Computation of the Multivariate Gaussian Rate-Distortion-Perception Function

Giuseppe Serra*, Photios A. Stavrou*, Marios Kountouris*†
*Communication Systems Department, EURECOM, Sophia-Antipolis, France
†Department of Computer Science and Artificial Intelligence, Andalusian Research Institute in Data Science and Computational Intelligence (DaSCI), University of Granada, Spain
{giuseppe.serra,fotios.stavrou,marios.kountouris}@eurecom.fr

***Abstract*—In this paper, we propose a generic method for computing the rate-distortion-perception function (RDPF) of a multivariate Gaussian source under tensorizable distortion and perception metrics. Through the assumption of a jointly Gaussian reconstruction, we establish that the optimal solution of the RDPF belongs to the vector space spanned by the eigenvector of the source covariance matrix. Consequently, the multivariate optimization problem can be expressed as a function of the scalar Gaussian RDPFs of the source marginals, constrained by global distortion and perception levels. Utilizing this result, we devise an alternating minimization scheme based on the block nonlinear Gauss–Seidel method. This scheme solves optimally the optimization problem while identifying the optimal stage-wise distortion and perception levels. Furthermore, the associated algorithmic embodiment is provided, along with the convergence and the rate of convergence characterization. Lastly, in the regime of "perfect realism", we provide the analytical solution for the multivariate Gaussian RDPF. We corroborate our findings with numerical simulations and draw connections to existing results.**

I. Introduction

The recently established field of rate-distortion-perception (RDP) theory has attracted major interest in the information theory community. Appearing simultaneously by Blau and Michaeli in [1] and Matsumoto in [2], [3], the RDP framework proposes a generalization of the classical rate-distortion (RD) theory, originally envisioned by Shannon [4]. Echoing the increasing body of research highlighting the limitations of solely focusing on distortion minimization in the reconstructed signals (see e.g., [5]–[9]), the RDP framework focuses on the concept of perceptual quality, which refers to the property of a sample to appear pleasing from a human perspective. This is enacted by extending the classical single-letter RD formulation, incorporating a divergence constraint between the source distribution and its estimation at the destination. The divergence constraint acts as a proxy for human perception, quantifying the satisfaction experienced when utilizing the data. Moreover, this divergence constraint may have multiple interpretations and can be seen as a semantic quality metric, measuring the relevance of the reconstructed source from the observer's perspective [10].

Multiple coding theorems have been developed for the RDP framework. Under the assumption of infinite common randomness between the encoder and decoder, Theis and Wagner in [11] prove a coding theorem for stochastic variable-length codes in both one-shot and asymptotic regimes. Chen *et. al.* in [12] derive coding theorems for the asymptotic regime, analyzing the three distinct operative conditions; when the encoder and the decoder share or not common randomness, and when both have private randomness. Originally in the context of the output-constrained RDF, but also valid for the "perfect realism" RDPF case, Saldi *et. al.* [13] provide coding theorems for when only finite common randomness between encoder and decoder is available. Similar results are also presented, specifically for the RDPF, by Wagner in [14].

Similarly to the classical RD theory, the mathematical embodiment of the RDP framework is represented by the rate-distortion-perception function (RDPF), which, as its classical counterpart, does not enjoy a general analytical solution. However, despite the general complexity, certain closed-form expressions have been developed for specific categories of sources. For instance, binary sources subject to Hamming distortion and total variation distance have closed-form expressions, as discussed in [1]. Similarly, for scalar Gaussian sources under mean squared-error (MSE) distortion closed-form expressions have been established for and squared Wasserstein-2 distance in [15] and for Kullback–Leibler (KL) divergence, Geometric Jensen-Shannon divergence, and Hellinger distance in [16]. In [17] the authors have recently derived closed-form parametric expressions using two distinct reverse water-filling algorithms for the case of the multivariate Gaussian source under MSE distortion and either Wasserstein-2 distance or KL divergence perception.

The complexity associated with deriving analytical solutions for the RDPF has prompted research into computational methods for its estimation. Toward this end, both Serra *et al.* in [18] and Chen *et. al.* in [19] propose algorithms for the computation of the RDPF for general discrete sources, studying respectively the cases where the perception constraint belongs to the family of f-divergences or Wasserstein distances, KL divergence, and total variation distance. Alternatives to these algorithmic approaches for the computation of the RDPF rely on data-driven solutions employing generative adversarial networks [1], [15], [20] which, despite providing a practical framework for data-driven codec optimization, are highly computational- and data-intensive and lack generalization capabilities.

