



HAL
open science

Goal-oriented compression for L_p -norm-type goal functions: Application to power consumption scheduling

Yifei Sun, Hang Zou, Chao Zhang, Samson Lasaulce, Michel Kieffer

► To cite this version:

Yifei Sun, Hang Zou, Chao Zhang, Samson Lasaulce, Michel Kieffer. Goal-oriented compression for L_p -norm-type goal functions: Application to power consumption scheduling. *Journal of The Franklin Institute*, 2024, 361 (10), pp.106926. 10.1016/j.jfranklin.2024.106926 . hal-04602339

HAL Id: hal-04602339

<https://hal.science/hal-04602339v1>

Submitted on 5 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Highlights

Goal-oriented compression for L_p -norm-type goal functions: Application to power consumption scheduling

Yifei Sun, Hang Zou, Chao Zhang, Samson Lasaulce, Michel Kieffer

- General framework for designing compression methods for the L_p norm minimization problem.
- Novel linear and nonlinear transformation schemes by taking into account the performance degradation in terms of the L_p norm induced by model reduction.
- Tailor the quantization rule to be goal-oriented by considering the impact of the precoding and the final use of the compressed data.
- Evaluation of the proposed coding schemes with a real dataset and show the significant performance improvement compared to existing conventional transformation and quantization techniques.

Goal-oriented compression for L_p -norm-type goal functions: Application to power consumption scheduling

Yifei Sun^a, Hang Zou^b, Chao Zhang^{c,d}, Samson Lasaulce^{c,e}, Michel Kieffer^a

^a*Universite Paris-Saclay - CNRS - CentraleSupélec - L2S, Gif-sur-Yvette, F-91192, France*

^b*Technology Innovation Institute, Abu Dhabi, UAE*

^c*Khalifa University, 6G Research Center, Abu Dhabi, 127788, UAE*

^d*School of Computer Science and Engineering, Central South
University, Changsha, 410083, China*

^e*CRAN, CNRS-Universite de Lorraine, Nancy, F-54000, France*

Abstract

Conventional data compression schemes aim at implementing a trade-off between the rate required to represent the compressed data and the resulting distortion between the original and reconstructed data. However, in more and more applications, what is desired is not reconstruction accuracy but the quality of the realization of a certain task by the receiver. In this paper, the receiver task is modeled by an optimization problem whose parameters have to be compressed by the transmitter. Motivated by applications such as the smart grid, this paper focuses on a goal function which is of L_p -norm-type. The aim is to design the precoding, quantization, and decoding stages such that the maximum of the goal function obtained with the compressed version of the parameters is as close as possible to the maximum obtained without compression. The numerical analysis, based on real smart grid signals, clearly shows the benefits of the proposed approach compared to the conventional distortion-based compression paradigm.

Keywords: Data compression, Goal oriented communications, Quantization, Learning, Neural networks, Precoding.

1. Introduction

With the development of new paradigms such as the industrial internet, the internet of things (IoT), the smart grid, or networked controlled systems for instance, networks become more and more distributed. Information exchanges between the different devices are necessary to implement cooperation or coordination and achieve a given goal. A huge amount of data is often generated and transmitted, such as in

the smart grid [1, 2]. Because of practical limitations in terms of communication and computational resources, it is important if not necessary to compress the exchanged data.

There exists a quite solid literature on the problem of data compression in the smart grid. Compared to lossless compression (see e.g., [3]), lossy compression achieves a higher compression ratio at the expense of degrading the accuracy of data. Reference [4] transforms power quality event data to 2D, showing the correlation between far away sample, which help to exploit the redundancy of the data. [5] compresses the estimated fundamental component and transient component by different compression techniques. A two-step compression is proposed in [6]. The dimension of the data is firstly reduced by using principal component analysis. The preserved components are further compressed by using compressed sampling. Due to the limit in terms of rate, quantization schemes [7] are often used to minimize distortion [8].

Most of the aforementioned lossy compression techniques try to minimize the mean square error (MSE) while compressing the data. Alternative distortion measures have been considered for a long time, see [7] and the references therein. Indeed, it has been known that the MSE may be not suited for assessing the performance of various image processing operations (*e.g.*, image segmentation or pattern recognition). Over the past years, the signal processing and digital communications communities have realized that the distortion-based data compression paradigm should be revisited [9, 10, 11, 12, 13, 14, 15] by adapting the compression scheme to the final use, task, or goal pursued by the receiver. The attempt to incorporate aspects such as semantics and effectiveness into communication theory is in fact not new. For instance, reference [16] adopts a probabilistic logic approach to study the semantic aspect whereas [17] develops a complexity theory approach of the goal-oriented communication problem. Concerning the signal processing and communication point of view, which is the one of interest for the present paper, the literature of goal-oriented communications is still in full development. For instance, a goal-oriented signal processing technique, based on a graph-based semantic language and a goal filtering method, is proposed and applied to specific goals in [13]. The goal-oriented/semantic communication concept has also been developed to design the next cellular communication generation (see e.g., [14, 18, 19, 20]).

Concerning the goal-oriented approach for the data compression problem, the literature is equally relatively small but contains some relevant works that the present work can be related to. For data compression also, the signal processing community, has also been aware well before the recent publications on the topic that large compression gains might be reaped when considering the final use of the signal or

image (see, *e.g.*, [21, 22, 23, 24] for the case of speech compression). The goal-oriented quantization problem is posed and formalized for the first time in [9] and is tackled in a deeper manner in [25], [26], [10], [27], and in [28]. In [29], the problem of goal-oriented quantization is studied for the control problem of rendez-vous in a multi-agent setting. As for the goal-oriented precoding or data preprocessing problem, it has been introduced independently in [10] and [30].

Compared to the most related works, the present paper provides significant progresses into the design of goal-oriented data source encoders. The paper presents a novel task-oriented compression scheme when the task can be modeled by an optimization problem, which is a very relevant model for communications and energy systems. Our approach precisely assumes the task can be represented by a function (to be maximized) whose variables are the decisions to be taken and whose parameters have to be compressed. The focus of this paper is on the L_p -norm minimization problem, which is an important problem for example for smart grid applications, and in particular, for the fundamental problem of power consumption scheduling (*e.g.*, to minimize the total consumed energy price, Joule losses, or the peak power). More specifically, for the goal-oriented precoding or transform stage, the best linear solution is developed and compared to a non-linear solution which is based on a convolution auto-encoder. In contrast with [10], the proposed precoding scheme exploits the structure of the decision function (an approximation of it to be precise). Also, a goal-oriented quantizer is used not to quantize the input signal (as in [29, 28]) but the goal-oriented precoder output. Additionally, the used quantizer is not assumed to be a set of uniform scalar quantizers as in [10] and also it is tailored to the goal, which is an L_p -norm-type utility function. We also investigate the problem of knowing how the two compression stages interact each other, to understand to what extent it is possible to accumulate the gains of the goal-oriented paradigm when applied to the two stages. Last but not least, the schemes are applied to real smart grid measurements, which leads to a detailed numerical performance analysis and discussing the design of goal-oriented precoding and quantization. The paper does not only provide implementable coding schemes but also provides quantitative elements behind the intuition that accounting for the impact of compression noise on the task is beneficial. Our approach allows one to provide *e.g.*, an analytical characterization of an approximation of the optimal precoder, a goal-oriented quantization algorithm which works for an arbitrary utility function and not only for the L_p -norm. A purely data-driven approach (based on neural networks) would not provide these structural elements which can be exploited both for interpretations and making implementation easier.

The rest of the paper is organized as follows. We formulate the problem in Section

2. In Section 3 a linear approximation (LT) is considered for the optimal decision function. A linear and nonlinear transformation (NLT) are proposed in Section 4 and Section 5 respectively. We describe the proposed goal-oriented quantizer in Section 6. An iterative approach optimizing the linear transformation stage and the goal-oriented quantization stage is proposed in Section 7. Section 8 provides the numerical performance analysis. The paper is concluded in Section 9.

2. Problem formulation

Consider the following utility function

$$u(x; \ell) = -\|x + \ell\|_p \quad (1)$$

where $x = [x_1, x_2, \dots, x_N]^T \in \mathbb{R}_+^N$ represent the vector of decision variables, $\ell = [\ell_1, \ell_2, \dots, \ell_N]^T \in \mathbb{R}_+^N$ is a vector of parameters, and $\|\cdot\|_p$ is the L_p -norm, *i.e.*, for any $v \in \mathbb{R}^n$, $\|v\|_p = (|v_1|^p + \dots + |v_n|^p)^{1/p}$ with $p \geq 1$. Assuming that the sum of the decision variables is lower-bounded as

$$\sum_{j=1}^N x_j \geq E \quad (2)$$

where $E > 0$ is a constant and that (1) has to be maximized, one obtains the following optimization problem

$$\begin{aligned} & \underset{x}{\text{maximize}} \quad u(x; \ell) \\ & \text{s.t.} \quad \sum_{j=1}^N x_j - E \geq 0 \\ & \quad \quad x_j \geq 0, \quad j = 1, \dots, N. \end{aligned} \quad (3)$$

The solution of (3) is denoted as $x^*(\ell)$. The optimization problem (3) can model typical resource allocation problems in several applications. For instance, in the power consumption scheduling problem of energy systems, ℓ represents the non-controllable part of the energy consumption of a given household over N time slots (*e.g.*, consumed by the lighting and cooking appliances, TV, computers), x represents the controllable part (*e.g.*, the desired state-of-charge of an electric vehicle (EV)) to be allocated over the N time slots by some decision-making entity (scheduler), the constraint (2) corresponds to the minimum amount of energy that has to be provided. The utility function thus corresponds to minus the L_p -norm of the total consumption vector. When p becomes large, maximizing the considering utility amounts to

minimizing the peak power. When $p = 2$ and ℓ and x are interpreted as currents instead of powers, then it corresponds to the Joule losses minimization problem.

