation of the IB problem, that is suitable in both the discrete and Gaussian case.

Other noticeable examples can be found in [21], [22], where the IB is proposed as a method to improve the generalization capabilities for DNN-based tasks, also in the presence of out-of-distribution (OOD) data. The IB principle has been applied also in conjunction with Graph Neural Networks (GNNs), in the so-called Graph Information Bottleneck framework [23], [24].

Some recent works testify the importance of the IB method for GOCs communications. Specifically, the approaches proposed in [25], [26], [27], [28], [29] can be exploited in Edge Intelligence settings, where one or more Edge Devices offload a specific learning/inference task towards the servers placed at the edge network. The work in [30] proposes an IB framework for task-oriented communications, which is based on a slight modification of the IB formulation to cope with OOD data. Reference [31] proposes a Robust Information Bottleneck (RIB) formulation to cope with digital communication schemes. In [13], the authors propose an optimal resource allocation framework based on the Gaussian IB, e.g., regression tasks, for Edge Machine Learning applications. The authors in [32] propose an approach to optimize the shared codebook in Type-Based Multiple Access (TBMA) based on the Information Bottleneck.

Except for [13], where the Gaussian assumption holds true, all the aforementioned works are based on the VIB, which retrieves a good approximation of the IB optimal solution, paying the cost of a considerable complexity. To reduce this complexity, [33], [34] propose an heuristic approximation of the IB principle, where the latent (and tunable) compact representation, together with the final decisions, are obtained by a framework that exploits multiple couples of Convolutional Encoders (CEs) and Convolutional Classifiers (CCs). However, this approach is highly based on empirical considerations, without any theoretical claim on the performance and accuracy that is possible to obtain by different compression architectures, or different feature sizes. Actually, although it is practically impossible to derive a closed-form expression to link the inference accuracy of a neural network with the input feature vectors, some sub-optimal and theoretically grounded criteria can be exploited. A possible attempt in this direction can be found in [35], where the authors study a multi-user Edge-Intelligence scenario, where the Edge Devices perform feature extraction and quantization prior to the transmission of the data towards a centralized edge-server for the final inference. Specifically, the authors in [35] propose to quantify performance by the

so called *discriminant-gain* (i.e., the symmetrized Kullback-Leibler divergence between two classes in the Euclidean feature space). The method we propose, if applied to a Goal-Oriented and Edge-Assisted Communication Scenario (which is out of the scope of this paper) would share some similarities with [35]: indeed, both the methods rely on Gaussian assumptions of the features, and their linear extraction by either PCA in [35], or Gaussian Information Bottleneck, as we propose. Specifically, the authors in [35] propose the discriminant-gain to optimally allocate resources, such as quantization bits, under energy and latency constraints. Actually, our proposal does not consider features quantization, which could be handled either by the approach proposed in [35], or simply by exploiting mean squared error (MSE) cost functions, as we will detail in Section III-A. Furthermore, differently from [35], our approach is naturally oriented to *dynamically* optimize also the feature's vector size, by means of Section III-B, Fig. 2 and [33], [34], [36].

Our Contribution. In this paper we propose a novel method to adapt the Gaussian IB (GIB) principle [12] to any Goal-Oriented inference task that exploits a (Deep) NN to come up with a decision. In a nutshell, our solution embeds an *opportunistic* regression step between a transformation of the input data and the output of the first layer of the NN trained to perform a task (for instance, image classification). This opportunistic regression is instrumental to exploit the IB method to extract the most relevant features from the data in a principled way, giving us the flexibility to trade the learning task accuracy with the feature's vector size. Differently from [33], [34], [36], the proposed solution can be implemented by training and exploiting a *single* DNN architecture, rather than a bank of DNNs, e.g., one for each specific size of the input features. The advantages of our formulation with respect to the Variational Information Bottleneck implementations [17], which is widely considered in the recent literature, are mainly related to the inference and training complexity:

- First of all, in classical VIB implementations, the size of the compressed representation is constrained by the encoder architecture. Thus, like the scheme proposed in [33], [34], [36], VIB requests different encoder/classifier architectures for different compression levels. Conversely, the proposed approach may possibly rely on a single (classifying) neural network, which follows a simple compression stage and is trained once and for all.
- Compared to the proposed OIB approach, the VIB encoding requests an extra neural network to produce a compact representation of the data, which has to be jointly trained with the classifying network. Furthermore, the VIB training procedure needs Monte-Carlo sampling (obtained through the reparametrization trick) to obtain an unbiased estimation of the gradient, thus further increasing the complexity with respect to the proposed OIB.

TABLE 1. Main Notation and Definitions.

Symbol	Definition
$\mathcal{D}_{tr}, \mathcal{D}_{test}$	Training and test-sets
$\mathbf{x}_n, \mathbf{y}_n$	Input and output for the n -th sample
$\tilde{\mathbf{x}}_n$	Gaussian transformation of the n -th sample
$\tilde{\mathbf{y}}_n$	Target of the n -th sample of the Gaussian regression task
n_x, n_y	Size of the input and output variables
\mathbf{z}_n	Compressed representation for the n -th sample
n_z	Size of the compressed representation
ρ	Compression ratio
\mathbf{A}_ρ	Compression matrix
Θ_ρ	Re-expansion matrix
Σ	Covariance matrix
$\Sigma_{\cdot }$	Conditional Covariance Matrix
$H(\cdot)$	Differential Entropy
$I(\cdot, \cdot)$	Mutual Information

The simulation results testify the effectiveness of the proposed approach, showing either a performance gain with respect to other compression strategies, or lower complexity, or both. Furthermore, the simulation results also show that the proposed approach allows to minimize the entropy of the compressed representations of the input data, with respect to competitive approaches. This means, that it is theoretically possible to search for compression schemes capable to represent the source with less bits, making our proposal particularly attractive for Edge-Assisted Goal-Oriented Communications.

Outline. The rest of this paper is organized as follows. In Section II, we give a general background about feature compression and we briefly recall the Information Bottleneck problem, with focus on the Gaussian case. Then, in Section III, we describe our opportunistic approach to adapt the Gaussian IB solution to general Inference Tasks, and in Section IV we present and discuss our simulation results. Finally, Section V draws the conclusion and sketches the possible future research directions.

II. BACKGROUND

This section provides a brief summary about feature extraction and compression, and the Information Bottleneck method. For the sake of clarity, in Table 1 we report the main notation and definitions used throughout the paper.

A. FEATURE EXTRACTION AND COMPRESSION

Feature-Extraction (FE) aims to extract features from the raw data that are relevant for the underlying inference task. Classical machine learning models (e.g., SVM, linear regression, etc.) envisage to perform Feature-Extraction as a separated data pre-processing stage. Conversely, when we deal with Deep Neural Networks, FE is directly embedded in the first layers, which allows us to retrieve parsimonious data representations containing all the necessary information to fulfil the final learning/inference task. In some cases, Feature Extraction can lead to a substantial dimensionality

reduction of the input data. In this case we talk about *Feature Compression*, which can be highly beneficial in all applications where it is required to share or communicate the learned representation, as in Edge Intelligence scenarios.

Specifically, let us consider a data-set $\mathcal{D}\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, where $\mathbf{x}_n \in \mathbb{R}^{n_x}$ are input variables and $\mathbf{y}_n \in \mathbb{R}^{n_y}$ are the associated labels. Feature-Compression is a pre-processing stage that reduces the size of the input data \mathbf{x}_n from n_x to n_z , by generating a new variable $\mathbf{z}_n \in \mathbb{R}^{n_z}$, whose compression-ratio is expressed by $\rho = n_x/n_z$. The size reduction aims to simplify the decision function of the learning task. Furthermore, a well-designed features extraction process should allow to capture only the information that is *relevant* for the specific inference task, with a possible improvement in terms of the model generalization capabilities [37].

In general, a feature extractor is modeled through a function $\mathbf{z}_\rho = f_\rho(\mathbf{x})$. When $f(\cdot)$ is linear with respect to \mathbf{x} , we end up with a Linear Feature Extractor (LFE)

$$\mathbf{z}_\rho = \mathbf{A}_\rho \mathbf{x}, \quad (1)$$

where $\mathbf{A}_\rho \in \mathbb{R}^{n_z \times n_x}$ is the *compression matrix*, for a specific compression ratio ρ . Typically, LFEs are employed before an inference/learning task based on non-linear models, e.g., a (Deep) NN [38], while non-linear FE are typically used jointly with a linear discriminant function.

The compression matrix \mathbf{A}_ρ can be designed according to different (possibly optimal) criteria. *Supervised* LFEs derive the projection matrix plugging the training labels \mathbf{y}_n in the design criteria, differently from *unsupervised* LFEs, like PCA [39], which take into account only the training input data, identifying for instance the most relevant features to reconstruct the original input \mathbf{x}_n from the compressed features $\mathbf{z}_{\rho,n}$.

In the sequel we will show how the Gaussian Information Bottleneck solution naturally leads to a supervised Feature-Compression process, which can be exploited to design parsimonious computational offloading operations in Edge-Assisted inference tasks. More specifically, the GIB does not only identify the best features to extract, for any given feature's size, but it also weight them in order to minimize a lower bound on the bits necessary for their representation.

B. THE INFORMATION BOTTLENECK METHOD

Let's suppose to have two random variables $\mathbf{x} \in \mathbb{R}^{n_x}$ and $\mathbf{y} \in \mathbb{R}^{n_y}$. Our goal is to characterize the information on \mathbf{x} that is relevant about \mathbf{y} . In a ML setting, \mathbf{x} and \mathbf{y} represent the input and the output of a specific inference task, respectively. The IB method aims to find a *probabilistic* compact representation \mathbf{z} of the input variable \mathbf{x} , while preserving a certain amount of information about the output of the inference task \mathbf{y} . This problem can be formulated in Lagrangian form [40] as follows [11]

$$\underset{p(\mathbf{z}|\mathbf{x})}{\text{minimize}} \quad I(\mathbf{z}, \mathbf{x}) - \beta I(\mathbf{z}, \mathbf{y}), \quad (2)$$

where $I(\cdot, \cdot)$ denotes the mutual information [41] function. More specifically, we look for a statistical mapping $p(\mathbf{z}|\mathbf{x})$

that, minimizing $I(\mathbf{x}, \mathbf{z})$ encourages a *probabilistic* compression on \mathbf{z} , while controlling through $I(\mathbf{z}, \mathbf{y})$ the amount of information retained by \mathbf{z} about the output \mathbf{y} .