A. Our Approach and Contributions

This work focuses on the design of algorithmic solutions for the computation of the (upper) bounds obtained assuming a multivariate Gaussian RDPF. To this end, in Section III-A

979-8-3503-8284-6/24/$31.00 ©2024 IEEE

we prove that, under convex and tensorizable distortion and divergence metrics and assuming jointly Gaussian reconstruction, the optimal solution of the multivariate Gaussian RDPF belongs to the space of the eigenvectors of the source covariance matrix. In other words, an optimal solution can be characterized as a set of design matrices commuting by pairs[1]. The resulting optimization problem can be solved optimally using the *block nonlinear Gauss–Seidel method* [22], an alternating minimization approach for which we develop an algorithmic embodiment (see Algorithm 1). For the specific algorithm, we show convergence (Theorem 1) and provide an upper bound on the worst-case convergence rate (Theorem 2). In Section III-B, we provide an application example of Algorithm 1 considering the case of MSE distortion and squared Wasserstein-2 divergence constraints. Leveraging the analytical results of this example, in Section III-C we characterize the closed-form solution of the multivariate Gaussian RDPF in the regime of *perfect realism* (Corollary 2). Specifically, this result provides the optimal stagewise distortion allocation, i.e., the distortion introduced on each dimension of the Gaussian reconstruction, according to what is interpreted as an *adaptive water-level* (see Fig. 2).

B. Notation

Given a Polish space $\mathcal{X}$, we denote by $(\mathcal{X}, \mathbb{B}(\mathcal{X}))$ the Borel measurable space induced by the metric, with $\mathcal{P}(\mathcal{X})$ denoting the set of probability measures defined thereon. For a random variable X defined on $(\mathcal{X}, \mathbb{B}(\mathcal{X}))$, we denote with $p_X \in \mathcal{P}(\mathcal{X})$ its probability measure and with μ_X and Σ_X its mean and covariance matrix, respectively. Given two random variables X and Y, $X \perp Y$ indicates their statistical independence. We denote the diagonal of a square matrix by $\operatorname{diag}(\cdot)$. Given a square matrix $A \in \mathbb{R}^d$, we indicate its positive definiteness (resp. positive semi-definiteness) with the notation $A \succ 0$ (resp. $A \succeq 0$) and the set of its eigenvalues with $\{\lambda_{A,i}\}_{i=1,\ldots,d}$.

II. RDPF on General Alphabets

We commence by providing the definition and some properties of the RDPF for general alphabets. Such preliminaries can be found for instance in [1].

Definition 1. *(RDPF) Let a source X be a random variable distributed according to $p_X \in \mathcal{P}(\mathcal{X})$. Then, the RDPF for a source $X \sim p_X$ under the distortion measure $\Delta : \mathcal{X}^2 \to \mathbb{R}_0^+$ and divergence function $d : \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}_0^+$ is defined as follows:*

$$R(D,P) = \min_{P_{\hat{X}|X}} I(X,\hat{X}) \tag{1}$$

$$\text{s.t.} \quad \mathbb{E}\Big[\Delta(X,\hat{X})\Big] \le D \tag{2}$$

$$d(p_X||p_{\hat{X}}) \le P \tag{3}$$

where the minimization is among all conditional distributions $P_{\hat{X}|X} : \mathcal{X} \to \mathcal{P}(\hat{\mathcal{X}})$.

[1]For the definition of matrix commutability, we refer the reader to [21, Section 0.7.7]

We point out the following remark on Definition 1.

Remark 1. *(On Definition 1) Following [1], it can be shown that* (1) *has some useful functional properties, under mild regularity conditions. In particular, [1, Theorem 1] shows that $R(D,P)$ is (i) monotonically non-increasing function in both $D \in [D_{\min}, D_{\max}] \subset [0,\infty)$ and $P \in [P_{\min}, P_{\max}] \subset [0,\infty)$; (ii) convex if the divergence $d(\cdot||\cdot)$ is convex in its second argument.*

III. Main Results

A. A Generic Alternating Minimization Approach

The main goal of this section is to provide a general yet simple algorithmic approach to compute the bounds obtained assuming a vector Gaussian RDPF, able to tackle a large set of divergence metrics. We begin with the characterization of (1) for jointly Gaussian random variables under distortion measure $\Delta(\cdot,\cdot)$ and divergence measure $d(\cdot||\cdot)$.