We assume that only an approximated version $\hat{\ell} \in \mathbb{R}^N$ of the non-controllable load vector ℓ is available to the scheduler which takes the decision x . This assumption can be not only motivated by the existence of limitations in terms of communication or computational resources but also for a need in terms of privacy. A lossy compression technique is thus implemented to remove some redundancy from the source signal ℓ . Considering $\hat{\ell}$ in place of ℓ in (3), the resulting solution becomes $x^*(\hat{\ell})$.

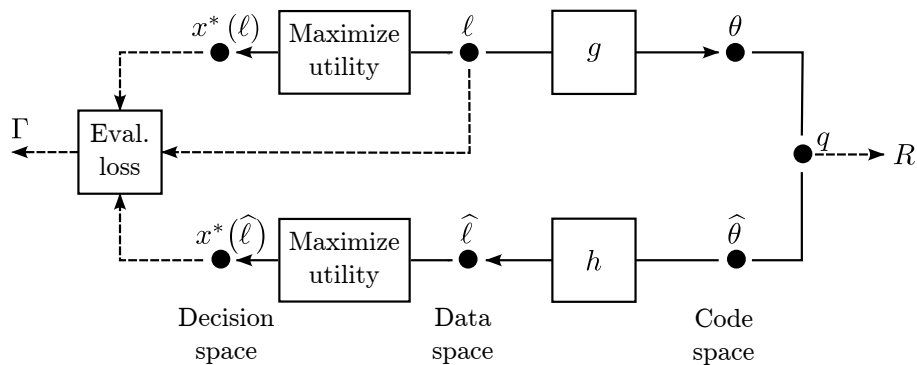


Figure 1: Coding scheme targeting optimality loss minimization

Instead of minimizing the reconstruction error as in conventional compression schemes, this paper aims to design a compression scheme for ℓ to mitigate the performance degradation resulting from the distortion (compression noise) introduced on $\hat{\ell}$, *i.e.*, to minimize the expected optimality loss

$$\Gamma = \mathbb{E}_{\ell} \left[\left| u(x^*(\ell); \ell) - u(x^*(\hat{\ell}); \ell) \right|^2 \right], \quad (4)$$

where the expectation is over ℓ . To define properly the arguments of Γ , it is first needed to define the different functions that need to be optimized. These functions appear in Figure 1. In this setting, the compression noise stems both from precoding/decoding and quantization. The precoding mapping

$$g: \begin{array}{l} \mathbb{R}_+^N \rightarrow \mathbb{R}^K \\ \ell \mapsto \theta \end{array} \quad (5)$$

transforms the parameter vector ℓ to a vector θ of dimension $K \leq N$, and the

dimensionality reduction induces some information loss. A quantizer

$$\begin{aligned} q: \mathbb{R}^K &\rightarrow \mathbb{R}^K \\ \theta &\mapsto \hat{\theta} \end{aligned} \tag{6}$$

follows the precoder, which constitutes a second source of the compression noise. The quantized vector $\hat{\theta}$ is transmitted and assumed received without error by the decision-making entity. A decoding function

$$\begin{aligned} h: \mathbb{R}^K &\rightarrow \mathbb{R}^N \\ \hat{\theta} &\mapsto \hat{\ell} \end{aligned} \tag{7}$$

is then used to obtain an estimate $\hat{\ell}$ of ℓ .

The goal is to design the precoding, quantization, and decoding functions that minimize the expected optimality loss (4)

$$(g^*, q^*, h^*) \in \arg \min_{(g, q, h)} \Gamma(g, q, h) = \mathbb{E}_\ell \left[\left| u(x^*(\ell); \ell) - u(x^*(\hat{\ell}); \ell) \right|^2 \right] \tag{8}$$

where $\hat{\ell} = h(q(g(\ell)))$.

Solving (8) directly is generally a hard task computationally speaking. This is the main motivation for searching for suboptimal solutions whose determination involve affordable complexity. In this paper, the functions g , q , and h are therefore optimized separately. More precisely, the precoding stage without quantization noise is designed. In contrast with conventional linear transformation such as the Karhunen-Loève Transform (KLT), the precoding stage aims at minimizing the optimality loss in terms of the utility function u . The KLT is known to be the best linear transformed in terms of MSE but this optimality no longer holds for the optimality loss. Second, the precoding/decoding scheme is fixed and a goal-oriented quantization scheme is proposed to mitigate the performance degradation brought by quantization noise. At last, the precoding/decoding stage and the quantization stage are optimized in an iterative manner.

3. Linear approximation of the optimal decision

The impact of quantization on the optimality loss, as defined by (4), can be seen to depend on the utility function u and the optimal decision function x^* . In [28], it has been proved formally (in the high-resolution regime) how the regularity and smoothness properties of these functions impact the optimality loss. To be able to

exploit the optimal decision function in the design of the precoding stage, we resort to a linear approximation of the former. In this section, we thus first characterize the solution of the optimization problem (3). Then, a linear approximation of this solution is provided to make the problem tractable.

3.1. Optimal decision function

For a given value of ℓ , the solution of (3) is provided by Proposition 1 and involves a water-filling approach, which is widely used [31, 32, 33].

Proposition 1. *Consider a value of the parameter vector ℓ , and assume, without loss of generality, that*

$$\ell_1 \leq \dots \leq \ell_j \leq \dots \leq \ell_N.$$

Then the components of $x^(\ell)$ are obtained by*

$$x_j^* = (\mu - \ell_j)^+, \quad (9)$$

where

$$\mu = \frac{1}{n^*} \left(E + \sum_{j=1}^{n^*} \ell_j \right)$$

indicates the water level, $(\cdot)^+ = \max(\cdot, 0)$, and n^* is the number of non-zero components of $x^*(\ell)$ evaluated as

$$n^* = \max_{n \leq N} \left\{ n : (n-1)\ell_n - \sum_{j=1}^{n-1} \ell_j \leq E \right\}. \quad (10)$$

Proof 1. *See Appendix A.1.*

3.2. Linear approximation

The solution provided by Proposition 1 has a non-linear dependency in ℓ , and this dependency is not explicit in general. To circumvent this difficulty we resort to a first order approximation of the optimal decision function. For this purpose, we evaluate the sensitivity of $x^*(\ell)$ with respect to ℓ , by considering the first-order Taylor expansion of x^* around ℓ

$$x^*(\ell + d\ell) = x^*(\ell) + \mathbf{H}(\ell)d\ell + o(d\ell) \quad (11)$$

where

$$\mathbf{H}(\ell) = \frac{\partial x^*}{\partial \ell^T}(\ell)$$

is the Jacobian matrix of $x^*(\ell)$ obtained using Proposition 2.

Proposition 2. Consider a value of the parameter vector ℓ such that n^* given by (10) remains constant over some neighborhood $\mathcal{N}(\ell)$ of ℓ . Without loss of generality, assume that

$$\ell_1 \leq \dots \leq \ell_{n^*} \leq \ell_{n^*+1} \leq \dots \leq \ell_N.$$

Then

$$x^*(\ell) = \mathbf{H}(\ell)\ell + b(\ell) \quad (12)$$

with

$$\mathbf{H}(\ell) = \begin{pmatrix} -1 + \frac{1}{n^*} & \frac{1}{n^*} & \dots & \frac{1}{n^*} & 0 & \dots & 0 \\ \frac{1}{n^*} & -1 + \frac{1}{n^*} & \ddots & \vdots & \vdots & & \vdots \\ \vdots & \ddots & \ddots & \vdots & \vdots & & \vdots \\ \frac{1}{n^*} & \dots & \frac{1}{n^*} & -1 + \frac{1}{n^*} & 0 & \dots & 0 \\ 0 & \dots & \dots & 0 & 0 & \dots & 0 \\ \vdots & & & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & \dots & 0 & 0 & \dots & 0 \end{pmatrix}$$

and

$$b(\ell) = \left(\frac{E}{n^*}, \dots, \frac{E}{n^*}, 0, \dots, 0 \right)^T.$$

Proof 2. See Appendix A.2.

4. Linear goal-oriented precoding

In this section, we investigate the problem of the determination of the best linear precoding or transform to be applied to ℓ , which is to approximate ℓ by a weighted sum of a given number of vectors of a certain basis; the number of vectors of this sum precisely corresponds to the number $K \in \{1, \dots, N\}$ defined by (5). It is known that, for a given number of basis vector, the KLT provides the best basis in the sense of the MSE $\mathbb{E}_\ell(\|\ell - \widehat{\ell}\|^2)$ (see, *e.g.*, [34]). Nevertheless, this optimality result is no longer true in the presence of an arbitrary goal and in particular for the optimality loss definition used in this paper. The motivation of this section is therefore to propose a linear precoding or transformation scheme which is matched to the L_p -norm function.