When the Lagrange multiplier $\beta \rightarrow \infty$ we get a more informative representation, while at the same time penalizing the compression. Otherwise, as $\beta \rightarrow 0$, we obtain a more compact representation in probabilistic sense, sacrificing the learning performance.

In the general case, it does not exist a closed-form solution for (2). In the discrete case, we must resort to an iterative fixed-point algorithm [11], [42], which suffers of a considerable computational complexity and some instability, due to the estimation of the mutual information. For the continuous case, it is possible to resort to sub-optimal VIB formulations, which are quite complex to be implemented in practice [25]. A valuable exception is represented by the Gaussian case [12], which admits an elegant closed form solution, as we recall in the following.

C. GAUSSIAN INFORMATION BOTTLENECK

If \mathbf{x} and \mathbf{y} are jointly characterized by a multivariate Gaussian distribution, i.e., $(\mathbf{x}, \mathbf{y}) \sim \mathcal{N}(0, \Sigma_{xy})$, the solution of (2) can be expressed in closed-form. More specifically, it can be proved that the optimal statistical mapping is linear, and it is given by [12]

$$\mathbf{z} = \mathbf{A}_\rho \mathbf{x} + \xi, \quad (3)$$

where $\xi \sim \mathcal{N}(0, \mathbf{I})$ is an additive Gaussian vector, statistically independent of \mathbf{x} and \mathbf{z} , and the compression matrix \mathbf{A}_ρ is built from the eigenvalues and left eigenvectors $\{\lambda_i, \mathbf{v}_i\}_{i=1, \dots, n_x}$ of $\Sigma_{x|y} \Sigma_x^{-1}$, where $\Sigma_{x|y}$ is the conditional covariance matrix of \mathbf{x} given \mathbf{y} , while Σ_x^{-1} is the inverse covariance matrix of \mathbf{x} . Interestingly, the solution relies on the same eigenvectors used in Canonical Correlation Analysis (CCA) [43]. More specifically, sorting the eigenvectors \mathbf{v}_i by the ascending values of the corresponding eigenvalues λ_i , for any fixed value of the Lagrange multiplier β , the matrix \mathbf{A}_ρ is obtained by the non-zero rows of

$$\mathbf{A} = \begin{cases} [\mathbf{0}^T; \mathbf{0}^T; \dots; \mathbf{0}^T], & 0 \leq \beta \leq \beta_1^c \\ [\alpha_1 \mathbf{v}_1^T; \mathbf{0}^T; \dots; \mathbf{0}^T], & \beta_1^c < \beta \leq \beta_2^c \\ [\alpha_1 \mathbf{v}_1^T; \alpha_2 \mathbf{v}_2^T; \mathbf{0}^T; \dots; \mathbf{0}^T], & \beta_2^c < \beta \leq \beta_3^c \\ \vdots & \vdots \\ [\alpha_1 \mathbf{v}_1^T; \alpha_2 \mathbf{v}_2^T; \dots; \alpha_n \mathbf{v}_n^T], & \beta_n^c < \beta \end{cases} \quad (4)$$

where $\beta_i^c = \frac{1}{1-\lambda_i}$, $\alpha_i = \sqrt{\frac{\beta(1-\lambda_i)-1}{\lambda_i r_i}}$, and $r_i = \mathbf{v}_i^t \Sigma_x \mathbf{v}_i$.

The structure of the matrix \mathbf{A} makes evident the trade-off between the features' compression and the learning performance. Indeed, considering that only the first $n_z = n_x/\rho$ non-zero rows are used for compression when $\beta_{n_z}^c < \beta \leq \beta_{n_z+1}^c$, low values of β induces the use of few CCA eigenvectors, thus reaching an extreme compression degree, with a consequent degradation in terms of learning performance. On the other hand, as β increases, we start adding more eigenvectors to the compression matrix \mathbf{A}_ρ ,




FIGURE 1. GIB architecture. \mathbf{x} is the input of the specific inference task (e.g., image classification), $\tilde{\mathbf{x}}$ represents the Gaussian Transformation of the data while $\tilde{\mathbf{y}}$ is the output of the first layer of the considered neural network.

thus obtaining a larger intermediate representation \mathbf{z}_ρ , with a higher amount of information with respect to \mathbf{y} .

D. RELATIONSHIPS BETWEEN GIB AND CCA

Looking at eq. (4) we note that, differently from CCA, the GIB solution characterizes also the loadings α_i . Furthermore, it considers a noise term ξ in the mapping. These differences arise since the IB aims to optimize the trade-off between *compression*, captured by the term $I(\mathbf{x}, \mathbf{z})$, and *relevance* with respect to the outcome, captured by $I(\mathbf{z}, \mathbf{y})$. Conversely, if we fix a priori the dimension n_z of the feature vector, it makes sense to search for the best linear mapping $\mathbf{z} = \mathbf{M}\mathbf{x} + \epsilon$ that maximizes $I(\mathbf{z}, \mathbf{y})$. Solving the optimization problem under Gaussian assumptions, turns out that the optimal transformation matrix \mathbf{M}^* is given by the eigenvectors of $\Sigma_{\mathbf{x}|\mathbf{y}}\Sigma_{\mathbf{x}}^{-1}$. Furthermore, $\Sigma_\epsilon = \mathbf{0}$ (cf. Appendix for the proof). Interestingly, the (accuracy) optimal solution does not depend on the loadings α_i of the eigenvectors (and is equivalent in this sense to usual CCA). This fact leads us to the following considerations:

- Under Gaussian assumptions, the projection on the CCA basis vector is the best linear projection in terms of mutual information $I(\mathbf{z}, \mathbf{y})$, i.e., accuracy-wise.
- Since the solution does not depend on the loadings α_i on (4) (see Appendix), if we focus only on the size of the compressed representation \mathbf{z} without care about its representational complexity, GIB and CCA become equivalent.

Thus, in performing feature extraction with a fixed number of components and a fixed encoding, if Multivariate Gaussian assumptions hold true, we can get rid of the noise term ϵ and project on the basis vectors of CCA. This is actually obvious, because the noise term of the GIB is not informative with respect to the task outcome \mathbf{y} .

Conversely, the GIB mapping in (3) gives us representations that are optimized in terms of minimum $I(\mathbf{x}, \mathbf{z})$, i.e., offering the possibility to design an encoding rule that minimizes the number of bits [41], in the spirit of rate-distortion theory arguments. While this aspect may be of relative importance in classical ML, which typically focuses only on the feature's size n_z , it turns out particularly useful in Edge Inference scenarios assisted by Goal-Oriented communications, where we want to save as much transmission resources (e.g., source-coding bits) as possible,

while targeting specific requirements in terms of learning performance [8].

III. PROPOSED APPROACH

Herein we present a strategy to adapt the GIB framework to a generic (non-linear, non-Gaussian) inference task. The proposed system model is depicted in Fig. 1. We start from the input \mathbf{x} of the inference task and then we apply a linear and invertible data transformation $h(\cdot)$ to obtain another vector $\tilde{\mathbf{x}} = h(\mathbf{x})$, which is assumed to be well approximated by a multivariate Normal, by Central Limit Theorem (CLT) arguments [44]. Without restriction of generality, in this manuscript we employ a (bidimensional) Discrete Fourier Transform (DFT) [45], which enjoys the appealing feature of fast and low complexity implementation by Fast Fourier Transform (FFT) algorithms. Thus, $\tilde{\mathbf{x}}$ represents the input for an *opportunistic* regression task, that we define on purpose to exploit the GIB framework. $\tilde{\mathbf{x}}$ is also the input to the overall classifying NN shown in Fig. 1, where \mathbf{L}_0 is an operator representing the mixing stage of the first NN layer (before the activation function).

In this work, we stick with a linear mapping, i.e., \mathbf{L}_0 is a matrix, because it preserves the joint Gaussianity of $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$: this is instrumental to exploit the GIB principle to infer $\tilde{\mathbf{y}}$ by a compressed version \mathbf{z}_ρ of $\tilde{\mathbf{x}}$. This is consistent with the use of a shallow neural networks, as well as of a convolutional neural networks [46].

Denoting the training-set as $\mathcal{D}_{\text{tr}} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^{N_{\text{tr}}}$, the proposed training procedure is essentially divided in three steps:

- 1) We firstly train the preferred inference network on the (transformed) training set $\{(\tilde{\mathbf{x}}_n, \mathbf{y}_n)\}_{n=1}^{N_{\text{tr}}}$.¹
- 2) On the trained model, we compute the transformation $\tilde{\mathbf{y}}_n = \mathbf{L}_0\tilde{\mathbf{x}}_n + \lambda\eta$, where $\eta \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{n_y})$, and λ is a regularization parameter to be chosen. This way, we construct a regression data-set $\hat{\mathcal{D}}_{\text{tr}} = \{(\tilde{\mathbf{x}}_n, \tilde{\mathbf{y}}_n)\}_{n=1}^{N_{\text{tr}}}$.
- 3) We use $\hat{\mathcal{D}}_{\text{tr}}$ to compute the sample-mean estimators of the CCA covariance matrices $\Sigma_{\tilde{\mathbf{x}}}$ and $\Sigma_{\tilde{\mathbf{x}}|\tilde{\mathbf{y}}}$, that the GIB exploits to build the compression matrix \mathbf{A}_ρ , as detailed in Section II-C and (4).²

¹ We stress that the method can be applied to any inference/classification network. Thus, herein we do not focus on any sort of optimality of the network, and we will mostly use convolutional NN networks, due to the image classification application we focus on to verify the results.

² Actually, step (2) could be avoided because $\tilde{\mathbf{y}} = \mathbf{L}_0\tilde{\mathbf{x}}$ and $\Sigma_{\tilde{\mathbf{x}}|\tilde{\mathbf{y}}}$ can be directly computed by the knowledge (estimate) of $\Sigma_{\tilde{\mathbf{x}}}$ and \mathbf{L}_0 [47].