Problem 1. *Given a Gaussian source $X \sim \mathcal{N}(\mu_X, \Sigma_X)$, $\Sigma_X \succ 0$, assume that the reconstructed random vector $\hat{X} \in \mathbb{R}^N$ is chosen such that the joint tuple $(X,\hat{X})$ is jointly Gaussian. Then, the reconstructed message admits a linear (forward) realization of the form $\hat{X} = AX + W$, where $A \in \mathbb{R}^{N\times N}$, $W \sim \mathcal{N}(\mu_W, \Sigma_W)$, $W \perp X$ and $\Sigma_W \succeq 0$, such that $\mu_{\hat{X}} = A\mu_X + \mu_W$ and $\Sigma_{\hat{X}} = A\Sigma_X A^T + \Sigma_W$. Moreover, we can cast* (1)-(3) *as follows:*

$$R(D,P) \le R^G(D,P) = \min_{\substack{A\in\mathbb{R}^{N\times N}, \Sigma_W \succeq 0 \\ \mathbb{E}[\Delta(X,\hat{X})] \le D \\ d(p_X||p_{\hat{X}}) \le P}} \frac{1}{2}\log\Big(\frac{|A\Sigma_X A^T + \Sigma_W|}{|\Sigma_W|}\Big). \tag{4}$$

Remark 2. *The upper bound $R(D,P) \le R^G(D,P)$ arises from the assumption of jointly Gaussian reconstruction and holds with equality only for specific distortion $\Delta(\cdot,\cdot)$ and perception $d(\cdot||\cdot)$ metrics. Such cases include for example* MSE *distortion and either the squared Wasserstein-2 distance [15] and the reverse KL-distance [16].*

Under the assumption of tensorizable fidelity metrics $\mathbb{E}[\Delta(\cdot,\cdot)]$ and $d(\cdot||\cdot)$, i.e.,

$$\mathbb{E}\Big[\Delta(X,\hat{X})\Big] \ge \sum_{i=1}^N g\Big(\mathbb{E}\Big[\Delta(X_i,\hat{X}_i)\Big]\Big)$$

$$d(p_X||p_{\hat{X}}) \ge \sum_{i=1}^N h\Big(d(p_{X_i}||p_{\hat{X}_i})\Big)$$

with $g(\cdot)$ and $h(\cdot)$ convex functions dependent on the fidelity metrics, applying [23, Lemma 2] in (4) leads to the following lower bound

$$R^G(D,P) \overset{(\star)}{\ge} \min_{\substack{\lambda_{A,i}, \lambda_{\Sigma_W,i} \\ \sum_{i=1}^N g(\mathbb{E}[\Delta(X_i,\hat{X}_i)]) \le D \\ \sum_{i=1}^N h(d(p_{X_i}||p_{\hat{X}_i})) \le P}} \sum_{i=1}^N \frac{1}{2}\log\Big(1 + \frac{\lambda_{A,i}{}^2 \lambda_{X,i}}{\lambda_{W,i}}\Big) \tag{5}$$

where $(\star)$ holds with equality if the triplet (A, Σ_W, Σ_X) commute by pairs. The sufficient condition that achieves the lower bound in (5) can be easily satisfied[2], hence one can replace inequality with equality without loss of generality. This observation does not generalize beyond i. i. d. random vectors.

Proposed alternating minimization method: To solve (5), we first introduce the (vector) optimization variables $\mathbf{D} = [D_i]_{i\in 1:N}$ and $\mathbf{P} = [P_i]_{i\in 1:N}$ such that

$$D_i = \mathbb{E}\Big[\Delta(X, \hat{X}_i)\Big], \quad P_i = d(p_{X_i}||p_{\hat{X}_i}), \quad \forall i \in 1:N.$$

Once the slack variables above are substituted in (5), we yield

$$R^G(D,P) = \min_{\substack{\mathbf{D},\mathbf{P} \\ \sum_{i=1}^N g(D_i)\leq D \\ \sum_{i=1}^N h(P_i)\leq P}} \sum_{i=1}^{N} R_i^G(D_i, P_i) \tag{6}$$

where $R_i^G(D_i, P_i)$ corresponds to the stagewise RDPF given by

$$R_i^G(D_i,P_i) = \min_{\substack{\lambda_{A,i},\lambda_{\Sigma_W,i} \\ D_i=\mathbb{E}\left[\Delta(X_i,\hat{X}_i)\right] \\ P_i=d(p_{X_i}||p_{\hat{X}_i})}} \tfrac{1}{2}\log\Big(1+\tfrac{\lambda_{A,i}{}^2\lambda_{X,i}}{\lambda_{W,i}}\Big).$$

We note that (6) defines three distinct "rate region" cases, which - when combined - describe the entirety of the RDPF. In particular, we distinguish the cases where either only the distortion constraint is active (hereinafter referred to as **Case I**), or only the perception constraint is active (hereinafter referred to as **Case II**), or both are active (hereinafter referred to as **Case III**). We remark that **Case III** is the most interesting one, as the other two cases easily follow from its computation.