In this section, to simplify the analysis, the effect of quantization noise is assumed to be negligible. Under these assumptions, one has that $\theta = g(\ell) = \mathbf{B}\ell$, and

$\widehat{\ell} = h(\theta) = \mathbf{B}^T \theta$, where $\mathbf{B} \in \mathbb{R}^{K \times N}$ is the precoding matrix and \mathbf{B}^T the decoding matrix. Then the expected optimality loss (4) becomes

$$\Gamma(\mathbf{B}) = \mathbb{E}_\ell \left[\left| u(x^*(\ell); \ell) - u(x^*(\mathbf{B}^T \mathbf{B} \ell); \ell) \right|^2 \right]. \quad (13)$$

Two notational remarks are in order at this point. First, notice that in the above, a small abuse of notation is employed for clarity. Since, the optimality loss function is only considered with respect to the precoding stage, the arguments q and h are removed from Γ , and since only linear precoding is considered, the precoding function g is replaced with the matrix \mathbf{B} . Second, motivated by practical considerations, in which exact statistics are not available but one has only access to a set of measurements, the notation $\widehat{\Gamma}_T$ will be used to refer to the empirical version of the optimality loss function Γ . Assuming that a dataset $\mathcal{L} = \{\ell^{(1)}, \ell^{(2)}, \dots, \ell^{(T)}\}$ of T realizations or samples of ℓ is available, the empirical optimality loss expresses as

$$\widehat{\Gamma}_T(\mathbf{B}) = \frac{1}{T} \sum_{i=1}^T \left| u(x^*(\ell^{(i)}); \ell^{(i)}) - u(x^*(\mathbf{B}^T \mathbf{B} \ell^{(i)}); \ell^{(i)}) \right|^2. \quad (14)$$

Using the linear form (12) introduced in Proposition 2, one obtains

$$\begin{aligned} \widehat{\Gamma}_T(\mathbf{B}) &= \frac{1}{T} \sum_{i=1}^T \left| u(\mathbf{H}(\ell^{(i)}) \ell^{(i)} + b(\ell^{(i)}); \ell^{(i)}) \right. \\ &\quad \left. - u(\mathbf{H}(\mathbf{B}^T \mathbf{B} \ell^{(i)}) \mathbf{B}^T \mathbf{B} \ell^{(i)} + b(\mathbf{B}^T \mathbf{B} \ell^{(i)}); \ell^{(i)}) \right|^2. \end{aligned} \quad (15)$$

To minimize $\widehat{\Gamma}_T$ we resort to a gradient descent algorithm. To simplify the computation procedure of the derivatives of $\widehat{\Gamma}_T$, it is assumed that a small variation of \mathbf{B} does not change the entries of $\mathbf{H}(\mathbf{B}^T \mathbf{B} \ell^{(i)})$ and $b(\mathbf{B}^T \mathbf{B} \ell^{(i)})$ for all $i = 1, \dots, T$. This assumption allows one to state the following proposition.

Proposition 3. *Consider $p \in \mathbb{N}^+$. The matrix containing the derivatives of $\widehat{\Gamma}_T(\mathbf{B})$ with respect to the components of \mathbf{B} is*

$$\nabla_{\mathbf{B}} \widehat{\Gamma}_T(\mathbf{B}) = \frac{1}{T} \sum_{i=1}^T C_i (\mathbf{B} \ell^{(i)} \beta_i^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T \beta_i \ell^{(i)T}) \quad (16)$$

where

$$C_i = 2 \left(u(x^*(\ell^{(i)}); \ell^{(i)}) - u(x^*(\widehat{\ell}^{(i)}); \ell^{(i)}) \right) \|x^*(\widehat{\ell}^{(i)}) + \ell^{(i)}\|_p^{1-p} \quad (17)$$

$$\beta_i = \underbrace{\left(x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)} \right) \odot \cdots \odot \left(x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)} \right)}_{p-1} \quad (18)$$

where \odot indicates Hadamard product.

Moreover, when $p \rightarrow +\infty$, (16) boils down to

$$\nabla_{\mathbf{B}} \widehat{\Gamma}_T(\mathbf{B}) = \frac{1}{T} \sum_{i=1}^T D_i \left(\mathbf{B} \ell^{(i)} s_{k(i)}^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T s_{k(i)} \ell^{(i)T} \right) \quad (19)$$

where

$$D_i = 2 \left(u \left(x^* \left(\ell^{(i)} \right); \ell^{(i)} \right) - u \left(x^* \left(\widehat{\ell}^{(i)} \right); \ell^{(i)} \right) \right) \quad (20)$$

$$s_{k(i)} = \begin{pmatrix} 0_{(k(i)-1) \times 1} \\ 1 \\ 0_{(N-k(i)) \times 1} \end{pmatrix} \quad (21)$$

and

$$k(i) = \arg \max_k x_k^* \left(\widehat{\ell}^{(i)} \right) + \ell_k^{(i)}.$$

Proof 3. See Appendix A.3

Using (16) or (19), a local minimization of $\widehat{\Gamma}_T(\mathbf{B})$ can be performed by using a gradient descent algorithm (see Algorithm 1) to obtain \mathbf{B}^* . The search is initialized with the K first vectors of the KLT built from the empirical covariance matrix obtained by using the vectors of the dataset \mathcal{L} .

Algorithm 1 Gradient descent search \mathbf{B}^*

Input: Initialize \mathbf{B} with the K first vectors of Karhunen-Loève transform

Input: Initialize $i = 0$,

while $i < it_{\max}$ and optimality loss reduced more than 0.01% **do**

 Compute the gradient: $\mathbf{G} \leftarrow \nabla_{\mathbf{B}} \widehat{\Gamma}_T(\mathbf{B})$

 Perform line search to get λ such that $\widehat{\Gamma}_T(\mathbf{B} - \lambda \mathbf{G}) < \widehat{\Gamma}_T(\mathbf{B})$

 Update: $\mathbf{B} = \mathbf{B} - \lambda \mathbf{G}$

$i \leftarrow i + 1$

end while

5. Nonlinear goal-oriented precoding

The use of linear transforms for data compression is largely motivated by complexity issues. Nevertheless, linear transforms are generally not optimal. They are optimal in terms of MSE when the input signal corresponds to realizations of a (vector) Gaussian random variable [34]. A natural question is thus to assess the benefits of a nonlinear transform in the presence of goal functions such as the L_p -norm. This is why we consider here a larger class of precoders for which the parameters can be learned. The aim of the precoder is to obtain a latent representation $\theta \in \mathbb{R}^K$ of ℓ as follows $\theta = g(\ell, \Phi)$, where Φ is a vector of parameters. Then, the decoder takes a possibly quantized version of the latent representation as input to obtain an estimate of ℓ as $\hat{\ell} = h(\theta, \Psi)$, where Ψ is a vector of parameters.

Auto-encoders appear as a natural tool to implement a nonlinear transform when the purpose is to perform model reduction. Auto-encoders have been previously developed for image compression [35, 36, 37] and are considered here to obtain a suited low-dimension representation of ℓ .

A convolutional auto-encoder consists of an encoding stage and a decoding stage as illustrated in Figure 2. The encoding stage is simply a concatenation of convolution layers followed by a fully connected layer in the end while the decoding stage implements the inverse structure of the encoding one.




Figure 2: Structure of the considered convolutional auto-encoder to evaluate the nonlinear goal-oriented precoding and decoding functions

The convolutional auto-encoder aims at searching a pair of (Φ^*, Ψ^*) minimizing the expected optimality loss

$$(\Phi^*, \Psi^*) \in \arg \min_{(\Phi, \Psi)} \mathbb{E}_\ell [|u(x^*(\ell); \ell) - u(x^*(h(g(\ell, \Phi), \Psi)); \ell)|^2]. \quad (22)$$

As in the previous section, we also make a small notational abuse, since Ψ and Φ are taken as the arguments of $\widehat{\Gamma}_T$.

Assume a set \mathcal{L} of realizations of ℓ is available, (Φ^*, Ψ^*) can be trained through the loss function defined as

$$\widehat{\Gamma}_T(\Phi, \Psi) = \frac{1}{T} \sum_{i=1}^T |u(x^*(\ell^{(i)}); \ell^{(i)}) - u(x^*(h(g(\ell^{(i)}, \Phi), \Psi)); \ell^{(i)})|^2. \quad (23)$$

More precisely, as for the encoding stage in our simulation, one convolutional layer involving two kernels with size 5 is considered. Zero padding is implemented to impose the output of each kernel having the same dimension as the input. The outputs of two kernels are reshaped to a vector of dimension $2N$. A fully connected layer between the $2N$ neurons and K neurons is considered to get the latent representation θ . The decoding part is inverse processing to reconstruct the data $\widehat{\ell}$ with dimension N .

6. Goal-oriented quantization

In this section, we assume that the precoding and decoding functions $g(\cdot)$ and $h(\cdot)$ are given. We focus on the optimization of the quantizer design in order to minimize the optimality loss for a given choice of $g(\cdot)$ and $h(\cdot)$.

For the choice of quantization schemes, vector quantization is considered in this paper. Targeting signal reconstruction quality, using uniform scalar quantizers followed by entropy coding constitutes an overwhelmingly popular solution; it is adopted in neural compression schemes [35]. Nevertheless, when the principal goal pursued is not to reconstruct the original data, it has been shown that element-wise uniform scalar quantization (namely, using a bank of scalar quantizers) could induce a significant performance degradation for the task to be executed ([25, 28]). This can be explained by the fact that the different elements of the vector of variables to be compressed can have a markedly different influence on the final task; therefore, treating these elements equally is generally suboptimal. As a consequence, we use vector quantization in our scheme for task-oriented communication. Moreover, vector quantization allows one to ensure a fixed-rate quantized output, in contrast with auto-encoders accounting for rate constraints such as [35]. At last, our choice is very well suited for the power consumption scheduling problem exploiting smart meter measurements, which have a relatively small dimension.