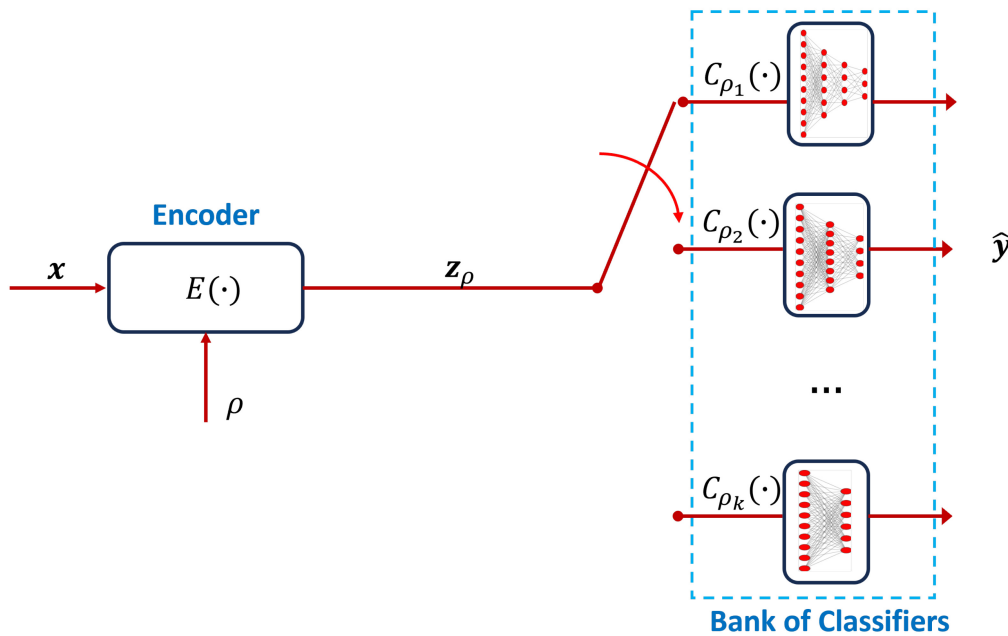
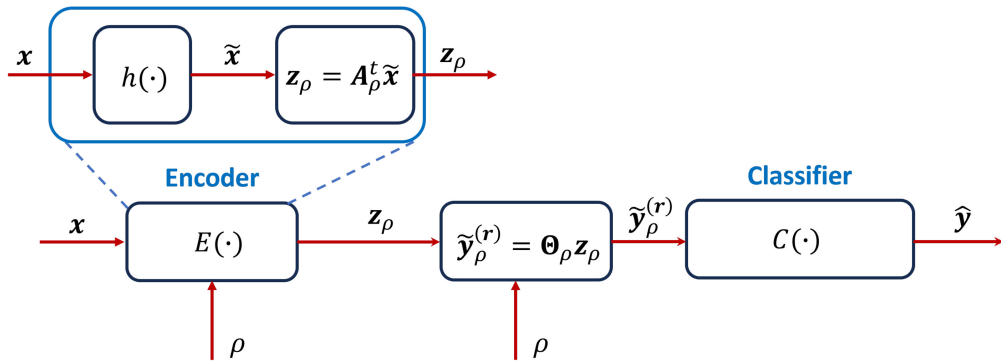


FIGURE 2. Alternative classification schemes. Note that although each LFE method shares the same linear structure of the encoder $E(\cdot)$, it is characterized by a specific (linear) transformation $h(\cdot)$ and compression matrix \mathbf{A}_ρ .

Specifically, the noise term we add at the output of \mathbf{L}_0 is necessary since the Gaussian Information Bottleneck fails in dealing with deterministic transformations [12]. Indeed, in case of deterministic mapping, the eigenvalues of $\Sigma_{x|y}\Sigma_x^{-1}$ would assume values in the set $[0, 1]$ [43], thus making infeasible the computation of the loads α_i in (4). This way, the solution is equivalent to the well-known Ridge Canonical Correlation Analysis (RCCA) [48], a regularized version of the classical CCA with a penalty on the L_2 norm of the projection vectors.

Once we learned how to compress features by the proposed Opportunistic IB (OIB), we can exploit this knowledge to classify data on the test-set $\mathcal{D}_{\text{test}}$, according to two different schemes, as detailed in the two following sub-sections. In particular, note that, for any possible compression ratio ρ , the associated matrix \mathbf{A}_ρ is simply obtained by (4), selecting the first n_z eigenvectors of a fixed

CCA. Thus, the compressive transformation is unique, and just the loading parameters α_i depends on (and has to be stored for) the specific compression ratio $\rho = n_x/n_z$, i.e., for the specific $\beta = \beta_{n_z+1}^c$.³ This is clearly attractive complexity wise.

A. CLASSIFICATION WITH FEATURE EXPANSION

The first and more appealing strategy, depicted in Fig. 2(a), envisages to use a *single* NN $C(\cdot)$ to perform the inference task, i.e., the same NN for any compression ratio ρ used by the OIB-based features' compression. To this end, we don't use the first layer \mathbf{L}_0 the original NN has been trained with, and we use the GIB-compressed representation \mathbf{z}_ρ to

³Note that, as detailed in the following, herein we can handle the classification based on a *size-tunable* compressed representation \mathbf{z}_ρ , by using a *single* NN, rather than a bank of encoders and classifiers, as it happens in [33], [34].

produce an estimate $\tilde{\mathbf{y}}_\rho^{(r)} = g_\rho(\mathbf{z}_\rho)$ of the $\tilde{\mathbf{y}} = \mathbf{L}_0\tilde{\mathbf{x}}$ that would have been actually obtained by using the first NN layer \mathbf{L}_0 . This *reconstructed* $\tilde{\mathbf{y}}_\rho^{(r)}$ is used to feed the *shortened* classification network $C(\cdot)$. For any fixed β , i.e., for any fixed compression ratio $\rho = n_x/n_z$, we assume the estimator $g_\rho(\mathbf{z}_\rho)$ to be linear, as expressed by

$$\tilde{\mathbf{y}}_\rho^{(r)} = \Theta_\rho \mathbf{z}_\rho. \quad (5)$$

Actually, (5) includes also the optimal Bayesian MMSE estimator, when $\tilde{\mathbf{y}}$ and \mathbf{z}_ρ are jointly Gaussian, as it can be safely assumed for the proposed OIB, thanks to the capability of the 2D-DFT transform to generate (almost) Gaussian distributed $\tilde{\mathbf{x}}$.

Thus, in this case, the Bayesian optimal reconstruction matrix Θ_ρ is given by the classical Linear (L)-MMSE estimator that, assuming zero-mean data, is expressed by [47]

$$\Theta_\rho = \mathbf{C}_{\tilde{\mathbf{y}}, \mathbf{z}_\rho} \mathbf{C}_{\mathbf{z}_\rho \mathbf{z}_\rho}^{-1}, \quad (6)$$

where $\mathbf{C}_{\tilde{\mathbf{y}}, \mathbf{z}_\rho}$ is the cross-correlation between $\tilde{\mathbf{y}}$ and \mathbf{z}_ρ , while $\mathbf{C}_{\mathbf{z}_\rho \mathbf{z}_\rho}^{-1}$ is the inverse correlation matrix of the compressed representation. In our practical setting, considering the training-set of the *opportunistic* regression sub-task, $\tilde{\mathcal{D}}_{\rho, \text{tr}} = \{(\mathbf{z}_{\rho, i}, \tilde{\mathbf{y}}_i)\}_{i=1}^{N_{\text{tr}}}$, we approximate the L-MMSE by its sample-mean counterpart, i.e., by the LS expression [47].

$$\hat{\Theta}_\rho = \left(\mathbf{Z}_{\rho, \text{tr}}^T \mathbf{Z}_{\rho, \text{tr}} \right)^{-1} \mathbf{Z}_{\rho, \text{tr}}^T \tilde{\mathbf{Y}}_{\text{tr}}, \quad (7)$$

where the design-matrix $\mathbf{Z}_{\rho, \text{tr}} = [\mathbf{z}_{\rho, 1}^T; \dots; \mathbf{z}_{\rho, 1}^T] \in \mathbb{R}^{N_{\text{tr}} \times n_z}$ collects the compressed representations of the training data, while $\tilde{\mathbf{Y}}_{\text{tr}} = [\tilde{\mathbf{y}}_{\rho, 1}^T; \dots; \tilde{\mathbf{y}}_{\rho, 1}^T] \in \mathbb{R}^{N_{\text{tr}} \times n_y}$ collects the associated outcomes.⁴

It is interesting to observe that MMSE estimation of $\tilde{\mathbf{y}}$ is also meaningful from a mutual information perspective, because for any estimator it holds true [41]

$$\mathbb{E}\{|\tilde{\mathbf{y}} - \tilde{\mathbf{y}}_\rho^{(r)}(\mathbf{z}_\rho)|^2\} \geq \frac{n_y}{2\pi e} e^{\frac{2H(\tilde{\mathbf{y}}|\mathbf{z}_\rho)}{n_y}}, \quad (8)$$

where $H(\tilde{\mathbf{y}}|\mathbf{z}_\rho)$ is the conditional entropy of $\tilde{\mathbf{y}}$ given \mathbf{z}_ρ . Thus, considering the mutual information $I(\mathbf{z}_\rho; \mathbf{y}) = H(\mathbf{y}) - H(\mathbf{y}|\mathbf{z}_\rho)$ of the NN output \mathbf{y} and the compact features \mathbf{z}_ρ , by the information processing inequality $I(\mathbf{z}_\rho; \tilde{\mathbf{y}}) \geq I(\mathbf{z}_\rho; \mathbf{y})$ [41], we can conclude that the MSE minimization in (8), corresponds to a maximization of an upper bound of the mutual information $I(\mathbf{z}_\rho; \mathbf{y})$, that is highly related to the final classification accuracy, which we actually target as the final classification performance.

This way, we pass to the classifying network $C(\cdot)$ a reconstructed $\tilde{\mathbf{y}}_{\text{rec}}$ that is optimized from the mutual information perspective: this would possibly lead to the best possible reconstruction from the (final) classification performance point of view (see the interplay of mutual information and cross-entropy [49]). The main steps of the proposed approach are summarized in Algorithm 1.