To find the optimal pair $(\mathbf{D}^*, \mathbf{P}^*)$ in (6) we resort to an application of an alternating minimization technique. Specifically, we define the following two subproblems of (6):

- For fixed $\mathbf{P}$, (6) simplifies to

$$\min_{\mathbf{D}} \sum_{i=1}^{N} R_i^G(D_i,P_i) \quad \text{s.t.} \quad \sum_{i=1}^{N} g(D_i) \leq D. \tag{7}$$

- For fixed $\mathbf{D}$, (6) simplifies to

$$\min_{\mathbf{P}} \sum_{i=1}^{N} R_i^G(D_i,P_i) \quad \text{s.t.} \quad \sum_{i=1}^{N} h(P_i) \leq P. \tag{8}$$

The solutions of optimization problems (7) and (8) are of central interest since their alternate application forms a minimization scheme that can optimally solve (6). This class of alternating minimization schemes is referred to as *block nonlinear Gauss-Seidel (GS) method* [22]. In the following, we prove the convergence to an optimal point of a GS scheme based on the solutions of (7) and (8).

Theorem 1. *(Convergence) Let the optimization problem* (6) *be defined for finite distortion and perception levels* (D, P). *Let* $(\mathbf{D}^{(0)}, \mathbf{P}^{(0)})$ *be an initial point and let the sequence* $\{(\mathbf{D}^{(n)}, \mathbf{P}^{(n)}) : n = 1, 2, \ldots\}$ *be the sequence obtained by the alternating optimization of problems* (7) *and* (8). *Then the sequence has a limit* $\lim_{n\to\infty}(\mathbf{D}^{(n)}, \mathbf{P}^{(n)}) = (\mathbf{D}^*, \mathbf{P}^*)$ *and the limit is an optimal solution of* (6).

Proof: See [23, Theorem 5]. ∎

Although the GS method simplifies the optimization of (6), subproblems (7) and (8) remain constrained optimization problems whose solution may not be easily approached. Therefore, to further simplify the optimization task, we define the unconstrained optimization problem associated with (6). Let $s = (s_D, s_P)$, with $s_D > 0$ and $s_P > 0$, be the vector of Lagrangian multipliers respectively associated with the distortion (s_D) and perception (s_P) constraints. Then, the Lagrangian function $L_{R^G}(s)$ associated with (6) is defined as

$$\begin{aligned}\min_{\mathbf{D},\mathbf{P}} L_{R^G}(\mathbf{D},\mathbf{P},s) &= \min_{\mathbf{D},\mathbf{P}} \sum_{i=1}^{N} R_i^G(D_i,P_i) \\ &+ s_D \sum_{i=1}^{N} g(D_i) + s_P \sum_{i=1}^{N} h(P_i).\end{aligned} \tag{9}$$

Similarly to the constrained case, the optimal pair $(\mathbf{D}^*, \mathbf{P}^*)$ in (9) can be characterized through an alternate minimization scheme. Hence, the associated subproblems are

- For fixed $\mathbf{P}$,

$$\min_{\mathbf{D}} \sum_{i=1}^{N} R_i^G(D_i,P_i) + s_D \sum_{i=1}^{N} g(D_i) \tag{10}$$

- For fixed $\mathbf{D}$,

$$\min_{\mathbf{P}} \sum_{i=1}^{N} R_i^G(D_i,P_i) + s_P \sum_{i=1}^{N} h(P_i). \tag{11}$$

Assume the Lagrangian multiplier vector s to be given and let $\mathbf{D}_s^*$ and $\mathbf{P}_s^*$ be the optimal solutions obtained from the Gauss-Seidel method for subproblems (10) and (11), respectively. Furthermore, let $D_s = \sum_{i=1}^N g(D_{s,i}^*)$ and $P_s = \sum_{i=1}^N h(P_{s,i}^*)$. Then, due Lagrangian duality [25], we can compute $R(\mathbf{D}_s^*, \mathbf{P}_s^*)$ as

$$R^G(D_s, P_s) = L_{R^G}(\mathbf{D}_s^*, \mathbf{P}_s^*, s) - s_D D_s - s_P P_s. \tag{12}$$

Remark 3. *The assumption of strictly positive Lagrangian multipliers* (s_D, s_P) *implies finite* (D, P) *levels in Theorem 1, therefore guaranteeing the convergence of Algorithm 1. However, under the assumption of bounded perception metric* $d(\cdot||\cdot)$, *the case* $s_P = 0$ *does not violate the assumptions of Theorem 1. In this regime, the perception constraint of Problem 1 is inactive, making the problem equivalent to the classical RD function problem.*

To characterize the worst-case performance of Algorithm 1, we provide an upper bound on its convergence rate.