A vector quantizer partitions the space \mathbb{R}^K of the encoded parameter $\theta = g(\ell)$

into several disjoint quantization regions $\mathcal{C}_1, \dots, \mathcal{C}_M$, *i.e.*,

$$\bigcup_{i=1}^M \mathcal{C}_i = \mathbb{R}^K$$

and $\mathcal{C}_i \cap \mathcal{C}_j = \emptyset$ for all $i \neq j$. The quantization function is such that

$$q(\theta) = r_i \iff \theta \in \mathcal{C}_i, \quad (24)$$

where $\mathcal{R} = \{r_1, \dots, r_M\}$ is the set of representatives associated with the quantization regions $\mathcal{C}_1, \dots, \mathcal{C}_M$ respectively. Contrary to conventional quantizers, which usually minimize the mean-square reconstruction error, in what follows, we search for a pair $(\mathcal{R}^*, \mathcal{C}^*)$ that minimizes the expected optimality loss

$$\begin{aligned} \mathbb{E}_\ell \left[\left| u(x^*(\ell); \ell) - u(x^*(\hat{\ell}); \ell) \right|^2 \right] = \\ \sum_{m=1}^M \int_{\ell \in \mathcal{L}_m} |u(x^*(\ell); \ell) - u(x^*(h(r_m)); \ell)|^2 \phi(\ell) d\ell, \end{aligned} \quad (25)$$

where $\phi(\ell)$ is the probability density function of ℓ and

$$\mathcal{L}_m \triangleq \{\ell \in \mathbb{R}^N | g(\ell) \in \mathcal{C}_m\}, 1 \leq m \leq M.$$

Finding jointly $(\mathcal{R}^*, \mathcal{C}^*)$ is not trivial. Moreover, the evaluation of (25) considering $\phi(\ell)$ is complex, even if $\phi(\ell)$ is perfectly known. Consequently, a practical algorithm is proposed which is similar to the decisional quantizer proposed in [26]. The main steps of this algorithm are detailed in what follows. As in Section 4, a set $\mathcal{L} = \{\ell^{(1)}, \ell^{(2)}, \dots, \ell^{(T)}\}$ of realizations of the parameter ℓ is used to approximate (25).

Assume that at iteration j , a set of representatives $\mathcal{R}^{(j)} = \{r_1^{(j)}, \dots, r_M^{(j)}\}$ is available.

1. From $\mathcal{R}^{(j)}$, partition the set \mathcal{L} as

$$\mathcal{L}_m^{(j)} = \left\{ \ell \in \mathcal{L} \mid \mathcal{E}(r_m^{(j)}; \ell) = \min_i \mathcal{E}(r_i^{(j)}; \ell) \right\}, 1 \leq m \leq M, \quad (26)$$

where

$$\mathcal{E}(r_m; \ell) = |u(x^*(\ell); \ell) - u(x^*(h(r_m)); \ell)|^2. \quad (27)$$

2. Update the set of representatives $\mathcal{R}^{(j+1)} = \{r_1^{(j+1)}, \dots, r_M^{(j+1)}\}$ as

$$r_m^{(j+1)} \in \arg \min_r \sum_{\ell \in \mathcal{L}_m^{(j)}} \mathcal{E}(r; \ell), 1 \leq m \leq M. \quad (28)$$

For any $\ell \in \mathbb{R}_+^K$, the goal-oriented quantizer output is then obtained as

$$q^{(j+1)}(g(\ell)) = r_m^{(j+1)}$$

with

$$m = \arg \min_{i=1, \dots, M} \mathcal{E}(r_i^{(j+1)}; \ell).$$

3. Evaluate the resulting estimate of the expected optimality loss as

$$\Gamma^{(j+1)} = \sum_{m=1}^M \sum_{\ell \in \mathcal{L}_m^{(k)}} |u(x^*(\ell); \ell) - u(x^*(h(q^{(j+1)}(g(\ell))))); \ell)|^2. \quad (29)$$

Like Linde-Buzo-Gray (LBG) algorithm [38], Algorithm 2 performs this iterative evaluation to determine a locally optimal goal-oriented quantizer. The set of representatives $\mathcal{R}^{(0)}$ is initialized considering, *e.g.*, randomly chosen elements of the set $g(\mathcal{L}) = \{g(\ell^{(1)}), g(\ell^{(2)}), \dots, g(\ell^{(T)})\}$. As LBG algorithm, Algorithm 2 may

Algorithm 2 Goal-oriented quantizer design algorithm

Input: Utility function $u(x; \ell)$

Output: \mathcal{R}^*

Initialization: $\mathcal{R}^{(0)} = \{r_1^{(0)}, \dots, r_M^{(0)}\}$, $j = 1$, $\Gamma^{(0)} = \infty$

while $j \leq j_{\max}$ and optimality loss used more than 0.01% **do**

 Evaluate $\mathcal{L}_m^{(j)}$, $1 \leq m \leq M$, from $\mathcal{R}^{(j-1)}$ using (26)

 Evaluate $r_m^{(j+1)}$, $m = 1, \dots, M$, using (28)

 Evaluate $\Gamma^{(j+1)}$ using (29)

$j = j + 1$

end while

only converge to a local minimum of the expected optimality loss. The convergence depends on the initialization of $\mathcal{R}^{(0)}$.

7. Iterative optimization of linear transform and goal-oriented quantization

Using Algorithm 1 and Algorithm 2, one obtains a pair of separately designed precoder/decoder pair and quantizer. In Section 4, Algorithm 1 neglects the impact of the quantization noise in the design of the linear precoder/decoder. This section presents an iterative design, accounting for the quantization noise in Algorithm 1, to further use the optimality loss.

The effect of quantization in Section 6 is to add some noise η to the output of the precoder

$$\begin{aligned}\widehat{\theta} &= q(\theta) \\ &= \theta + \eta.\end{aligned}$$

We assume that η is Gaussian with mean $\mu_\eta \in \mathbb{R}^K$ and covariance matrix Σ_η . Considering the linear precoder/decoder pair, the reconstructed signal becomes

$$\begin{aligned}\widehat{\ell} &= \mathbf{B}^T \widehat{\theta} \\ &= \mathbf{B}^T (\theta + \eta) \\ &= \widetilde{\ell} + \mathbf{B}^T \eta\end{aligned}\tag{30}$$

where $\widetilde{\ell} = \mathbf{B}^T \mathbf{B} \ell = \mathbf{B}^T \theta$ is reconstructed from θ in absence of quantization.

Consider the j -th iteration of an iterative optimization algorithm of the precoder/decoder pair and of the quantizer. Assume that a precoding matrix $\mathbf{B}^{(j)}$, and a set of representatives $\mathcal{R}^{(j)}$ have been obtained from Algorithms 1 and 2. The characteristics of the quantization noise can be evaluated using the set \mathcal{L} of realizations of ℓ as

$$\begin{aligned}\widehat{\mu}^{(j)} &= \frac{1}{T} \sum_{i=1}^T \eta^{(i)} \\ \widehat{\Sigma}^{(j)} &= \frac{1}{T} \sum_{i=1}^T (\eta^{(i)} - \widehat{\mu}_\eta) (\eta^{(i)} - \widehat{\mu}_\eta)^T\end{aligned}$$

where

$$\eta^{(i)} = q\left(\mathbf{B}^{(j)} \ell^{(i)}\right) - \mathbf{B}^{(j)} \ell^{(i)} \quad .\tag{31}$$

At iteration $j + 1$, the impact of the quantization noise can be introduced by introducing for each $\ell^{(i)} \in \mathcal{L}$, $\bar{\kappa}$ noisy reconstructions of $\ell^{(i)}$

$$\widehat{\ell}^{(i,\kappa)} = \left(\mathbf{B}^{(j)}\right)^{\text{T}} \mathbf{B}^{(j)} \ell^{(i)} + \left(\mathbf{B}^{(j)}\right)^{\text{T}} \eta^{(i,\kappa)} \quad (32)$$

where $\eta^{(i,\kappa)}$, $1 \leq \kappa \leq \bar{\kappa}$ are realizations of independent and identically distributed Gaussian random vectors with mean $\widehat{\mu}^{(j)}$ and covariance matrix $\widehat{\Sigma}^{(j)}$. The noisy reconstructions $\widehat{\ell}^{(i,\kappa)}$ are then used in the evaluation of the gradient (16) or (19) to obtain a precoder/decoder optimization accounting for the quantization noise. As a result, a matrix $\mathbf{B}^{(j+1)}$ is obtained. Algorithm 2 is then applied to get $\mathcal{R}^{(j+1)}$. Then $\widehat{\mu}^{(j+1)}$ and $\widehat{\Sigma}^{(j+1)}$ can be evaluated.

The iterative optimization process is summarized in Algorithm 3. The computational complexity of linear precoding part is $O(TN^2K)$, the complexity of the goal-oriented quantization part is $O(N_{rep}(TNK + C_r))$, where N_{rep} is the overall number of representatives and C_r represents the complexity of computing the representatives of the L_p norm optimization problem.