⁴Note that we are implicitly assuming to cope with an over-determined system (i.e., $N_{\text{tr}} > n_z$) with $\mathbf{z}_{\rho, \text{tr}}$ full-rank matrix, which guarantee the uniqueness of the estimation. Otherwise, although non-unique, the L-MMSE can still be computed via Singular Value Decomposition [47].

Algorithm 1 Opportunistic Information Bottleneck

Input:

Training-set: $\mathcal{D}_{\text{tr}}(\mathbf{x}_n, \mathbf{y}_n)_{n=1}^{N_{\text{tr}}}$

inference network $\mathcal{M} = \mathbf{L}_0 \cup [\mathbf{L}_1, \dots, \mathbf{L}_k]$

The size of the compressed representation n_z .

Output:

The optimal compression transformation \mathbf{A}_ρ

Process:

- 1: Compute a gaussian transformation $h(\cdot)$ of the input data and obtain $\{\tilde{\mathbf{x}}_n\}_{n=1}^{N_{\text{tr}}}$
 - 2: Train the preferred inference network \mathcal{M} on the training-set $\{(\tilde{\mathbf{x}}_n, \mathbf{y}_n)\}_{n=1}^{N_{\text{tr}}}$
 - 3: Run the network to compute $\tilde{\mathbf{y}}_n = \mathbf{L}_0\tilde{\mathbf{x}}_n + \lambda\eta$ and build the regression data-set $\tilde{\mathcal{D}}_{\text{tr}} = \{(\tilde{\mathbf{x}}_n, \tilde{\mathbf{y}}_n)\}_{n=1}^{N_{\text{tr}}}$.
 - 4: Compute the Gaussian IB transformation \mathbf{A}_ρ to infer the regression data-set $\tilde{\mathcal{D}}_{\text{tr}}$ from a compressed representation $\mathbf{z} \in \mathbb{R}^{n_z}$.
 - 5: Collect the transformation for each training sample on the design matrix $\mathbf{Z}_{\rho, \text{tr}} = [\mathbf{z}_{\rho, 1}^T; \dots; \mathbf{z}_{\rho, 1}^T] \in \mathbb{R}^{N_{\text{tr}} \times n_z}$, and the associated outcomes in $[\tilde{\mathbf{y}}_{\rho, 1}^T; \dots; \tilde{\mathbf{y}}_{\rho, 1}^T] \in \mathbb{R}^{N_{\text{tr}} \times n_y}$.
 - 6: compute the L-MMSE estimator using Eq. (6)
 - 7: **return** $\mathbf{A}_\rho, \hat{\Theta}_\rho$.
-

B. CLASSIFICATION ON THE COMPRESSED REPRESENTATION

An alternative, although less flexible and appealing scheme, can be obtained assuming to use the OIB-compressed features \mathbf{z}_ρ to directly feed a classifying NN with a (smaller) input-size n_z , as shown in Fig. 2. In this case, once we learned the OIB compression matrix \mathbf{A}_ρ by the procedure described in the previous section, we have to train (and exploit) a specific (different) classification network $C_\rho(\cdot)$ for each possible compression factor ρ , ending-up with a *bank* of NN classifiers, similarly to what proposed in [33] to decode the output of a matched bank of encoders, which played the role of the OIB-encoders we introduce herein. This further training of the bank of *reduced* networks on the compressed features of the training set, may clearly help to improve the classification performance of the system, which exploits a dedicated NN for each specific compression ratio ρ . The price to be paid is a much higher complexity, both for the training phase, as well as for the overall architecture, which requires a bank of multiple NNs to perform the final classification. In this sense we consider this option less attractive, although it is interesting to investigate the merits and limits of the single-NN architecture we described before.

C. COMPLEXITY ANALYSIS

We detail herein the complexity analysis of the proposed algorithm both for classification training and test. As far as the training phase is concerned, we bear four main costs:

- 1) The training cost of the preferred inference network, which depends on the network architecture and on the data-set.

- 2) The transformation cost associated with the 2D-FFT of the training-set, which is equal to $\mathcal{O}(N_{tr}n_x \log_2 n_x)$.
- 3) The computation cost to determine the GIB compression matrix \mathbf{A}_ρ , which is dominated by the estimation of the conditioned data covariance matrix $\Sigma_{x|y}$ and the computation of the Singular Value Decomposition (SVD) of $\Sigma_{x|y}\Sigma_x^{-1}$, leading to an overall computational complexity $\mathcal{O}(n_x^3 + n_x^2 N_{tr})$.
- 4) The computation cost of the L-MMSE estimator, which requests $\mathcal{O}(n_z^2 N_{tr})$ operations.

Since the computation cost of the L-MMSE estimator and of the compression matrix \mathbf{A}_ρ dominate the cost of the 2D-FFT, the overall computation complexity is similar to that one of a classical PCA, which also requests SVD. Furthermore, the complexity is drastically reduced with respect to the Linear Feature Extractor based on Mutual Information [38], which also exploits a preliminary PCA computation, followed by a gradient based procedure to retrieve the compressed features. During the inference procedure, all the methods based on linear feature extraction need for a computational complexity which is $\mathcal{O}(n_x n_z + n_z n_{\tilde{y}})$ due to the compression and reconstruction. On the other hand, The Variational Information Bottleneck needs an extra NN to perform compression, whose computational complexity would be dramatically higher with respect to performing a matrix multiplication. This confirms that the proposed method would be particularly useful in Edge Inference scenarios, where devices are typically characterized by limited computational capabilities.

IV. EXPERIMENTS

In this section we discuss our experimental results to demonstrate the effectiveness of the proposed approach. Firstly, we describe the data-sets employed in our simulations. Then, for the sake of comparisons, we briefly summarize four alternative FE strategies available in the literature. Finally, we assess the performance of the proposed method.

A. DATA-SET AND NN TRAINING

We considered three data-sets in our experiments:

- The GTSRB data-set [50], composed of 1213 RGB traffic sign images divided in 43 different classes with a size $n_x = 32 \times 32 \text{ px} \times 3 \text{ colors}$. In order to obtain a reliable estimation of the covariance matrices involved in the OIB framework (and in PCA), we considered a data-augmentation factor equal to 5, obtained by applying random rotations to the original images. This procedure led us to deal with an augmented data-set of 6.065 images, that we split in 4.852 for the training-set and 1.213 for the test-set.
- The Euro-SAT data-set [51], composed of 2.7×10^4 RGB images captured by the Sentinel-2 satellite divided in 10 different classes. Also this data-set has images with size $n_x = 32 \times 32 \text{ px} \times 3 \text{ colors}$. We considered 21.600 images for the training-set and 5.400 images for the test-set.

- The MNIST data-set [52], which is composed of 7×10^4 grey-scale images of handwritten digits divided in 10 different classes, with a size $n_x = 28 \times 28 = 784 \text{ px}$. The images are split in 6×10^4 for the training-set and 1×10^4 for the test-set.

We focused on these well known data-sets, because they lead to a training phase of the proposed architecture, with reasonable complexity and computational time. Indeed, as clarified in Section III, the proposed OIB needs to estimate data covariance matrices (and an inverse), e.g., $\Sigma_{x|y}$, and Σ_x^{-1} , with a complexity that scales cubical with the image size: thus, for larger images the computational burden could become problematic for a standard PC, as the one we used to test the proposed approach.⁵ Regarding the training procedure, we employed the well-know *Adam* optimizer [53]. We considered 30 epochs in all our simulations, a learning rate $lr = 10^{-3}$ and a batch-size $|\mathcal{B}| = 32$.

The validity of the multivariate Normal assumptions under the proposed OIB strategy, have been successfully verified by the *Henze-Zirkler* test for Multivariate Normality [54], that has been applied on the 2D-DFT outputs $\tilde{\mathbf{x}}$ that feed the OIB stage.

B. COMPETITIVE APPROACHES

We compare the OIB framework with four alternative FE strategies:

- The well-know Principal Component Analysis (PCA) [39], which is typically used for *unsupervised* LFE, i.e., without taking into account relevance to the specific inference task \mathbf{y} . In this case, for any compression ration $\rho = n_x/n_z$, the features are obtained projecting the data \mathbf{x} on the first n_z eigenvectors of the estimated covariance matrix $\tilde{\Sigma}_x$.
- The Mutual Information Based LFE (MIB-LFE) [38], which proposes to improve PCA by embedding the full PCA on a lower dimensional space by an orthogonal projection matrix \mathbf{A} , which is learnt (together with the classifying NN) by maximizing an approximated bound of the mutual information $I(\mathbf{z}_\rho, \mathbf{y})$. The optimization procedure is based on gradient ascent and Gram-Schmidt ortho-normalization. This method, as the proposed OIB, represents a theoretically principled, task-oriented, supervised LFE, whose effectiveness has been experimentally confirmed and validated in many works [55].
- The nonlinear-FE framework proposed in [33], [34], [36], that exploits a bank of encoders and classifiers, each characterized by a different compression ratio ρ . Specifically, each encoder produces \mathbf{z}_ρ by down-sampling the input image \mathbf{x} by a cascade of convolutional/max-pooling layers and it is jointly trained with an associated CC, which performs the final classification.

⁵The reduction of the complexity to train the system, and to estimate the covariance and precision matrices by exploiting the data-structure and parallel computations, will be investigated in future works.

- The Variational Information Bottleneck (VIB) method [17], which is a well-know *approximation* of the IB principle, based on a variational bound of the IB cost function in (2). Actually, also the VIB can be implemented by considering an architecture with an encoder to extract the features \mathbf{z}_ρ and a NN decoder to classify them [17], [25]. The main disadvantage of the VIB is that, since the output size of the encoder is fixed, similarly to the approach described in [33], it is necessary to train a bank of encoders/classifiers, each one associated to a specific compression factor.