Theorem 2. *(Upper bound on the Convergence Rate) Let* $\{(\mathbf{D}^{(n)}, \mathbf{P}^{(n)})\}_{n=0,\ldots,T}$ *be the sequence of iterations generated by Algorithm 1 in* T *iterations and let* $(\mathbf{D}^*, \mathbf{P}^*)$ *be a*

[2] This is because the matrices (A, Σ_W) are design variables and can be chosen such that they have the *same eigenvectors* as Σ_X, for details see, e.g., [24, Proposition 1].

minimizer of $L_{R^G}(\cdot,\cdot)$. Then, there exists positive and finite constant C such that Alg. 1 guarantees that

$$L_{R^G}(\mathbf{D}^{(T)},\mathbf{P}^{(T)}) - L_{R^G}(\mathbf{D}^*,\mathbf{P}^*) \leq \frac{C}{T} \tag{13}$$

i.e., the asymptotic rate of convergence of Alg. 1 is upper bounded by $\mathcal{O}\big(\frac{1}{T}\big)$ (sublinear convergence rate).

Proof: See [23, Theorem 6]. ■

Algorithm 1 Algorithm of Theorem 1

Require: source distribution $p_X = \mathcal{N}(\mu_X, \Sigma_X)$ with $\Sigma_X \succ 0$; Lagrangian parameters $s = (s_D, s_P)$ with $s_D > 0$ and $s_P > 0$; error tolerances ϵ; initial point $(\mathbf{D}^{(0)}, \mathbf{P}^{(0)})$.

```
ω ← +∞; n ← 1;
while ω > ϵ do
    D^(n) ← Solution Problem (10) for (P^(n−1), s_D)
    P^(n) ← Solution Problem (11) for (D^(n), s_P)
    ω ← ||(D^(n), P^(n)) − (D^(n−1), P^(n−1))||_2
    n ← n + 1
end while
```

Ensure: $D = \sum_{i=1}^N g(D_i^{(n)})$, $P = \sum_{i=1}^N h(P_i^{(n)})$, $R^G(D,P) = \sum_{i=1}^N R_i^G(D_i^{(n)}, P_i^{(n)})$.

In Algorithm 1 we implement the alternating minimization scheme of Theorem 1 using the unconstrained formulation previously discussed, which allows for the computation of any multivariate Gaussian RDPF of the form characterized in Problem 1 as long as we can have a characterization of the problem for the univariate case.

B. Application of the Alternating Minimization Approach

In this subsection, we apply the theoretical results of Subsection III-A to the specific case of MSE distortion and squared Wasserstein-2 perception constraints. We remark that similar specializations can be developed for other divergence measures whenever the scalar RDPF characterization is available (see e.g., [16]), as shown in the numerical examples of Section IV. We start by solving the subproblems (7) and (8). To this end, we leverage the RDPF provided in [15] to characterize the function $R_i(D_i, P_i)$, which is the stagewise RDPF for the i^{th} dimension, under MSE distortion metric and W_2^2 perception metric, for a Gaussian source $X_i \sim \mathcal{N}(0, \lambda_{\Sigma_X,i})$. Furthermore, using the tensorization of the squared Wasserstein-2 distance (for details see [23, Proposition 3]), the functional form of as the auxiliary optimization variables $\mathbf{D} = [D_i]_{i\in 1:N}$ and $\mathbf{P} = [P_i]_{i\in 1:N}$ is defined as follows:

$$D_i = \mathbb{E}\Big[||X_i - \hat{X}_i||^2\Big] = (1-\lambda_{A,i})^2\lambda_{\Sigma_X,i} + \lambda_{\Sigma_W,i}$$
$$P_i = \mathrm{W}_2^2(p_{X_i}, p_{\hat{X}_i}) = \lambda_{\Sigma_X,i} - \sqrt{{\lambda_{A,i}}^2\lambda_{\Sigma_X,i} + \lambda_{\Sigma_W,i}}$$

where $\hat{X}_i \sim \mathcal{N}(0, \lambda_{\Sigma_{X_i},i})$ is the stagewise linear realization of the form $\hat{X}_i = \lambda_{A,i}X_i + W_i$ with $W_i \sim \mathcal{N}(0, \lambda_{\Sigma_W,i})$. Instead, the tensorization functions $g(\cdot)$ and $h(\cdot)$, introduced in (6), are equal to the identity function, i.e. $g(\cdot) = h(\cdot) = \mathrm{id}(\cdot)$.

We derive the optimal solution of subproblem (7) for the described case in the following theorem.

Theorem 3. *Let the Lagrangian multiplier $s_D > 0$ be given. Then, for fixed $\mathbf{P}$, the optimal stagewise distortions levels $\mathbf{D}^*(\mathbf{P}) = [D_i^*(P_i)]_{i\in i:N} \in \mathcal{S}$ achieving the minimum of (10) are given by*

$$\begin{aligned} D_i^* &= P_i + 2\sqrt{\lambda_{\Sigma_X,i}}\Big(\sqrt{\lambda_{\Sigma_X,i}} - \sqrt{P_i}\Big) \\ &+ \left(\frac{1}{2s_D} - \sqrt{4\lambda_{\Sigma_X,i}(\sqrt{\lambda_{\Sigma_X,i}} - \sqrt{P_i})^2 + \frac{1}{4s_D^2}}\right). \end{aligned} \tag{14}$$

Proof: See [23, Theorem 7]. ■

We now move to the solution of subproblem (8), as described in the following theorem.

