

Analyzing α -divergence in Gaussian Rate-Distortion-Perception Theory

Martha V. Sourla*, Giuseppe Serra†, Photios A. Stavrou†, Marios Kountouris†‡

*School of Electrical & Computer Engineering, Technical University of Crete, Chania, Greece

†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

msourla@tuc.gr, {giuseppe.serra, fotios.stavrou}@eurecom.fr, mariosk@ugr.es

Abstract—The problem of the estimation of the information rate-distortion-perception function (RDPF), which is a relevant information-theoretic quantity in goal-oriented lossy compression and semantic information reconstruction, is investigated here. Specifically, we study the RDPF tradeoff for Gaussian sources subject to a mean-squared error (MSE) distortion and a perception measure that belongs to the family of α -divergences. Assuming a jointly Gaussian RDPF, which forms a convex optimization problem, we characterize an upper bound for which we find a parametric solution. We show that evaluating the optimal parameters of this parametric solution is equivalent to finding the roots of a reduced exponential polynomial of degree α . Additionally, we determine which disjoint sets contain each root, which enables us to evaluate them numerically using the well-known bisection method. Finally, we validate our analytical findings with numerical results and establish connections with existing results.

Index Terms—rate-distortion-perception, goal-oriented semantic communication, α divergence

I. INTRODUCTION

Rate-distortion-perception (RDP) theory, which has led to a surge of research recently, considers the problem of lossy compression under perceptual constraints on the reconstructed samples, generalizing the classical rate-distortion (RD) theory. Its mathematical representation is given by the rate-distortion-perception function (RDPF), whose properties are first studied in [1]–[3]. The rationale behind this approach is rooted in the observation that minimizing the distortion does not necessarily result in visually pleasing human perception results. Interestingly, the divergence constraints can be interpreted as semantic quality metrics, which quantify the relevance and significance of the reconstructed source from the observer’s perspective in goal-oriented semantic communication [4] (see also [5]).

The RDPF, like the conventional RDF, is devoid of an analytical solution for generic sources. Nonetheless, there exist some efforts to obtain closed-form expressions for specific sources under several commonly used metrics related to distortion and perception. Some of these efforts are reported in [2], [6], focusing on discrete binary sources under Hamming distance as a distortion measure and total-variation distance as a perception measure and also assuming continuous Gaussian sources under MSE distortion and either Wasserstein-2 distance, Hellinger distance, Geometric Jensen-Shannon divergence or (direct and reverse) Kullback-Leibler

(KL) divergence perception constraints. A generic method to compute the RDPF for continuous alphabet sources (not necessarily Gaussian), when the distribution of the source is constrained to be the same as the distribution of the output is studied in [7].

The key contribution of this paper is a closed-form solution to the RDPF of a Gaussian distributed source subject to MSE distortion and an α divergence perception constraint. Alpha divergences, which are first proposed by Chernoff [8] and thoroughly investigated by Amari [9], [10], are unique canonical divergences that reside at the intersection of the f -divergences and Bregman divergences in a manifold of positive measures. Since α -divergences are convex on both arguments, the RDPF forms a convex optimization problem. By constraining the reconstruction to be jointly Gaussian with the source, we define the jointly Gaussian RDPF (JG-RDPF) and show that it is an upper bound on the RDPF. Next, we provide a parametric solution to the JG-RDPF, for which we need to find the roots of a reduced exponential polynomial of degree α , derived from the perception constraint. We demonstrate that the reduced polynomial has a unique global stationary point and that there exist two unique roots that belong to disjoint sets for every value of α . The two roots can be approximated numerically using the bisection method. Subsequently, we provide a new parametric closed-form expression for the Pearson divergence ($\alpha = 2$), which is a specific instance of α divergence. We verify our theoretical findings with numerical evaluations.