We stress that the first two LFE methods share a similar structure with the proposed OIB in Fig. 2(d). Indeed, while $h(\cdot)$ stands for 2D-DFT in the OIB, it represents a lossless PCA for the other two methods. Then, $\mathbf{A}_\rho = [\mathbf{I}_{n_z}, \mathbf{0}]$ for unsupervised PCA, while \mathbf{A}_ρ is a full $n_x \times n_z$ mixing matrix for MIB-LFE [38]. Consequently, like the OIB, also the first two methods admit a single-NN classification strategy, while the nonlinear-FEs in [33], [34], [36] and the VIB are structurally designed with a bank-of-NNs, as shown in Fig. 2(d)(b).

C. CLASSIFICATION BY A SINGLE-NN

A single-NN is trained by the procedure described in Section III. Thus, for any different compression ratio ρ , the input of this NN is obtained from the compact representation \mathbf{z}_ρ , by the LS-estimator in (7). Employing CLT arguments, we can assume that, also for unsupervised PCA and MIB-LFE, the compact representation \mathbf{z}_ρ and $\tilde{\mathbf{y}}$ are jointly-Gaussian. Thus, L-MMSE expansion from \mathbf{z}_ρ to $\tilde{\mathbf{y}}$ turns out to be optimal, as for the OIB. Conversely, given the non-linear structure of the compressive encoder in [33] and of the encoders in the VIB [17], we cannot assume the joint Gaussianity of \mathbf{z}_ρ and $\tilde{\mathbf{y}}$, and an L-MMSE reconstruction of the CC input would be sub-optimal.

To explore the generality and potentials of the proposed approach, we did tests by exploiting both Shallow and Convolutional Neural Networks (SNNs and CNNs). Actually a SNN has been tested only on the MNIST data-set since, differently from the GTSRB data-set, it can be reliably classified also by simple SNN models, thanks to the low complexity of the underlying classification task.

Tables 2, 3 and 4 report the details for baseline NNs that we employed to implement the overall classification network shown in Fig. 2, without exploiting the proposed OIB approach. Tables also report the Multiplication and Accumulation Complexity (MAC), associated to each layer. All the NNs employ a $\text{relu}(\cdot)$ activation function.

We remind that the goal of the proposed OIB is to implement classification by a *size-tunable* feature extraction, optimally trading performance for complexity. The *single-network* architectures summarized in Fig. 2(d), fulfil this goal by introducing a 2D-DFT processing, and replacing layer \mathbf{L}_0 by the compressing matrix \mathbf{A}_ρ and the expanding matrix Θ_ρ . Table 6 reports the MACs of the proposed architecture, for different compression ratios $\rho = n_x/n_z$,

TABLE 2. Architecture of the Shallow network used to classify the MNIST data-set.

Layer	Input Features	Output Features	Complexity [MACs]
\mathbf{L}_0	784	256	200704
\mathbf{L}_1	256	128	32768
\mathbf{L}_2	128	64	8192
\mathbf{L}_3	64	16	1024
\mathbf{L}_4	16	10	160
Total Complexity [MACs]			242848

TABLE 3. Architecture of the Convolutional Neural Network used to classify the GTSRB data-set. Each Convolutional stage considers a 2×2 stride.

Layer	Type	Output Shape	Kernel	[MACs]
\mathbf{L}_0	Strided Conv2D	$16 \times 16 \times 1$	9×9	37969
\mathbf{L}_1	Strided Conv2D	$8 \times 8 \times 32$	3×3	16200
\mathbf{L}_2	Strided Conv2D	$4 \times 4 \times 64$	3×3	225792
\mathbf{L}_3	Conv2D	$4 \times 4 \times 128$	3×3	294912
\mathbf{L}_4	Linear	43	N/D	88064
Total Complexity [MACs]				662937

TABLE 4. Architecture of the Convolutional Neural Network used to classify the Euro-SAT data-set. The first stage considers a 2×2 stride, while the max-pooling stages consider a 2×2 window.

Layer	Type	Output Shape	Kernel	[MACs]
\mathbf{L}_0	Strided Conv2D	$32 \times 32 \times 1$	3×3	24300
\mathbf{L}_1	Conv2D+MaxPool	$16 \times 16 \times 64$	3×3	518400
\mathbf{L}_2	Conv2D+MaxPool	$4 \times 4 \times 128$	3×3	14450688
\mathbf{L}_3	Conv2D+MaxPool	$2 \times 2 \times 256$	3×3	10616832
\mathbf{L}_4	Linear	256	N/D	1048576
\mathbf{L}_5	Linear	10	N/D	2560
Total Complexity [MACs]				26661356

TABLE 5. Saving in terms of computational complexity for both the UE and the ES for the shallow NN architecture (MNIST data-set).

n_z	Comp. [MACs]	Class. [MACs]	Comp. Save [%]
10	18080	44704	74.13
20	25920	47264	69.84
30	33760	49824	65.56
40	41600	52384	61.27
50	49440	54944	56.99
60	57280	57504	52.70
70	65120	60064	48.42
80	72960	62624	44.13
90	80800	65184	39.85
100	88640	67744	35.56

distinguishing between the compressive FE and the final classification, which exploits only layers \mathbf{L}_1 - \mathbf{L}_4 of the original network.

The FE cost is due to the 2D-FFT and the subsequent multiplication by \mathbf{A}_ρ , while the *classification cost* is due to multiplication with the matrix Θ_ρ , and the computations through the NN layers \mathbf{L}_1 - \mathbf{L}_4 .⁶

⁶The proposed architecture is particularly valuable when the FE is performed at a given device, while the final classification at a central server equipped with more powerful processing capabilities, such as it happens in cloud-based applications, and Edge Machine Learning [13], [36], [56] frameworks.

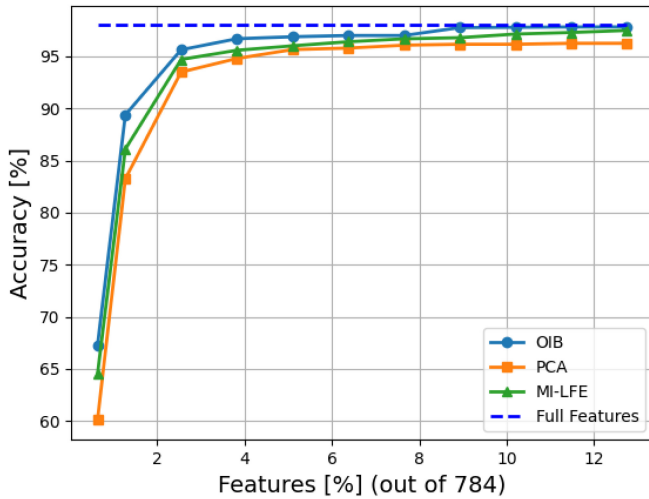


FIGURE 3. Performance comparison with feature's re-expansion and a single-NN architecture.

Fig. 3 plots the accuracy performance of the single-NN, as a function of the number n_z of extracted features, i.e., of the compression ratio $\rho = 784/n_z$. It is clear how the proposed approach outperforms both the MIB-LFE [38] and the unsupervised PCA compression [39]. PCA, as expected, shows the worst performance because it retrieves features that are the most informative for image *reconstruction* purposes, without taking into account their relevance for the actual classification task. However, even though [38] would allow to extract features in a supervised and task-oriented manner, its adaptation to such a single-NN architecture, with tunable compression ratio, produces higher performance degradation with respect to the proposed OIB.

D. ANALYSIS OF COMPRESSION/RELEVANCE TRADE-OFFS

In the previous section, we testified the effectiveness of our method as a supervised linear feature extractor. However, as already pointed out in Section II-D, the Information Bottleneck method, that the OIB takes inspiration from, is focused on a compression/relevance trade-off. The compression term is associated to the mutual information $I(\tilde{\mathbf{x}}, \mathbf{z})$ (cf eq. (2)), that the IB aims to minimize in order to retrieve a compact representation of the input, although still *relevant* for the outcome of the learning task. Taking in mind the relationship $I(\tilde{\mathbf{x}}, \mathbf{z}) = H(\mathbf{z}) - H(\mathbf{z}|\tilde{\mathbf{x}})$ we note that minimizing the IB cost is somehow related to reducing the entropy of the compressed representation $H(\mathbf{z})$, which represents a lower-bound on the minimum number of bits required to encode the source without loss of information [41]. Thus, in Edge-Assisted Goal-Oriented Communications scenarios, where we aim to transmit the minimum amount of data to pursue a *Goal*, minimizing latency and energy consumption [8], it makes sense to analyze the task performance (e.g., the




FIGURE 4. $H(\mathbf{z})$ as a function of the number of extracted components for PCA, CCA and OIB.

correct classification rate) as a function of the entropy of the compressed representation.

Furthermore, although the loading coefficients α_i in (4) of the GIB solution don't carry any improvement in terms of learning performance with respect to any other loading set, as clarified in Section II-D, they turn out useful in a data compression perspective. To better highlight this fact, we compare our OIB approach with PCA and CCA. Specifically, CCA projects the inputs $\tilde{\mathbf{x}}$ in the same directions of the OIB solution, giving however the same load to each component (i.e., $\{\alpha_i\}_{i=1}^{n_z} = 1$). We did our comparisons on the MNIST data-set, trained on the shallow network architecture considered in the previous section. The network considered for PCA has been trained on the original images of the data-set, while for CCA and OIB we used the same network architecture trained on the 2D-DFT of the images. In Fig. 4 we show the entropy $H(\mathbf{z})$ as a function of the number of components for the considered approaches. As expected, OIB leads to a lower entropy on the compressed representation with respect to the competitors. This means that it is possible to find an encoding rule that allows to minimize the number of bits to represent \mathbf{z} , and associated savings in terms of transmission resources.

In Fig. 5 we compare the three compression philosophies considering the informativeness of the compressed representation with respect to the classification task with the same entropy amount $H(\mathbf{z})$. Using Central Limit Theorem arguments, we can assume that the compressed representation \mathbf{z} has a multivariate Normal distribution. Thus, we computed the entropy considering the formula $H(\mathbf{z}) = \frac{1}{2} \log((2\pi e)^{n_z} |\Sigma_{\mathbf{z}}|)$ [41], where the covariance matrix has been estimated by sample-mean.⁷ As expected, our OIB approach performs better with respect to the competitive approaches, since it is the only one specifically focused on the compression maximization under learning performance

⁷Note, since the differential entropy is scale-dependant, to make fair comparisons we normalized the power of \mathbf{z} for all the considered approaches.