Theorem 4. *Let the Lagrangian multiplier $s_P > 0$ be given. Then, for fixed $\mathbf{D}$, the optimal stagewise perception levels $\mathbf{P}^*(\mathbf{D}) = [P_i^*(D_i)]_{i\in i:N} \in \mathcal{S}$ achieving the minimum of (11) can be characterized as the zeros of the vector function $T(\cdot): \mathbb{R}^N \to \mathbb{R}^N$ where each component is defined as*

$$T_i(x) \triangleq \frac{\partial R_i^G(D_i, P_i)}{\partial P_i}\bigg|_{(D_i, x_i)} + s_P. \tag{15}$$

Proof: See [23, Theorem 8]. ■

Corollary 1. *Let $T_i : \mathbb{R}^N \to \mathbb{R}$ be the i^{th} component of the vector function $T(\cdot)$ defined in Theorem 4. Then, T_i is a continuous and non-decreasing function on $\mathbb{R}$. Furthermore, T_i has at least one root in $\mathcal{S}$.*

Proof: See [23, Corollary 2]. ■

Even though a closed-form solution for subproblem (8) cannot be directly derived, the optimal $\mathbf{P}^*$ can be found as zeros of the functions $\{T_i\}_{i\in 1:N}$. Corollary 1 guarantees that the roots of $\{T_i\}_{i\in 1:N}$ can be numerically approximated using root-finding methods, e.g., bisection method [26, Sec. 2.1].

C. *RDPF under* perfect realism *regime*

Using the results of Theorem 3, we can characterize the optimal distortion levels in the *perfect realism* regime [12], [14], i.e., $\mathbf{P} = \mathbf{0}$.

Corollary 2. *Consider the optimization problem (6) for perception level $P = 0$. Then, for a given Lagrangian multiplier $s_D > 0$, the optimal solution $\mathbf{D}^* = [D_i^*]_{i\in 1:N}$ is given by*

$$D_i^* = 2\lambda_{\Sigma_X,i} + \frac{1}{2s_D} - \sqrt{4\lambda_{\Sigma_X,i}^2 + \frac{1}{4s_D^2}} \tag{16}$$

such that the distortion level $D = \sum_{i=1}^N D_i^*$.

Proof: (16) is obtained from (14) for $P_i = 0$. ■

We stress the following technical remarks for Corollary 2.

Remark 4. *The optimal solution* $\mathbf{D}^*$ *is well defined in the limit* $s_D \to 0$*, since* $\lim_{s_D \to 0} D_i^* = 2\lambda_{\Sigma_X,i}$.

Remark 5. *In the water-filling solution of the classical multivariate Gaussian RD function, the optimal solution* $\mathbf{D}^*_{RD} = [D^*_{i,RD}]_{i \in 1:N}$ *for* $s_D > 0$ *is shown to be*

$$D^*_{i,RD} = \min(w(s_D), \lambda_{\Sigma_X,i}) \qquad w(s_D) = \frac{1}{2s_D}$$

where water-level $w(s_D)$ *is independent of the marginal and the* $\min(\cdot)$ *operation is required to guarantee that* $D^*_{i,RD}$ *belongs to the constraint set. Heuristically, one can interpret it as the* i*-th source component being discarded in the reconstruction whenever* $w(s_D) \geq \lambda_{\Sigma_X,i}$*, thus upper bounding the maximum distortion observed in the* i*-th component. On the other hand, the solution identified in* (16)*, and in general the results of Theorem 3, can be seen as an adaptive water-level. Indeed, in* (16)*,* $w(s_D)$ *shows a dependency to* $\lambda_{\Sigma_X,i}$ *that guarantees that all source components are present in the reconstructed signal.*

IV. Numerical Results

In this section, we demonstrate numerical simulations using Algorithm 1. All the numerical experiments have been conducted considering a multivariate Gaussian source $X \sim \mathcal{N}(0, \Sigma_X)$ with $\Sigma_X = \mathrm{diag}([1, 3, 5])$.