Algorithm 3 Iterative optimization of the linear transformation and quantization

Input: Initial matrix $\mathbf{B}^{(0)}$ (KL basis), utility function $u(x; \ell)$, initial $\mathcal{R}^{(0)}$, $\mu^{(0)}$, and $\Sigma^{(0)}$;

Output: \mathbf{B}^* and \mathcal{R}^* ;

Initialization: $j = 1$;

while $j \leq j_{\max}$ and optimality loss reduced more than 0.01% **do**

 Evaluate $\mathbf{B}^{(j)}$ using Algorithm 1 and (32)

 Evaluate $\mathcal{R}^{(j)}$ using Algorithm 2

 Evaluate $\mu^{(j)}$ and $\Sigma^{(j)}$

 Evaluate $\Gamma^{(j)}$ using (29)

$j = j + 1$

end while

8. Numerical performance analysis

In this section, we conduct a comprehensive numerical analysis to provide insights into the benefits of goal-oriented compression. Our focus is on the power consumption scheduling problem in which the decision maker aims to find a controllable consumption vector x minimizing the L_p norm given a (perfect) day-ahead forecast of the non-controllable vector ℓ (the case of considering the forecast noise

can be treated as an extension of this work). Due to the communication resource limitation, the decision maker has the sole knowledge of an approximated version of ℓ , namely, $\hat{\ell}$. Our goal is to mitigate the performance degradation induced by the deviation between $\hat{\ell}$ and ℓ . To assess the influence of the compression loss on the goal, we define the relative squared optimality loss (RSOL) of the compression scheme with respect to the ideal case as follows:

$$\rho_C(\%) = 100 \frac{\sum_{i=1}^T \left(U_{\text{perfect}}^{(i)} - U_C^{(i)} \right)^2}{\sum_{i=1}^T \left(U_{\text{perfect}}^{(i)} \right)^2}, \quad (33)$$

where $U^{(i)}$ is obtained from the realization $\ell^{(i)}$, C represents the compression scheme, and the performance of the ideal case is obtained by assuming that the controllable consumption ℓ is perfectly known by the receiver (that is, the decision-maker).

To evaluate our approach in practical scenarios, consumption profiles from the Ausgrid [39] database are used for all the considered schemes. The data from Ausgrid consist of daily energy consumptions sampled every half hour for one year (from 01 – 07 – 2012 to 30 – 06 – 2013) for 300 users, thus the dimension of the data is $48 \times (365 \times 300)$, that is $N = 48$. Data have been randomly shuffled. Then 80 % of the data have been used to train the methods and the remaining 20 % for evaluation. The energy need in terms of controllable consumption is set to $E = 50$ kWh. As the system consists of several stages including precoding, quantization and decoding, we present the simulation results from different aspects for a comprehensive illustration.

8.1. Precoding

We first evaluate the benefit of using goal-oriented precoding schemes in the scenario without quantization noise. A KLT-based precoding scheme and an auto-encoder (AE) serve as references. The AE scheme is an adaptation to the compression for 1D vectors of the Tensorflow learned data compression network [40], itself based on the [35]. The architecture consists of three parts: an encoder that transforms the input vector into a lower-dimensional latent variable; a decoder that reconstructs the original source vector from this latent representation; a prior and entropy model between the encoder and decoder to model the marginal distribution of the latent variable and efficiently encode it to minimize the average code length. In this first part, the AE is trained to minimize the reconstruction mean-square error (MSE) as commonly used. When $K = 1$, Figure 3 clearly demonstrates the improvement provided by the proposed goal-oriented transformation schemes over the conventional KLT and the AE. The nonlinear transformation scheme obtained from our neural




Figure 3: Relative squared optimality loss (%) v.s. the exponent of L_p -norm (p) when $K = 1$. Both goal-oriented linear transform (LT) and goal-oriented non-linear transform (NLT) have a lower RSOL than the Karhunen-Loeve transform (KLT) and the auto-encoder (AE transform). Among all schemes, the proposed NLT solution provides the best result. The influence of p on the RSOL value shows also the influence of the goal function on the benefits of the goal-oriented approach.

network architecture outperforms the solution given by Algo. 1, since the activation function brings nonlinearity to the scheme. Moreover, one can observe that the RSOL increases with p . This can be explained by the fact that the denominator of RSOL defined by (33) usually decreases when p increases. For the remaining simulations, we fix $p = \infty$, corresponding to the peak power minimization problem.

To evaluate the influence of dimensionality reduction on the goal-oriented transformation schemes, we plot the RSOL w.r.t. the value of K . For the KLT, the orthogonal basis is obtained as eigenvectors of the covariance matrix, and the basis vectors are sorted from the most to the least important in terms of minimizing reconstruction error. One can observe that the proposed NLT scheme outperforms other techniques. Moreover, while the AE structure leverages the nonlinear capabilities of neural networks, its lack of task-specific focus results in greater performance degradation when using a smaller K . This loss in performance can be substantially mitigated by employing a latent space with larger dimensions.




Figure 4: Relative squared optimality loss (%) v.s. the dimension of the precoding output signal space / latent space (K).

Additionally, to interpret the obtained gain in a more intuitive manner, we choose one representative realization of ℓ from Ausgrid, and display the series of ℓ , and $x^*(\hat{\ell}) + \ell$ with different transformation methods in Figure 5. By approximating ℓ with a single basis vector, one can observe that the shape of $\hat{\ell}$ in goal-oriented transformations is more similar to ℓ compared with the KLT. As a result, there are fewer fluctuations in $x^*(\hat{\ell}) + \ell$, making it easier to minimize the L_p norm using the proposed LT and NLT schemes.

8.2. Quantization

After assessing the impact of transformation, we further study the goal-oriented quantization with given transformation schemes. We use the above-mentioned two transformation schemes, followed by a quantizer, and compare the performance with different quantization techniques. When the precoding process reduces the dimension of ℓ to $K = 1$, we compare three quantization techniques, namely, uniform quantization, LBG and GOQ. Figure 6 represents the RSOL against the number of quantization bits. When implementing LT and uniform quantization, we use the hardware-limited task-based quantization (HLTB) algorithm proposed by [10]. Figure 6 shows the GOQ exploits the quantization resources more efficiently such that the optimality loss can be reduced significantly, especially with non-linear transformation.

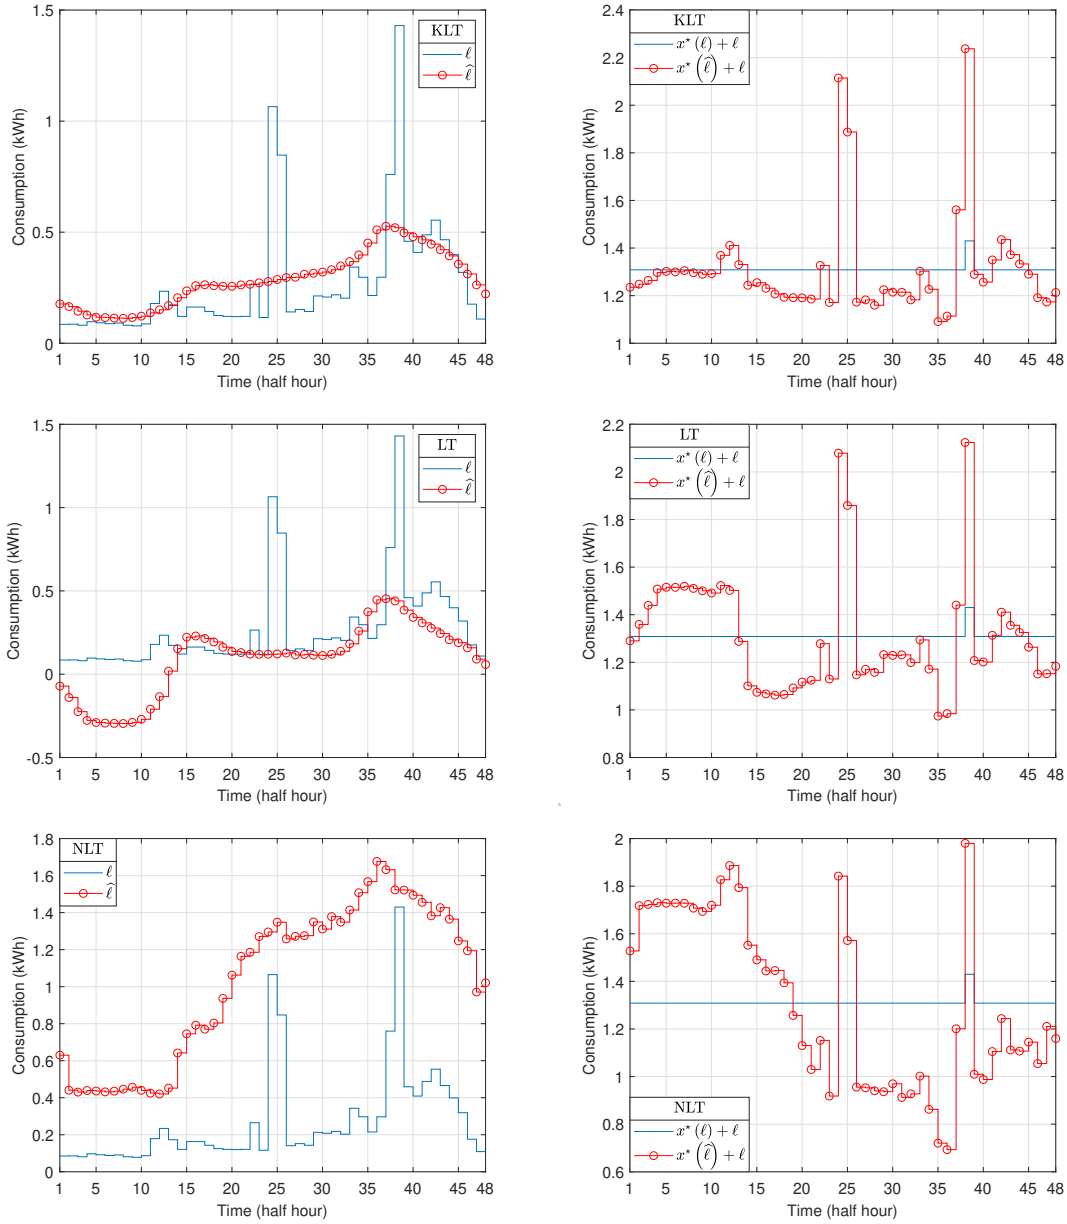


Figure 5: Daily energy consumption ℓ of one user and its approximations with the corresponding utility entries obtained by three different methods (when $K = 1$). NLT has the best performance of minimization of peak value.