FIGURE 5. Accuracy as a function of $H(\mathbf{z})$ for PCA, CCA and OIB.

constraints. As in the previous case, the unsupervised PCA approach shows the worst performance. Regarding CCA, taking in mind that the loadings do not affect the task performance, if we consider the same compressed size n_z we would have the same performance of the OIB approach. However, as shown in Fig. 4, CCA leads to a solution characterized by a higher complexity of representation (captured by $H(\mathbf{z})$) for a fixed number of components. To reach the same complexity with the OIB transformation, we have to consider more components, with a consequent improvement in terms of accuracy, when comparing the schemes at the same entropy.

E. CLASSIFICATION WITH A BANK-OF-NNS

In this section we test the performance using the architecture in Fig. 2, which exploits multiple classifying NNs, each one with an input size that is matched to the size n_z of the extracted features \mathbf{z}_ρ , without re-expanding them to $\tilde{\mathbf{y}}_\rho^{(r)}$. For the CE compression strategy, we employed the joint CE/CC training procedure described in [33]. To perform fair comparisons, for the CE compression strategy we employed a bank of CEs, whose MACs complexity is similar to the OIB method. This has been done considering a mixed architecture composed of two convolutional layers followed by a fully connected layer, used to adjust the output size to the desired size n_z . The number of channel at the output of the first layer has been adjusted in order to match the desired complexity. The architecture of the encoder for the sizes n_z is reported in Table 6. In the table we also report the comparison between the computational complexity of the proposed encoder and the Opportunistic Information Bottleneck.

We assessed our performance considering the GTSRB data-set [50] and we compare the proposed approach with the Method proposed in [38] and with the well-know Variational Information Bottleneck (VIB) [17]. Indeed, these two methods based on Information Theory arguments [41], and they are specifically tailored to classify directly on the compressed features without require the re-expansion.

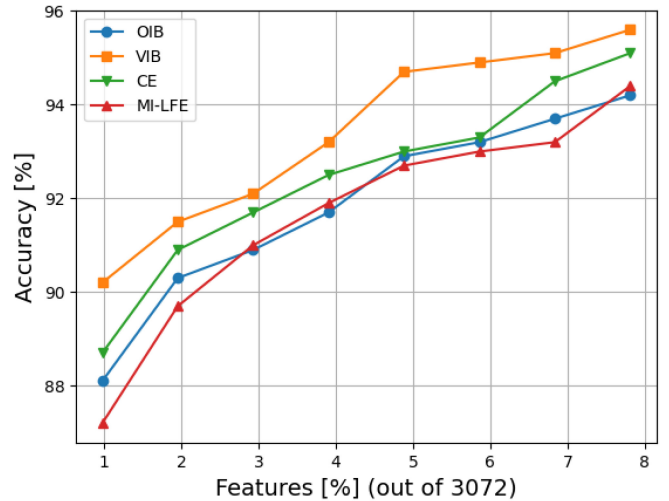


FIGURE 6. Performance comparison with the *bank-of-NN* architecture.

Fig. 6, which shows the accuracy versus the percentage of compressed features, witnesses that the Variational Information Bottleneck, performs slightly better with respect to the other methods. VIB represents an alternative way to train a Convolutional Encoder, which is based on a more principled approximation of the Information Bottleneck principle [17]. Furthermore, differently from what we will show in Section IV-F, the VIB performs better with respect to the OIB framework. This represents a benefit of the joint training procedure of the compression and classification networks based on the Variational Bound [17], which represents a theoretically-grounded approximation of the IB objective function. Conversely, there are no substantial differences between the proposed approach, the CE-based compression and the Linear Feature Extractor based on mutual information. As far as the CE-based compression is concerned, the CEs in [33], can be interpreted as a (non-linear) heuristic mimicking of the Information Bottleneck principle: somehow this fact is confirmed by the similar performance with respect to the proposed OIB, which is also a sub-optimal approximation of the IB, although more theoretically principled. The proposed OIB in such a multi-NN implementation, performs similarly to the MIB-LFE in [38], for any compression ratio. Note, that in this case we are not re-expanding the feature size. i.e., we are working for each compression ratio with the structure underlying the original MIB-LFE design. Thus, in this case OIB and MIB-LFE performs similarly because they share a quite similar structure and both the methods propose a way to approximate the maximization of the mutual information between the compact representation \mathbf{z}_ρ and the learning task output \mathbf{y} . Specifically, [38] proposes an approximation of the $I(\mathbf{z}_\rho, \mathbf{x})$ that analytically relies on a statistical independence assumptions of the data, which is only approximated by the full-PCA pre-processing for non-Gaussian features. Further approximations are introduced in [38] to compute the *negentropy* in the objective function. In the proposed OIB, we opportunistically embedded in the

TABLE 6. Variable encoder architecture for the different compression factors. $C_{1,2}$ represents the two convolutional layers, while L_1 is the linear layer. For the convolutional layers we considered 3×3 convolutions and an input image with size $32 \times 32 \times 3$ px.

n_z	Output shape C_1	Output shape C_2	Output shape L_1	OIB [MACs]	CE [MACs]
30	$15 \times 15 \times 15$	$7 \times 7 \times 1$	30	101408	105386
60	$15 \times 15 \times 28$	$7 \times 7 \times 1$	60	193568	196917
90	$15 \times 15 \times 41$	$7 \times 7 \times 1$	90	285728	288448
120	$15 \times 15 \times 54$	$7 \times 7 \times 1$	120	377888	379979
150	$15 \times 15 \times 67$	$7 \times 7 \times 1$	150	470048	471509
180	$15 \times 15 \times 80$	$7 \times 7 \times 1$	180	562208	563040
210	$15 \times 15 \times 93$	$7 \times 7 \times 1$	210	654368	654571
240	$15 \times 15 \times 106$	$7 \times 7 \times 1$	240	746528	746102

original inference task, another IB sub-problem, that we are able to optimally solve in closed form by the GIB. The main assumption here is that the good performance of the surrogate linear regression task solved by the GIB, induces also an almost optimal performance of the associated classification task. This intuitive explanation, which has been partially supported by information theoretic arguments in Section III-A, is confirmed by the comparable performance with the MIB-LFE approach. The nice part of the OIB is that it is a bit easier to be trained and furthermore, shows better performance in the single-NN architecture described in the previous section.

F. COMPARISONS WITH THE VARIATIONAL INFORMATION BOTTLENECK

To further prove the effectiveness of the proposed approach, we compare it with the well-known Variational Information Bottleneck (VIB) method [17]. We made our experimental validation considering the network architectures reported in Table 3 and Table 4, trained on the GTSRB [50] and the Euro-SAT data-sets [51]. VIB is a well-know approximation of the Information Bottleneck principle, based on a variational bound of the IB cost function, which can be easily implemented considering an encoder/decoder architecture.

As already pointed out, the main disadvantage of the VIB is that, since the output size of the encoder is fixed, similarly to the approach described in [33], it is necessary to train a bank of encoders/classifiers, each one associated to a specific compression factor ρ . Thus, to fairly compare the two approaches, we adapted the VIB formulation to work with the single-architecture setting, considering the following training procedure:.

- 1) We trained the CNN reported in Table 3 on the GTSRB data-set the CNN reported in Table 4 on the Euro-SAT data-set.
- 2) We trained the VIB networks considering the procedure described in [17] for all the possible number of features. We implemented an encoder with the following structure: Conv2D(F) \rightarrow Linear($256 \times F, 2 \times n_z$), where F is the number of output filters. The number of filters has been changed for the different compression factors to maintain the same complexity,




FIGURE 7. Comparisons between the proposed approach and the Variational Information Bottleneck on the GTSRB data-set.

in terms of MACs, of the OIB compression. As far as the decoder is concerned, we implemented the same architecture reported in Tables 3 and 4, ignoring the layer L_0 . We considered a linear layer between the encoder and the decoder in order to expand the encoder to the expected input size of the layer L_1 , i.e., $16 \times 16 \times 1$ for the network reported in Table 3 and $32 \times 32 \times 1$ for the network reported in Table 4.

- 3) After the end of the training procedure, we got rid of the decoder and we performed the same procedure reported in Section III in order to get the LS estimators Θ_ρ , which allows us to reconstruct the input of the intermediate layer L_1 from the compressed representation obtained through the VIB encoder.
- 4) We re-trained the classification networks with re-expansion for each possible compression factor ρ .

In order to fairly compare the proposed OIB with the VIB, we also re-trained the OIB classification network for each possible size of n_z (i.e., for each ρ), by re-expanding the compressed features \mathbf{z}_ρ with the proper LS estimator $\hat{\Theta}_\rho$.

In Fig. 7 we report the average accuracy versus the percentage of compressed features on the test-set of the GTSRB data-set. We firstly note that, as expected, the re-training




FIGURE 8. Comparisons between the proposed approach and the Variational Information Bottleneck on the Euro-SAT data-set.

procedure on the expanded features allows to highly improve the performance for the proposed OIB, which slightly outperforms the classification accuracy of the VIB. However, a classical VIB approach, with a tunable FE size, naturally leads to a higher complexity of the system, since it is necessary to consider a classification architecture with multiple NNs, each one trained on a compressed \mathbf{z}_ρ with a specific number of features n_z (i.e., a specific $\rho = n_x/n_z$). To avoid such implementation complexity, another possible approach is to consider a single classifying NN that is trained on a set of $\tilde{\mathbf{y}}_\rho^{(r)}$, which are re-expanded versions of \mathbf{z}_ρ , having different compression ratios, i.e., $\rho \in \{\rho_1, \rho_2, \dots, \rho_k\}$ and, consequently, a different number of features n_z . Specifically, during the training phase, we randomly compressed the training data considering a uniform discrete distribution for the possible values of ρ . Then, we tested this *average* architecture, both for the VIB and the OIB design, at different (fixed) levels of compression. Fig. 7 shows that such *average* classifying NN suffers some performance degradation with respect to the classification based on the multiple-NNs architectures. However, it also grants a noticeable performance improvement with respect to the single-NN architecture without re-training and, noticeably, also in this case the quite simple and elegant OIB framework outperforms the VIB. The same considerations hold also for the Euro-SAT data-set, as witnessed by Fig. 8.