1) RDPF Curves: In Fig. 1, we report the Gaussian RDPF under MSE distortion metric and either squared Wasserstein-2, squared Hellinger distance or KL divergence perception metric. Focusing on the Wasserstein-2 case (Fig. 1a), we compare the computed curve with the classical RD curve (black line), where in the latter the perception measure has been computed post hoc using the same divergence metric. The result confirms that, for bounded divergence measure, the RD curve delineates the boundary between the regions of *Case I* and *Case III* and can be obtained as an extreme case of RDPF surface. Additionally, the surface region of *Case I* can be retrieved by rigid translation of the boundary curve, as for any (D, P) point in the region the perception constraint is not active, turning the RDPF problem into the classical RD problem. The same observations can be extended for the case of the RDPF under Hellinger distance H^2 perception metric (Fig. 1b). However, in the case of KL divergence perception (Fig. 1c), a different behavior emerges in the limit case of $s_P \to 0$, i.e $P \to \infty$, due to the unbounded nature of the KL divergence. Therefore, the boundary between *Case III* and *Case I* cannot be computed and, following Remark 3, we impose a finite perception level P by setting $s_P \geq 10^{-3}$.

2) Adaptive Water-Level: In Fig. 2, we examine the per-dimension levels of distortion D_i^* and perception P_i^* between the source X and its reconstruction $\hat{X}$, comparing Algorithm 1 with the classical RD water-filling solution. The comparison is conducted considering a target distortion level $D = 6$ while varying the target perception level P. In the least constrained case ($P \approx 2$), the distortion level D_i^* closely follows the classical water-filling solution, whereas, in the more constrained case ($P \approx 0$), D_i^* is not distributed according to a uniform water-level and instead adapts to each marginal. Furthermore, as previously mentioned in Remark 5, values of distortions $D_1^* \geq \sigma^2_{X_1}$ can be observed for the first dimension for the latter case.

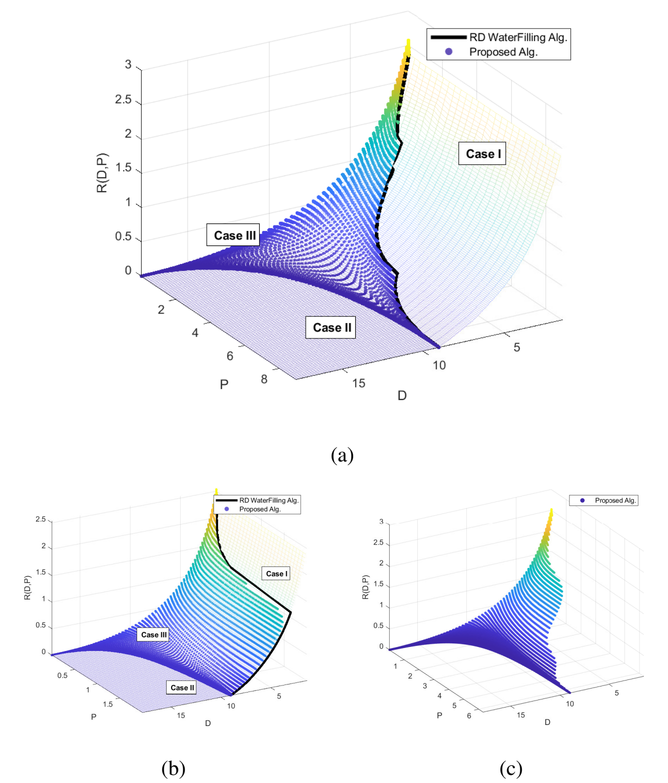

Fig. 1: $R(D, P)$ for a Gaussian source $X \sim \mathcal{N}(0, \Sigma_X)$ with $\Sigma_X = \mathrm{diag}([1, 3, 5])$ under MSE distortion metric and (a) W_2^2 perception metric, (b) H^2 Hellinger distance, and (c) KL divergence.

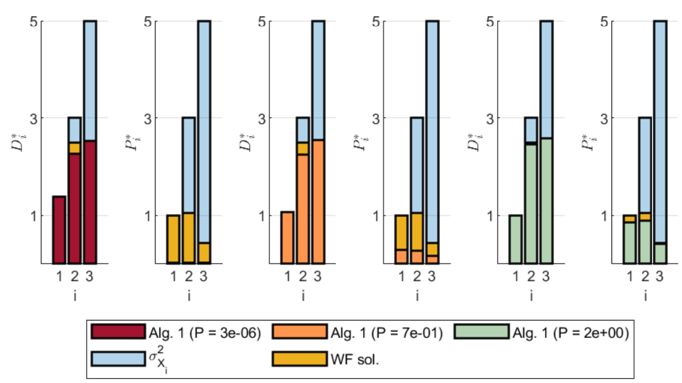

Fig. 2: Comparison of the per-dimension distortion D_i^* and perception P_i^* levels between the water-filling solution and Alg. 1, for $D = 6$.

Acknowledgement

This work is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101003431).

References

[1] Y. Blau and T. Michaeli, "Rethinking lossy compression: The rate-distortion-perception tradeoff," in *International Conference on Machine Learning*. PMLR, 2019, pp. 675–685.