Another important issue to discuss here is the choice of the reduced dimension K for a given total bit budget. Intuitively, when the number of bits tends to infinity, it is not necessary to reduce the dimension in the precoding stage since each element of the transformed vector can be quantized with negligible errors. Conversely, when the total bit budget is very small, a significant reduction in dimensionality is expected, as these bits need to transmit the most important features. Based on these observations, we explore the tradeoff between information loss induced by dimensionality reduction and information loss induced by quantization with finite bit budget. To determine the optimal dimension K that minimizes the optimality loss, we implement the proposed linear transformation scheme and the GOQ method. Figure 7 illustrates that the RSOL generally first decreases and then increases as K becomes larger, indicating the existence of an optimal dimension for the quantization input. This can be explained by the fact that we are operating at a fixed budget on the total number of quantization bits. Therefore, when K increases, the number of bits per dimension decreases, hence the observed phenomenon. This analysis highlights the importance of selecting an appropriate K for dimensionality reduction in the context of quantization, balancing the reduction in information loss due to dimensionality reduction with the information loss introduced by quantization. In addition, we address the rate-relative square optimality loss tradeoff in Fig. 8 for different values of K . For the AE, the compromise is tuned by properly weighting the MSE and the rate in the loss function. The network structure has to be retrained for each K and each value of the weights. In contrast with the AE structure, our method can achieve a given level of RSOL with much less transmission bits.

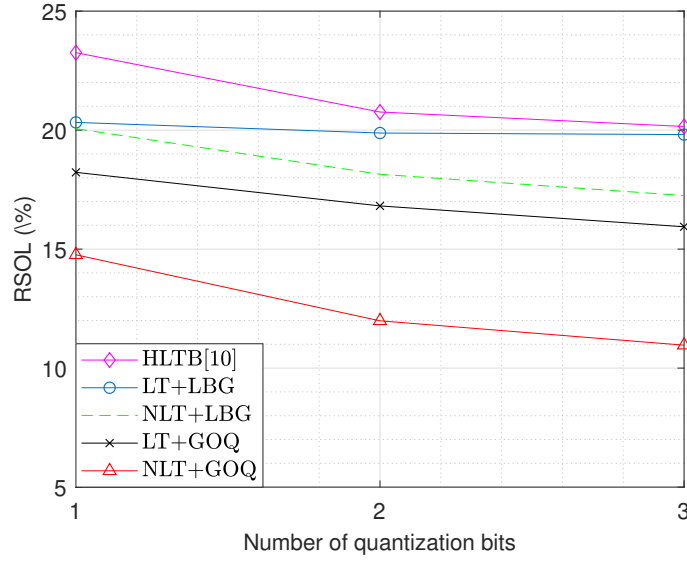


Figure 6: Evolution of the relative optimality loss (%) with quantization bit constraint ($N = 48$). RSOL is reduced with larger budget of quantization bits. Among these methods, NLT followed by a GOQ has the best performance.




Figure 7: Relative optimality loss v.s. the dimension of encoded space (K) with a fixed quantization bit constraint for linear transformation and goal-oriented quantization




Figure 8: Relative squared optimality loss v.s. the total rate as a function for different latent space dimensions (K) for the linear transform followed by the goal-oriented vector quantizer and for the auto-encoder structure adapted from [40].

8.3. Iterative algorithm

At last, we assess the performance of the proposed iterative algorithm by optimizing the transformation scheme and the quantization rule in an alternative way. By taking into account the interplay between quantization noise and the transformation scheme, Figure 9 shows that the iterative algorithm outperforms the aforementioned method using LT and GOQ once. Although the iterative algorithm provides only a marginal improvement beyond the initial optimization, it demonstrates the potential for further enhancement through careful refinement of the transformation scheme and the quantization rule. Regarding the computation cost, while iterative optimization has its merits in situations where resources are abundant and performance maximization is required, a single iteration suffices in resource-limited scenarios, providing a pragmatic balance between performance and computational efficiency.

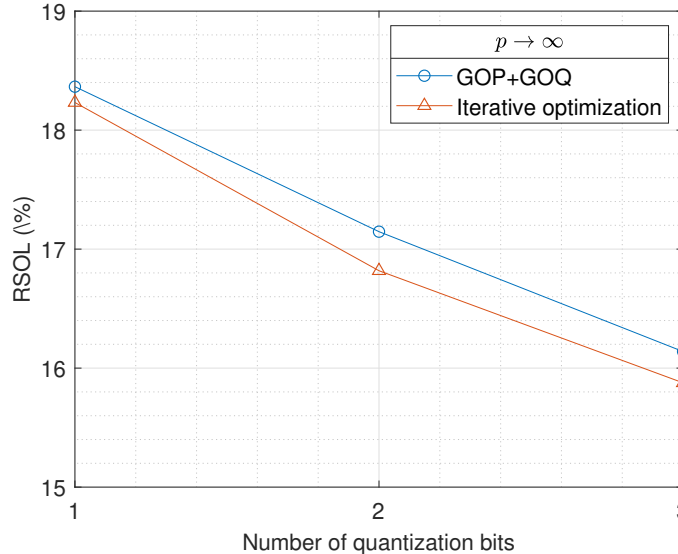


Figure 9: Comparing the result of the iterative optimization algorithm to the result without optimizing ($K = 1$, $p = \infty$)

9. Conclusion

In this paper, the task of the receiver is modeled by an optimization problem. The goal of the receiver is to minimize an L_p -norm performance metric. For taking its decision, the receiver is assumed to have only access to a compressed version of the parameters of the function (minus an L_p -norm) to be maximized. The problem of designing a goal-oriented precoder which is followed by a goal-oriented quantizer is addressed. By adopting a non-joint design approach for these two stages and making appropriate approximations, the problem becomes tractable. We provide both an interpretable linear transform which exploits the knowledge of the utility function u and optimal decision, and another non-linear transform based on neural networks. Compared to the KLT, the proposed linear transformation yields significant performance gains in terms of relative optimality loss. It is also seen to what extent the CNN-based nonlinear transformation performs better than the proposed linear transformation. By moving from the KLT to the CNN-based goal-oriented precoding, the relative optimality loss has been to drop from 20% to values as small as 1 – 2%, which shows all merits of matching the coding scheme to the goal. It is also seen under which conditions, the benefit from using a goal-oriented quantization can accumulate with those from the goal-oriented precoding stage. All these positive results show the interest of adapting, possibly on the fly, the coding scheme to the

task of the receiver. The proposed approach might be extended to other settings of practical interest. For instance, when the goal function is not known but only realizations of it are available. This would naturally lead to a reinforcement learning perspective for the design of goal-oriented encoders. Also, in this paper, the communication channel between the encoder and decoder is assumed to be perfect, which leaves space for improvements in the presence of communication noise as considered in [41, 42].

Appendix A. Proofs

Appendix A.1. Proof of Proposition 1

We want to prove that the solution x^* which maximizes (1) under the constraints (2) and (3) is a water-filling solution. For this, we first notice that (1) is a concave function and then apply KKT conditions, which are necessary and sufficient optimality conditions since (2) and (3) are affine constraints.

Consider a given vector ℓ and assume without loss of generality that

$$\ell_1 \leq \ell_2 \leq \dots \leq \ell_j \leq \ell_{j+1} \leq \dots \leq \ell_N. \quad (\text{A.1})$$

The considered optimization problem (3) is a convex problem since the Lp-norm is convex and the constraint functions are affine. Introducing the Lagrangian

$$\mathcal{L} = - \sum_{k=1}^N (x_k + \ell_k)^p - \delta \left(E - \sum_{k=1}^N x_k \right) - \sum_{k=1}^N \lambda_k x_k \quad (\text{A.2})$$

and applying KKT conditions, one obtains, as $p \geq 1$

$$\frac{\partial \mathcal{L}}{\partial x_k} = -p(x_k + \ell_k)^{p-1} + \delta - \lambda_k, \quad (\text{A.3})$$

and

$$\begin{cases} -p(x_k^* + \ell_k)^{p-1} + \delta^* - \lambda_k^* & = 0 \\ E - \sum_{i=1}^N x_i^* & = 0 \\ -x_k^* & \leq 0 \\ \lambda_k^* x_k^* & = 0. \end{cases} \quad (\text{A.4})$$