G. SENSITIVITY ANALYSIS

This section analyzes the sensitivity of the proposed method to the neural network structure, focusing on convolutional and fully connected neural networks. Let us consider the case of data classification with feature re-expansion (through the L-MMSE estimator).

We firstly remark that we are approximating the first linear layer of our network \mathbf{L}_0 , by a (low-rank) matrix factorization expressed by $\tilde{\mathbf{L}}_0 = \mathbf{A}_\rho \Theta_\rho$. Thus, the robustness




FIGURE 9. Accuracy on the test-set of the MNIST data-set for different number of components n_z on shallow and convolutional neural networks.

of our algorithm is related to the error we induce by said approximation, which naturally gets worse using smaller sizes n_z (i.e., lower-rank approximation) for the compressed representation.

However, it is important to note that the possible classification degradation depends also on another factor, i.e., the capability of the last network layers $[\mathbf{L}_1, \dots, \mathbf{L}_n]$ to be robust, or capable to adapt, to additive noise ignited in the first network layer \mathbf{L}_0 .

To clarify this point, we show in the following the simulation results obtained exploiting both fully connected and convolutional neural networks, evaluating the accuracy degradation with respect to the originally trained (full) NN, for different number n_z of extracted features. Specifically, we used the MNIST data-set, and the system has been retrained with the procedure detailed in Section III-A. Results are shown for a set of possible feature's size $n_z \in [5, 10, 15, \dots, 70]$.

For the fully connected neural network, the first layer is characterized by a matrix $\mathbf{L}_0 \in \mathbb{R}^{784 \times 196}$, while for the convolutional architecture we employed a 2×2 strided convolution to down-sample the input data from $28 \times 28 = 784$ px to 196 px, ending up with an equivalent (block diagonal) matrix $\mathbf{L}_0 \in \mathbb{R}^{784 \times 196}$.

As expected, Fig. 9 shows that the OIB-approximated NN provides better classification accuracy as we increase the size n_z of the compressed representation \mathbf{z}_ρ , for any fixed architecture and (re-) training strategy, due to the fact that the low-rank approximation $\tilde{\mathbf{L}}_0 = \mathbf{A}_\rho \Theta_\rho$ gets better. Furthermore, Fig. 9 shows that the convolutional neural network seems to be more sensitive to the approximation of $\tilde{\mathbf{L}}_0$. This higher sensitivity of CNN can be explained by considering that reconstructing $\tilde{\mathbf{y}}$ by $\tilde{\mathbf{L}}_0 \tilde{\mathbf{x}}$, has an effect similar to adding noise on the inputs of the next network layer \mathbf{L}_1 , and is well known that CNNs have robustness issues in this case [57]. However, this *in-layer* noise

effect can be mitigated by re-training the last network layers $[\mathbf{L}_1, \dots, \mathbf{L}_n]$ on the reconstructed data, i.e., on the output of \mathbf{L}_0 , paying the price to introduce some extra training complexity. Actually, by this retraining, the CNN performance for the proposed classification data-set are equivalent to those of the fully-connected architecture, as witnessed the blue curve in Fig. 9. Although, this is just a simple example, it is reasonable that properly retraining also other NN architectures, it is possible to make the approximation error introduced by OIB quite insensitive to the specific NN, as the NN layers $[\mathbf{L}_1, \dots, \mathbf{L}_n]$ should be capable to capture by re-training the mutual information $I(\tilde{\mathbf{y}}_{\rho}^{(r)}, \hat{\mathbf{y}})$ if they were capable to capture $I(\tilde{\mathbf{y}}, \hat{\mathbf{y}})$.

V. CONCLUSION AND FUTURE WORK

We presented a new approach to opportunistically exploit the closed-form solution of the Gaussian Information Bottleneck (GIB) in a general inference task, to enable a tunable and effective supervised feature extraction, highlighting the specific merit of the GIB with respect to CCA, as the best compact representation of the (same) extracted features. Experimental results on an image classification task testifies the effectiveness of the proposed approach and encourages to further investigate this research line. In particular, the proposed formulation seems particularly attractive any time there is the need, or the opportunity, to perform inference by a single-NN, with a number of features that may change in time as a function of the system resources, such as hardware, energy, storage, communications, and computations. In this view, future works may include the employment of the proposed compression strategy in dynamic resource allocation strategies for edge-assisted goal-oriented communications scenarios (see, e.g., [33], [34], [36]). Furthermore, it may be interesting to apply the OIB formulation to different inference tasks, to better assess its generality. Finally, the investigation of possible interplays between OIB and Variational IB formulations represents another appealing research line.

APPENDIX

MAXIMIZATION OF MUTUAL INFORMATION UNDER GAUSSIAN ASSUMPTIONS

Let us define the following optimization problem

$$\begin{aligned} \max_{\Phi_s} \quad & I(\mathbf{z}, \mathbf{y}) \\ \text{s.t.} \quad & \mathbf{z} = \mathbf{M}, \Sigma_{\epsilon} \end{aligned} \quad (9)$$

where \mathbf{x} and \mathbf{y} are characterized by a multivariate Normal distribution and ϵ is also Normal and statistically independent from both \mathbf{x} and \mathbf{y} . Since $I(\mathbf{z}, \mathbf{y}) = H(\mathbf{z}) - H(\mathbf{z}|\mathbf{y})$, recalling the expression of the entropy for Multivariate Gaussian variables [41], the objective function becomes

$$I(\mathbf{z}, \mathbf{y}) = \log(|\mathbf{M}\Sigma_{\mathbf{x}}\mathbf{M}^t + \Sigma_{\epsilon}|) - \log(|\mathbf{M}\Sigma_{\mathbf{x}|\mathbf{y}}\mathbf{M}^t + \Sigma_{\epsilon}|).$$

Proceeding as in [12], nulling the derivative $\frac{\partial I(\cdot, \cdot)}{\partial \mathbf{M}}$ we obtain the following equation

$$\begin{aligned} [\mathbf{M}\Sigma_{\mathbf{x}}\mathbf{M}^t + \Sigma_{\epsilon}]^{-1}\mathbf{M}\Sigma_{\mathbf{x}} - [\mathbf{M}\Sigma_{\mathbf{x}|\mathbf{y}}\mathbf{M}^t + \Sigma_{\epsilon}]^{-1}\mathbf{M}\Sigma_{\mathbf{x}|\mathbf{y}} &= 0 \\ [\mathbf{M}\Sigma_{\mathbf{x}|\mathbf{y}}\mathbf{M}^t + \Sigma_{\epsilon}][\mathbf{M}\Sigma_{\mathbf{x}}\mathbf{M}^t + \Sigma_{\epsilon}]^{-1}\mathbf{M} &= \mathbf{M}\Sigma_{\mathbf{x}|\mathbf{y}}\Sigma_{\mathbf{x}}^{-1}. \end{aligned} \quad (10)$$

Thus, $\mathbf{M}\Sigma_{\mathbf{x}|\mathbf{y}}\Sigma_{\mathbf{x}}^{-1}$ must reside in the space generated by the rows of \mathbf{M} , and it is composed by left eigenvectors of $\Sigma_{\mathbf{x}|\mathbf{y}}\Sigma_{\mathbf{x}}^{-1}$. This means we can write the optimal transformation matrix as $\mathbf{M} = \mathbf{W}\mathbf{V}$, where the rows of \mathbf{V} are composed by the left eigenvectors of $\Sigma_{\mathbf{x}|\mathbf{y}}\Sigma_{\mathbf{x}}^{-1}$, while $\mathbf{W} = \text{diag}(w_i), i = 1, \dots, n_x$ is the loadings matrix.

By definition of left eigenvector, $\mathbf{V}\Sigma_{\mathbf{x}|\mathbf{y}}\Sigma_{\mathbf{x}}^{-1} = \mathbf{D}\mathbf{V}$, where $\mathbf{D} = \text{diag}(\lambda_i), i = 1, \dots, n_x$. Substituting in eq (9) we obtain the following derivations

$$\begin{aligned} [\mathbf{W}\mathbf{D}\mathbf{V}\Sigma_{\mathbf{x}}\mathbf{V}^t\mathbf{W}^t + \Sigma_{\epsilon}][\mathbf{W}\mathbf{V}\Sigma_{\mathbf{x}}\mathbf{V}^t\mathbf{W}^t + \Sigma_{\epsilon}]^{-1}\mathbf{W} &= \mathbf{W}\mathbf{D} \\ [\mathbf{W}\mathbf{D}\mathbf{R}\mathbf{W}^t + \Sigma_{\epsilon}][\mathbf{W}\mathbf{R}\mathbf{W}^t + \Sigma_{\epsilon}]^{-1}\mathbf{W} &= \mathbf{W}\mathbf{D}, \end{aligned}$$

where $\mathbf{R} = \mathbf{V}\Sigma_{\mathbf{x}}\mathbf{V}^t$. Now, pre-multiplying by \mathbf{W}^{-1} , post-multiplying by $\mathbf{W}^{-1}(\mathbf{W}\mathbf{R}\mathbf{W}^t + \Sigma_{\epsilon})\mathbf{W}$ and re-arranging, we end up with the following equation

$$\mathbf{W}^{-1}\Sigma_{\epsilon}\mathbf{W} = \mathbf{D}\mathbf{W}^{-1}\Sigma_{\epsilon}\mathbf{W}. \quad (11)$$

Assuming that \mathbf{D} is not the identity matrix (which would imply that \mathbf{x} and \mathbf{y} are independent), eq. (11) is verified if and only if $\Sigma_{\epsilon} = \mathbf{0}$, for any possible loading matrix \mathbf{W} , which does not influence the optimization.