[2] R. Matsumoto, "Introducing the perception-distortion tradeoff into the rate-distortion theory of general information sources," *IEICE Comm. Express*, vol. 7, no. 11, pp. 427–431, 2018.

[3] ——, "Rate-distortion-perception tradeoff of variable-length source coding for general information sources," *IEICE Comm. Express*, vol. 8, no. 2, pp. 38–42, 2019.

[4] C. E. Shannon, "Coding theorems for a discrete source with a fidelity criterion," *Institute of Radio Engineers, National Convention Record*, vol. 4, pp. 142–163, 1993.

[5] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 105–114.

[6] T. R. Shaham and T. Michaeli, "Deformation aware image compression," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[7] S. Kudo, S. Orihashi, R. Tanida, and A. Shimizu, "GAN-based image compression using mutual information maximizing regularization," in *Picture Coding Symposium*, 2019, pp. 1–5.

[8] D. Minnen, J. Ballé, and G. D. Toderici, "Joint autoregressive and hierarchical priors for learned image compression," in *Advances in Neural Information Processing Systems*, vol. 31, 2018.

[9] L.-H. Chen, C. G. Bampis, Z. Li, A. Norkin, and A. C. Bovik, "ProxIQA: A proxy approach to perceptual optimization of learned image compression," *IEEE Trans. Image Proc.*, vol. 30, pp. 360–373, 2021.

[10] M. Kountouris and N. Pappas, "Semantics-empowered communication for networked intelligent systems," *IEEE Commun. Mag.*, vol. 59, no. 6, pp. 96–102, 2021.

[11] L. Theis and A. B. Wagner, "A coding theorem for the rate-distortion-perception function," in *International Conference of Learning Representations (ICLR)*, 2021, pp. 1–5.

[12] J. Chen, L. Yu, J. Wang, W. Shi, Y. Ge, and W. Tong, "On the rate-distortion-perception function," *IEEE Journal on Selected Areas in Information Theory*, pp. 1–1, 2022.

[13] N. Saldi, T. Linder, and S. Yüksel, "Output constrained lossy source coding with limited common randomness," *IEEE Transactions on Information Theory*, vol. 61, no. 9, pp. 4984–4998, 2015.

[14] A. B. Wagner, "The rate-distortion-perception tradeoff: The role of common randomness," *arXiv preprint arXiv:2202.04147*, 2022.

[15] G. Zhang, J. Qian, J. Chen, and A. Khisti, "Universal rate-distortion-perception representations for lossy compression," *Advances in Neural Information Processing Systems*, vol. 34, pp. 11 517–11 529, 2021.

[16] G. Serra, P. A. Stavrou, and M. Kountouris, "On the computation of the Gaussian rate-distortion-perception function," 2023. [Online]. Available: https://arxiv.org/abs/2311.09190

[17] J. Qian, "On the rate-distortion-perception tradeoff for lossy compression," Ph.D. dissertation, McMaster University, October 2023, http://hdl.handle.net/11375/28976.

[18] G. Serra, P. A. Stavrou, and M. Kountouris, "Computation of rate-distortion-perception function under f-divergence perception constraints," in *Proc. IEEE Int. Symp. Inf. Theory*, 2023, pp. 531–536.

[19] C. Chen, X. Niu, W. Ye, S. Wu, B. Bai, W. Chen, and S.-J. Lin, "Computation of rate-distortion-perception functions with Wasserstein barycenter," *arXiv preprint arXiv:2304.14611*, 2023.

[20] O. Kirmemis and A. M. Tekalp, "A practical approach for rate-distortion-perception analysis in learned image compression," in *2021 Picture Coding Symposium (PCS)*, 2021, pp. 1–5.

[21] R. Horn and C. Johnson, *Matrix Analysis*. Cambridge University Press, 2012.

[22] L. Grippo and M. Sciandrone, "On the convergence of the block nonlinear Gauss–Seidel method under convex constraints," *Operations Research Letters*, vol. 26, no. 3, pp. 127–136, 1998.

[23] G. Serra, P. A. Stavrou, and M. Kountouris, "On the computation of the Gaussian rate-distortion-perception function," *IEEE Journal on Selected Areas in Information Theory*, pp. 1–1, 2024.

[24] P. A. Stavrou and M. Skoglund, "Asymptotic reverse waterfilling algorithm of NRDF for certain classes of vector Gauss–Markov processes," *IEEE Trans. Autom. Control*, vol. 67, no. 6, pp. 3196–3203, 2022.

[25] R. T. Rockafellar, *Convex analysis*. Princeton University Press, 1970, vol. 18.

[26] R. L. Burden, J. D. Faires, and A. M. Burden, *Numerical analysis*. Cengage learning, 2015.