Getting rid of the slack variables $\{\lambda_k^*\}$, one obtains

$$p(x_k^* + \ell_k)^{p-1} = \delta^* \quad (\text{A.5})$$

$$x_k^* = \left(\left(\frac{\delta^*}{p} \right)^{\frac{1}{p-1}} - \ell_k \right)^+ \quad (\text{A.6})$$

where $(a)^+ = \max(a, 0)$. This is therefore a water-filling solution with water level $\mu = \left(\frac{\delta^*}{p} \right)^{\frac{1}{p-1}}$, thus

$$x_j^* = (\mu - \ell_j)^+. \quad (\text{A.7})$$

Let n^* be the number of non-zero entries of x^* . According to (A.1), one has $x_j^* > 0$ for $j = 1, \dots, n^*$ and $x_j^* = 0$ for $j > n^*$. Consequently, as

$$\sum_{j=1}^{n^*} x_j^* = E, \quad (\text{A.8})$$

one obtains

$$\sum_{j=1}^{n^*} (\mu - \ell_j) = E, \quad (\text{A.9})$$

leading to

$$\mu = \frac{1}{n^*} \left(E + \sum_{j=1}^{n^*} \ell_j \right). \quad (\text{A.10})$$

The value of n^* is then obtained as the largest value of n such that $\mu - \ell_n > 0$

$$\begin{aligned} n^* &= \arg \max_n \\ &\text{s.t. } \frac{1}{n} \left(E + \sum_{j=1}^n \ell_j \right) - \ell_n > 0. \end{aligned}$$

Appendix A.2. Proof of Proposition 2

As seen in Appendix A.1, $x^*(\ell) = [x_1^*, x_2^*, \dots, x_N^*]^T$ is the function of ℓ . As $n^* \in \{1, \dots, N\}$, $2^N - 1$ different subsets of \mathbb{R}_+^N may be defined as

$$\mathcal{M}_{\mathcal{I}} = \{ \ell \in \mathbb{R}_+^N \mid x_k^*(\ell) > 0, k \in \mathcal{I}, x_j^*(\ell) = 0, j \notin \mathcal{I} \}$$

where

$$\mathcal{I} \subseteq \{1, 2, \dots, N\}, \mathcal{I} \neq \emptyset$$

is the set of indexes of the non-zero entries of x^* .

In each region $\mathcal{M}_{\mathcal{I}}$, a linear expression of the solution of the optimization problem (3) can be obtained. Consider some \mathcal{I} and the associated region $\mathcal{M}_{\mathcal{I}}$. We assume

without loss of generality that the last $N - n^*$ elements of x^* are null for all $\ell \in \mathcal{M}_T$. From (A.7) and (A.10), one has

$$x_j^*(\ell) = \begin{cases} \frac{\sum_{i=1}^{n^*} \ell_i + E}{n^*} - \ell_j & j \leq n^* \\ 0 & j > n^* \end{cases} \quad (\text{A.11})$$

the Jacobian matrix of $x^*(\ell)$ is

$$\mathbf{H}(\ell) = \begin{pmatrix} -1 + \frac{1}{n^*} & \frac{1}{n^*} & \cdots & \frac{1}{n^*} & 0 & \cdots & 0 \\ \frac{1}{n^*} & -1 + \frac{1}{n^*} & \cdots & \vdots & \vdots & & \vdots \\ \vdots & \vdots & & \vdots & \vdots & & \vdots \\ \frac{1}{n^*} & \cdots & \cdots & -1 + \frac{1}{n^*} & 0 & \cdots & 0 \\ 0 & \cdots & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & & & \vdots & \vdots & & \vdots \\ 0 & \cdots & \cdots & 0 & 0 & \cdots & 0 \end{pmatrix}. \quad (\text{A.12})$$

Using (A.11) and the expression of $\mathbf{H}(\ell)$, one obtains

$$x^*(\ell) = \mathbf{H}(\ell) \ell + b(\ell).$$

Appendix A.3. Proof of Proposition 3

From (14), when $p \in \mathbb{N}^+$

$$\frac{\partial \widehat{\Gamma}_T}{\partial \mathbf{B}}(\mathbf{B}) = -\frac{1}{T} \sum_{i=1}^T 2 \left(u(x^*(\ell^{(i)}); \ell^{(i)}) - u(x^*(\widehat{\ell}^{(i)}); \ell^{(i)}) \right) \frac{\partial \mathcal{G}_i}{\partial \mathbf{B}}$$

where

$$\begin{aligned} \mathcal{G}_i &= u(x^*(\widehat{\ell}^{(i)}); \ell^{(i)}) \\ &= \left(\mathbf{1}^T \cdot \underbrace{\left(x^*(\widehat{\ell}^{(i)}) + \ell^{(i)} \right) \odot \cdots \odot \left(x^*(\widehat{\ell}^{(i)}) + \ell^{(i)} \right)}_p \right)^{\frac{1}{p}} \end{aligned}$$

with $\mathbf{1} = [1, \dots, 1]^T$ the vector of N ones. Then

$$\begin{aligned}
d\mathcal{G}_i &= \frac{1}{p} \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} * \\
&\quad d \left(\underbrace{\mathbf{1}^T \left(x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)} \right) \odot \dots \odot \left(x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)} \right)}_p \right) \\
&= \frac{1}{p} \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} * p * \beta_i^T d \left(x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)} \right) \\
&= \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} \beta_i^T d \left((\mathbf{H}_i \mathbf{B}^T \mathbf{B} + \mathbf{I}) \ell^{(i)} + b_i \right) \\
&= \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} \beta_i^T \left(\mathbf{H}_i d\mathbf{B}^T \mathbf{B} \ell^{(i)} + \mathbf{H}_i \mathbf{B}^T d\mathbf{B} \ell^{(i)} \right)
\end{aligned}$$

where $\mathbf{H}_i = \mathbf{H} \left(\widehat{\ell}^{(i)} \right)$, $b_i = b \left(\widehat{\ell}^{(i)} \right)$, and

$$\beta_i = \underbrace{\left(x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)} \right) \odot \dots \odot \left(x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)} \right)}_{p-1}.$$

As $d\mathcal{G}_i$ is a scalar, one has

$$\begin{aligned}
d\mathcal{G}_i &= \text{Tr} (d\mathcal{G}_i) \\
&= \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} \left(\text{Tr} \left(\beta_i^T \mathbf{H}_i d\mathbf{B}^T \mathbf{B} \ell^{(i)} \right) + \text{Tr} \left(\beta_i^T \mathbf{H}_i \mathbf{B}^T d\mathbf{B} \ell^{(i)} \right) \right) \\
&= \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} \left(\text{Tr} \left(\mathbf{B} \ell^{(i)} \beta_i^T \mathbf{H}_i d\mathbf{B}^T \right) + \text{Tr} \left(\ell^{(i)} \beta_i^T \mathbf{H}_i \mathbf{B}^T d\mathbf{B} \right) \right) \\
&= \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} \text{Tr} \left(\left(\mathbf{B} \ell^{(i)} \beta_i^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T \beta_i \ell^{(i)T} \right)^T d\mathbf{B} \right).
\end{aligned}$$

Then, one can deduce

$$\frac{\partial \mathcal{G}_i}{\partial \mathbf{B}} = \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p} \left(\mathbf{B} \ell^{(i)} \beta_i^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T \beta_i \ell^{(i)T} \right)$$

and

$$\frac{\partial \widehat{\Gamma}_T}{\partial \mathbf{B}} (\mathbf{B}) = -\frac{1}{T} \sum_{i=1}^T C_i \left(\mathbf{B} \ell^{(i)} \beta_i^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T \beta_i \ell^{(i)T} \right) \quad (\text{A.13})$$

where

$$C_i = 2 \left(u \left(x^* \left(\ell^{(i)} \right); \ell^{(i)} \right) - u \left(x^* \left(\widehat{\ell}^{(i)} \right); \ell^{(i)} \right) \right) \|x^* \left(\widehat{\ell}^{(i)} \right) + \ell^{(i)}\|_p^{1-p}. \quad (\text{A.14})$$

When $p \rightarrow +\infty$, from (14), we have

$$\frac{\partial \widehat{\Gamma}_T}{\partial \mathbf{B}}(\mathbf{B}) = -\frac{1}{T} \sum_{i=1}^T 2 \left(u(x^*(\ell^{(i)}); \ell^{(i)}) - u(x^*(\widehat{\ell}^{(i)}); \ell^{(i)}) \right) \frac{\partial \mathcal{F}_i}{\partial \mathbf{B}}$$

with

$$\mathcal{F}_i = \max_j \left(\widehat{x}_j^{(i)} + \ell_j^{(i)} \right)$$

and $x^*(\widehat{\ell}^{(i)}) = \left[\widehat{x}_1^{(i)}, \widehat{x}_2^{(i)}, \dots, \widehat{x}_N^{(i)} \right]^T$.

Furthermore, we assume for i -th vector, the $k(i)$ -th elements is maximum. Then

$$\begin{aligned} \mathcal{F}_i &= s_{k(i)}^T \left(x(\widehat{\ell}^{(i)}) + \ell^{(i)} \right) \\ &= s_{k(i)}^T \left((\mathbf{H}_i \mathbf{B}^T \mathbf{B} + \mathbf{I}) \ell^{(i)} + b_i \right) \end{aligned}$$

with

$$s_{k(i)} = \begin{pmatrix} 0_{(k(i)-1) \times 1} \\ 1 \\ 0_{(N-k(i)) \times 1} \end{pmatrix}.$$

We have

$$\begin{aligned} d\mathcal{F}_i &= s_{k(i)}^T (\mathbf{H}_i d\mathbf{B}^T \mathbf{B} \ell^{(i)} + \mathbf{H}_i \mathbf{B}^T d\mathbf{B} \ell^{(i)}). \\ \text{Tr}(d\mathcal{F}_i) &= \text{Tr} \left((\mathbf{B} \ell^{(i)} s_{k(i)}^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T s_{k(i)} \ell^{(i)T})^T d\mathbf{B} \right). \end{aligned}$$

Consequently,

$$\frac{\partial \mathcal{F}_i}{\partial \mathbf{B}} = \mathbf{B} \ell^{(i)} s_{k(i)}^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T s_{k(i)} \ell^{(i)T}$$

and

$$\frac{\partial \widehat{\Gamma}_T}{\partial \mathbf{B}}(\mathbf{B}) = -\frac{1}{T} \sum_{i=1}^T D_i (\mathbf{B} \ell^{(i)} s_{k(i)}^T \mathbf{H}_i + \mathbf{B} \mathbf{H}_i^T s_{k(i)} \ell^{(i)T}) \quad (\text{A.15})$$

with

$$D_i = 2 \left(u(x^*(\ell^{(i)}); \ell^{(i)}) - u(x^*(\widehat{\ell}^{(i)}); \ell^{(i)}) \right). \quad (\text{A.16})$$

References

- [1] M. P. Tcheou, L. Lovisolo, M. V. Ribeiro, E. A. B. da Silva, M. A. M. Rodrigues, J. M. T. Romano, P. S. R. Diniz, The compression of electric signal waveforms for smart grids: State of the art and future trends, IEEE Transactions on Smart Grid 5 (1) (2014) 291–302.