REFERENCES

- [1] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, "Generative adversarial networks: An overview," *IEEE Signal Process. Mag.*, vol. 35, no. 1, pp. 53–65, Jan. 2018.
- [2] K. R. Chowdhary, "Natural language processing," in *Fundamentals of Artificial Intelligence*. New Delhi, India: Springer, 2020, pp. 603–649.
- [3] C. Szegedy, A. Toshev, and D. Erhan, "Deep neural networks for object detection," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, Red Hook, NY, USA, 2013, pp. 2553–2561.
- [4] Z. Li et al., "Edge-assisted v2x motion planning and power control under channel uncertainty," *IEEE Trans. Veh. Technol.*, vol. 72, no. 7, pp. 9641–9646, Jul. 2023.
- [5] E. C. Strinati et al., "6G: The next frontier: From holographic messaging to artificial intelligence using subterahertz and visible light communication," *IEEE Veh. Technol. Mag.*, vol. 14, no. 3, pp. 42–50, Sep. 2019.
- [6] H. Kurunathan, H. Huang, K. Li, W. Ni, and E. Hossain, "Machine learning-aided operations and communications of unmanned aerial vehicles: A contemporary survey," *IEEE Commun. Surveys Tuts.*, vol. 26, no. 1, pp. 496–533, 1st Quart., 2024.
- [7] Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, "Edge intelligence: Paving the last mile of artificial intelligence with edge computing," *Proc. IEEE*, vol. 107, no. 8, pp. 1738–1762, Aug. 2019.
- [8] E. Calvanese Strinati and S. Barbarossa, "6G networks: Beyond Shannon towards semantic and goal-oriented communications," *Comput. Netw.*, vol. 190, May 2021, Art. no. 107930.
- [9] *IMT Traffic Estimates for the Years 2020 to 2030*, ITU-Rec. M. 2370-0, Int. Telecommun. Union, Geneva, Switzerland, 2015.
- [10] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948.
- [11] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annu. Allerton Conf. Commun., Control Comput.*, 1999, pp. 368–377.

- [12] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, "Information bottleneck for gaussian variables," *J. Mach. Learn. Res.*, vol. 6, pp. 165–188, Dec. 2005.
- [13] F. Pezone, S. Barbarossa, and P. Di Lorenzo, "Goal-oriented communication for edge learning based on the information bottleneck," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Singapore, 2022, pp. 8832–8836.
- [14] A. Zaidi, I. Estella-Aguerrri, and S. Shamai, "On the information bottleneck problems: Models, connections, applications and information theoretic views," *Entropy*, vol. 22, no. 2, p. 151, 2020.
- [15] J. Wang et al., "Information bottleneck-based interpretable multitask network for breast cancer classification and segmentation," *Med. Image Anal.*, vol. 83, Jan. 2023, Art. no. 102687.
- [16] F. Zhang, Y. Zheng, J. Wu, X. Yang, and X. Che, "Multi-rater label fusion based on an information bottleneck for fundus image segmentation," *Biomed. Signal Process. Control*, vol. 79, Jan. 2023, Art. no. 104108.
- [17] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, "Deep variational information bottleneck," in *Proc. Int. Conf. Learn. Represent.*, 2017, pp. 1–19.
- [18] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*.
- [19] M. M. Mahvari, M. Kobayashi, and A. Zaidi, "Scalable vector gaussian information bottleneck," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2021, pp. 37–42. [Online]. Available: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85115047762&doi=10.1109>
- [20] I. Estella Aguerri and A. Zaidi, "Distributed information bottleneck method for discrete and gaussian sources," in *Proc. Int. Zurich Seminar Inf. Commun. (IZS)*, 2018, pp. 35–39.
- [21] A. A. Alemi, I. Fischer, and J. V. Dillon, "Uncertainty in the variational information bottleneck," 2018, *arXiv:1807.00906*.
- [22] K. Ahuja et al., "Invariance principle meets information bottleneck for out-of-distribution generalization," in *Proc. 35th Conf. Neural Inf. Process. Syst.*, 2021, pp. 3438–3450.
- [23] T. Wu, H. Ren, P. Li, and J. Leskovec, "Graph information bottleneck," in *Proc. 34th Conf. Neural Inf. Process. Syst.*, 2020, pp. 20437–20448.
- [24] J. Yu, T. Xu, Y. Rong, Y. Bian, J. Huang, and R. He, "Graph information bottleneck for subgraph recognition," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–19.
- [25] J. Shao, Y. Mao, and J. Zhang, "Learning task-oriented communication for edge inference: An information bottleneck approach," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 1, pp. 197–211, Jan. 2022.
- [26] J. Shao, Y. Mao, and J. Zhang, "Task-oriented communication for multidevice cooperative edge inference," *IEEE Trans. Wireless Commun.*, vol. 22, no. 1, pp. 73–87, Jan. 2023.
- [27] L. D. Chamain, S. Qi, and Z. Ding, "End-to-end image classification and compression with variational autoencoders," *IEEE Internet Things J.*, vol. 9, no. 21, pp. 21916–21931, Nov. 2022.
- [28] F. Liu, W. Tong, Z. Sun, and C. Guo, "Task-oriented semantic communication systems based on extended rate-distortion theory," 2022, *arXiv:2201.10929*.
- [29] O. Alhussein, M. Wei, and A. Akhavan, "Dynamic encoding and decoding of information for split learning in mobile-edge computing: Leveraging information bottleneck theory," 2023, *arXiv:2309.02787*.
- [30] H. Li et al., "Task-oriented communication with out-of-distribution detection: An information bottleneck framework," 2023, *arXiv:2305.12423*.
- [31] S. Xie, S. Ma, M. Ding, Y. Shi, M. Tang, and Y. Wu, "Robust information bottleneck for task-oriented communication with digital modulation," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 8, pp. 2577–2591, Aug. 2023.
- [32] M. Zhu, C. Feng, C. Guo, N. Jiang, and O. Simeone, "Information bottleneck-inspired type based multiple access for remote estimation in IoT systems," *IEEE Signal Process. Lett.*, vol. 30, pp. 403–407, 2023.
- [33] F. Binucci, P. Banelli, P. Di Lorenzo, and S. Barbarossa, "Adaptive resource optimization for edge inference with goal-oriented communications," *EURASIP J. Adv. Signal Process.*, vol. 2022, no. 1, p. 123, 2022.
- [34] F. Binucci, P. Banelli, P. Di Lorenzo, and S. Barbarossa, "Dynamic resource allocation for multi-user goal-oriented communications at the wireless edge," in *Proc. 30th Eur. Signal Process. Conf. (EUSIPCO)*, Belgrade, Serbia, 2022, pp. 697–701.
- [35] D. Wen et al., "Task-oriented sensing, computation, and communication integration for multi-device edge AI," *IEEE Trans. Wireless Commun.*, vol. 23, no. 3, pp. 2486–2502, Mar. 2024.
- [36] F. Binucci, P. Banelli, P. D. Lorenzo, and S. Barbarossa, "Multi-user goal-oriented communications with energy-efficient edge resource management," *IEEE Trans. Green Commun. Netw.*, vol. 7, no. 4, pp. 1709–1724, Dec. 2023.
- [37] C. M. Bishop and N. M. Nasrabadi, *Pattern Recognition and Machine Learning*, vol. 4. New York, NY, USA: Springer, 2006.
- [38] J. M. Leiva-Murillo and A. Artés-Rodríguez, "Maximization of mutual information for supervised linear feature extraction," *IEEE Trans. Neural Netw.*, vol. 18, no. 5, pp. 1433–1441, Sep. 2007.
- [39] H. Abdi and L. J. Williams, "Principal component analysis," *Wiley Interdiscipl. Rev., Comput. Statist.*, vol. 2, no. 4, pp. 433–459, 2010.
- [40] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [41] T. M. Cover, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 1999.
- [42] Y. Ugur, I. E. Aguerri, and A. Zaidi, "A generalization of blahut-arimoto algorithm to compute rate-distortion regions of multiterminal source coding under logarithmic loss," in *Proc. IEEE Inf. Theory Workshop (ITW)*, 2017, pp. 349–353. [Online]. Available: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85046335084&doi=10.1109>
- [43] D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor, "Canonical correlation analysis: An overview with application to learning methods," *Neural Comput.*, vol. 16, no. 12, pp. 2639–2664, 2004.
- [44] S. Bandyopadhyay and S. N. Lahiri, "Asymptotic properties of discrete fourier transforms for spatial data," *Sankhyā: Indian J. Statist., Ser. A*, vol. 71, pp. 221–259, Aug. 2009.
- [45] J. S. Lim, *Two-Dimensional Signal and Image Processing*. Englewood Cliffs, NJ, USA: Prentice Hall, 1990.
- [46] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016, [Online]. Available: <http://www.deeplearningbook.org>.
- [47] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1993.
- [48] H. D. Vinod, "Canonical ridge and econometrics of joint production," *J. Econometr.*, vol. 4, no. 2, pp. 147–166, 1976.
- [49] K. G. Larkin, "Reflections on Shannon information: In search of a natural information-entropy for images," 2016, *arXiv:1609.01117*.
- [50] J. Stallkamp et al., "The German traffic sign recognition benchmark: A multi-class classification competition," in *Proc. Int. Joint Conf. Neural Netw.*, 2011, pp. 1453–1460.
- [51] P. Helber, B. Bischke, A. Dengel, and D. Borth, "EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 12, no. 7, pp. 2217–2226, Jul. 2019.
- [52] L. Deng, "The MNIST database of handwritten digit images for machine learning research," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 141–142, Nov. 2012.
- [53] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, San Diego, CA, USA, 2015, pp. 1–15.
- [54] N. Henze and B. Zirkler, "A class of invariant consistent tests for multivariate normality," *Commun. Statist. Theory Methods*, vol. 19, no. 10, pp. 3595–3617, 1990.
- [55] N. Chumerin and M. M. Van Hulle, "Comparison of two feature extraction methods based on maximization of mutual information," in *Proc. 16th IEEE Workshop Mach. Learn. Signal Process. (MLSP)*, Maynooth, Ireland, 2006, pp. 343–348.
- [56] M. Merluzzi, P. Di Lorenzo, and S. Barbarossa, "Wireless edge machine learning: Resource allocation and trade-offs," *IEEE Access*, vol. 9, pp. 45377–45398, 2021.
- [57] D. Hendrycks and T. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–16. [Online]. Available: <https://openreview.net/forum?id=HJz6tiCqYm>

Open Access funding provided by 'Università degli Studi di Perugia' within the CRUI CARE Agreement