Notation: Let \mathcal{X} be a Euclidean space (possibly finite-dimensional), and $\mathbb{B}(\mathcal{X})$ the Borel σ -algebra on \mathcal{X} . A random variable X defined on some probability space $(\Omega, \mathcal{F}, \mathcal{P})$ is a map $X : \Omega \mapsto \mathcal{X}$, where $(\mathcal{X}, \mathbb{B}(\mathcal{X}))$ is a measurable space. Given a continuous and twice differentiable function $f(x)$, the notation $f'(x)$ and $f''(x)$ denotes its first and second derivatives w.r.t. its argument.

II. PRELIMINARIES

We first state the formal definition and some preliminary results of the RDPF following [2]. Subsequently, we introduce the α -divergence measure used in the paper.

A. RDPF

We first provide the mathematical definition of the RDPF for general alphabets.

Definition 1: (RDPF) Let the information source X be a random variable on $(\mathcal{X}, \mathbb{B}(\mathcal{X}))$ with distribution $p_X \in \mathcal{P}(\mathcal{X})$. Let $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_0^+$ be a (measurable) distortion function satisfying $d(x, \hat{x}) = 0$ iff $x = \hat{x}$ and let $D : \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\hat{\mathcal{X}}) \rightarrow \mathbb{R}_0^+$ be a divergence function with $D(p_X, p_{\hat{X}}) = 0$ iff $p_X = p_{\hat{X}}$. The RDPF is defined as:

$$R(D, P) \triangleq \inf_{p_{\hat{X}|X}} I(X, \hat{X}) \quad (1)$$

$$\text{s.t. } \mathbb{E} [d(X, \hat{X})] \leq D, \quad (2)$$

$$D(p_X || p_{\hat{X}}) \leq P \quad (3)$$

where the infimum is over all conditional distributions $p_{\hat{X}|X} : \mathcal{X} \rightarrow \mathcal{P}(\hat{\mathcal{X}})$ and $I(X, \hat{X})$ denotes the mutual information [11] between the source X and the reconstructed source \hat{X} .

The following remark highlights the general properties of the RDPF.

Remark 1: (On Definition 1) The optimization problem (1) enjoys useful properties, under mild regularity conditions of the perception metric $D(\cdot || \cdot)$. In particular, [2, Theorem 1] shows that $R(D, P)$ is (i) monotonically non-increasing function in both arguments; (ii) convex if the divergence $D(\cdot || \cdot)$ is convex in its second argument.

B. Alpha Divergence

The α -divergence $D_\alpha(p||q)$ between two (normalized) probability distributions p and q , using the Amari notation [10], is defined as

$$D_{\alpha_A}(p||q) \triangleq \frac{4}{1 - \alpha_A^2} \left(1 - \int p(x)^{\frac{1-\alpha_A}{2}} q(x)^{\frac{1+\alpha_A}{2}} dx \right),$$

for $\alpha_A \in \mathbb{R} \setminus \{\pm 1\}$.¹ The case with $\alpha_A = \pm 1$ is derived from limit $\alpha_A \rightarrow \pm 1$. Interestingly, D_{α_A} and $D_{-\alpha_A}$ are dual in the sense that $D_{\alpha_A}(p||q) = D_{-\alpha_A}(q||p)$ (reference duality).

An alternate definition of α -divergence between two (normalized) probability distributions p and q (see e.g., [12]), which is used in the remainder of the paper, is given as follows:

$$D_\alpha(p||q) = \frac{1}{\alpha(\alpha - 1)} \left(\int_{-\infty}^{\infty} p(x)^\alpha q(x)^{1-\alpha} dx - 1 \right), \quad (4)$$

for $\alpha \in \mathbb{R} \setminus \{0, 1\}$. Since $\alpha = (1 - \alpha_A)/2$, the reference duality is expressed by $D_{1-\alpha}(p||q) = D_\alpha(q||p)$. We also have $D_\alpha(p||q) \geq 0$ with equality iff $p = q$, and that $D_\alpha(\cdot || \cdot)$ is not symmetric (i.e., $D_\alpha(p||q) \neq D_\alpha(q||p)$) except for $\alpha = 0.5$. Furthermore, $D_\alpha(p||q)$ is convex with respect to both arguments p and q .