- [2] M. Aiello, G. A. Pagani, The smart grid’s data generating potentials, in: 2014 Federated Conference on Computer Science and Information Systems, 2014, pp. 9–16.
- [3] J. Kraus, P. Štěpán, L. Kukačka, Optimal data compression techniques for smart grid and power quality trend data, in: 2012 IEEE 15th International Conference on Harmonics and Quality of Power, 2012, pp. 707–712.
- [4] Ö. N. Gerek, D. G. Ece, Compression of power quality event data using 2d representation, *Electric Power Systems Research* 78 (6) (2008) 1047–1052.
- [5] M. Zhang, K. Li, Y. Hu, A high efficient compression method for power quality applications, *IEEE Transactions on Instrumentation and Measurement* 60 (6) (2011) 1976–1985.
- [6] M. R. Chowdhury, S. Tripathi, S. De, Adaptive multivariate data compression in smart metering internet of things, *IEEE Transactions on Industrial Informatics* 17 (2) (2021) 1287–1297. doi:10.1109/TII.2020.2981382.
- [7] R. Gray, D. Neuhoff, Quantization, *IEEE Transactions on Information Theory* 44 (6) (1998) 2325–2383. doi:10.1109/18.720541.
- [8] T. Berger, J. Gibson, Lossy source coding, *IEEE Transactions on Information Theory* 44 (6) (1998) 2693–2723. doi:10.1109/18.720552.
- [9] C. Zhang, N. Khalfet, S. Lasaulce, V. Varma, S. Tarbouriech, Payoff-oriented quantization and application to power control, in: 2017 15th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt), IEEE, 2017, pp. 1–6.
- [10] N. Shlezinger, Y. C. Eldar, M. R. Rodrigues, Hardware-limited task-based quantization, *IEEE Transactions on Signal Processing* 67 (20) (2019) 5223–5238.
- [11] P. Popovski, O. Simeone, F. Boccardi, D. Gündüz, O. Sahin, Semantic-effectiveness filtering and control for post-5g wireless connectivity, *Journal of the Indian Institute of Science* 100 (2020) 435–443.
- [12] Q. Lan, D. Wen, Z. Zhang, Q. Zeng, X. Chen, P. Popovski, K. Huang, What is semantic communication? A view on conveying meaning in the era of machine intelligence, *Journal of Communications and Information Networks* 6 (4) (2021) 336–371.

- [13] M. Kalfa, M. Gok, A. Atalik, B. Tegin, T. M. Duman, O. Arikan, Towards goal-oriented semantic signal processing: Applications and future challenges, *Digital Signal Processing* 119 (2021) 103134.
- [14] E. C. Strinati, S. Barbarossa, 6G networks: Beyond Shannon towards semantic and goal-oriented communications, *Computer Networks* 190 (2021) 107930.
- [15] D. Gündüz, Z. Qin, I. E. Aguerri, H. S. Dhillon, Z. Yang, A. Yener, K. K. Wong, C.-B. Chae, Beyond transmitting bits: Context, semantics, and task-oriented communications, *IEEE Journal on Selected Areas in Communications* 41 (1) (2022) 5–41.
- [16] Y. Bar-Hillel, R. Carnap, Semantic information, *The British Journal for the Philosophy of Science* 4 (14) (1953) 147–157.
- [17] O. Goldreich, B. Juba, M. Sudan, A theory of goal-oriented communication, *Journal of the ACM (JACM)* 59 (2) (2012) 1–65.
- [18] M. Kountouris, N. Pappas, Semantics-empowered communication for networked intelligent systems, *IEEE Communications Magazine* 59 (6) (2021) 96–102.
- [19] C. Chaccour, W. Saad, M. Debbah, Z. Han, H. V. Poor, Less data, more knowledge: Building next generation semantic communication networks, *arXiv preprint arXiv:2211.14343* (2022).
- [20] T. M. Getu, G. Kaddoum, M. Bennis, Making sense of meaning: A survey on metrics for semantic and goal-oriented communication, *IEEE Access* (2023).
- [21] S. Roucos, R. Schwartz, J. Makhoul, A segment vocoder at 150 b/s, in: *ICASSP'83. IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 8, IEEE, 1983, pp. 61–64.
- [22] F. K. Soong, A phonetically labeled acoustic segment (plas) approach to speech analysis-synthesis, in: *International Conference on Acoustics, Speech, and Signal Processing*, IEEE, 1989, pp. 584–587.
- [23] Y. Hirata, S. Nakagawa, A 100bit/s speech coding using a speech recognition technique, in: *First European Conference on Speech Communication and Technology*, 1989.
- [24] C. V. Lopes, A. Chadha, A 40 bps speech coding scheme, in: *GLOBECOM'03. IEEE Global Telecommunications Conference (IEEE Cat. No. 03CH37489)*, Vol. 4, IEEE, 2003, pp. 2223–2226.

- [25] C. Zhang, S. Lasaulce, M. Hennebel, L. Saludjian, P. Panciatici, H. V. Poor, Decision-making oriented clustering: Application to pricing and power consumption scheduling, *Applied Energy* 297 (2021) 117106.
- [26] H. Zou, C. Zhang, S. Lasaulce, L. Saludjian, P. Panciatici, Decision-oriented communications: Application to energy-efficient resource allocation, in: 2018 6th International Conference on Wireless Networks and Mobile Communications (WINCOM), IEEE, 2018, pp. 1–6.
- [27] N. Shlezinger, Y. C. Eldar, Deep task-based quantization, *Entropy* 23 (1) (2021) 104.
- [28] H. Zou, C. Zhang, S. Lasaulce, L. Saludjian, H. V. Poor, Goal-oriented quantization: Analysis, design, and application to resource allocation, *IEEE Journal on Selected Areas in Communications* 41 (1) (2023) 42–54.
- [29] A. Mostaani, T. X. Vu, S. Chatzinotas, B. Ottersten, Task-oriented data compression for multi-agent communications over bit-budgeted channels, *IEEE Open Journal of the Communications Society* 3 (2022) 1867–1886.
- [30] Y. Sun, H. Zou, S. Lasaulce, M. Kieffer, L. Saludjian, A new approach of data pre-processing for data compression in smart grids, in: 2019 International Conference on Wireless Networks and Mobile Communications (WINCOM), IEEE, 2019, pp. 1–6.
- [31] O. Beaude, S. Lasaulce, M. Hennebel, J. Daafouz, Minimizing the impact of EV charging on the electricity distribution network, in: 2015 European Control Conference (ECC), 2015, pp. 648–653.
- [32] M. Shinwari, A. Youssef, W. Hamouda, A water-filling based scheduling algorithm for the smart grid, *IEEE Transactions on Smart Grid* 3 (2) (2012) 710–719.
- [33] W. Yu, G. Ginis, J. Cioffi, Distributed multiuser power control for digital subscriber lines, *IEEE Journal on Selected Areas in Communications* 20 (5) (2002) 1105–1115.
- [34] S. Mallat, *A wavelet tour of signal processing*, Elsevier, 1999.
- [35] J. Ballé, V. Laparra, E. P. Simoncelli, End-to-end optimized image compression, *arXiv preprint arXiv:1611.01704* (2016).

- [36] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, N. Johnston, Variational image compression with a scale hyperprior, arXiv preprint arXiv:1802.01436 (2018).
- [37] J. Lee, S. Cho, S.-K. Beack, Context-adaptive entropy model for end-to-end optimized image compression, arXiv preprint arXiv:1809.10452 (2018).
- [38] Y. Linde, A. Buzo, R. Gray, An algorithm for vector quantizer design, IEEE Transactions on communications 28 (1) (1980) 84–95.
- [39] Ausgrid, Solar home half-hour data - 1 july 2012 to 30 june 2013 (2014).
URL <https://www.ausgrid.com.au/Industry/Innovation-and-research/Data-to-share>
- [40] Tensorflow learned data compression, https://www.tensorflow.org/tutorials/generative/data_compression.
- [41] J. Dai, P. Zhang, K. Niu, S. Wang, Z. Si, X. Qin, Communication beyond transmitting bits: Semantics-guided source and channel coding, IEEE Wireless Communications (2022).
- [42] Y. Zhang, H. Zhao, K. Cao, L. Zhou, Z. Wang, Y. Liu, J. Wei, A highly reliable encoding and decoding communication framework based on semantic information, Digital Communications and Networks (2023).