1) *The Role of α :* The value of the α parameter may have a significant effect on the resulting approximate distribution q . For $\alpha \leq 0$, the α -divergence enforces q to have low density wherever p has low density (zero-forcing). On the other hand, when $\alpha \geq 1$, the divergence is inclusive, i.e., it enforces $q > 0$ wherever $p > 0$, hence avoiding zero probability density in regions of the space in which p has high density. When

¹The coefficient $4/(1 - \alpha_A^2)$ guarantees that any D_{α_A} gives the same Fisher information metric when p and q are infinitesimally close.

$\alpha \in (0, 1)$, the resulting q distribution is intermediate between the two extreme possibilities. In particular, when $\alpha \rightarrow 0$, we should expect that the approximate distribution q is more centered in the main mode of p (mean-seeking). By contrast, when $\alpha \rightarrow 1$, q is expected to cover the target distribution p more and capture more modes (mode-seeking).

2) *Special cases:* Several widely used divergences are actually α -divergences with specific values of α . For instance, when $\alpha \rightarrow 1$ we obtain the KL divergence, when $\alpha \rightarrow 0$ we obtain the reverse KL divergence, when $\alpha \rightarrow -1$ we obtain the inverse Pearson divergence, when $\alpha = 1/2$ we obtain the Hellinger distance, and when $\alpha = 2$ the Pearson divergence. For $\alpha = 0$, the α -connection [10] is the Levi-Civita connection under the Fisher-Rao metric. Note that the Rényi- α and the Tsallis- α divergences are closely related divergences (but not special cases).

3) *Gaussian Case:* Under the assumption that p and q belong to the set of Gaussian distributions, i.e., $p = \mathcal{N}(\mu, \sigma^2)$ and $q = \mathcal{N}(\nu, \rho^2)$, $D_\alpha(p||q)$ can be analytically characterized as

$$D_\alpha(p||q) = \frac{1}{\alpha(1 - \alpha)} (1 - H_\alpha(p, q))$$

$$H_\alpha(p, q) = \frac{\rho^\alpha \sigma^{1-\alpha}}{\sqrt{\alpha \rho^2 + (1 - \alpha) \sigma^2}} e^{-\frac{\alpha(1-\alpha)(\mu-\nu)^2}{2(\alpha \rho^2 + (1-\alpha)\sigma^2)}}. \quad (5)$$

However, (5) is guaranteed to be real only when $\alpha \in [0, 1]$. For $\alpha \notin [0, 1]$, additional conditions have to be imposed to guarantee that H_α is a real valued function, i.e., $\alpha \rho^2 + (1 - \alpha) \sigma^2 > 0$, which ensures the existence of the integral in (4).

For $\alpha > 1$, the condition becomes $\sigma^2 < \frac{\alpha}{\alpha-1} \rho^2$, indicating that the variance of p should be smaller than the variance of q multiplied by a constant factor. As $\alpha \rightarrow \infty$, this constant factor vanishes, and the constraint on the variances simplifies to $\sigma^2 \leq \rho^2$. On the other hand, as $\alpha \rightarrow 1^+$ the constraint disappears, i.e., $\sigma^2 < \infty$. For $\alpha < 0$, due to the duality of α divergence, we can apply the conclusions from the case $\alpha > 1$ by swapping the roles of σ and ρ . The constraint shifts into $\sigma^2 > \frac{\alpha}{\alpha-1} \rho^2$. When $\alpha \rightarrow -\infty$, it becomes $\sigma^2 \geq \rho^2$ and when $\alpha \rightarrow 0^-$, the constraint disappears, $\sigma^2 > 0$.

III. MAIN RESULTS

In this section, we present our main results. Consider the RDPF problem in Definition 1 assuming that the information source is Gaussian distributed, i.e., $X \sim \mathcal{N}(\mu, \sigma^2)$, the perception measure (3) belongs to the family of α -divergences, i.e., $D(\cdot || \cdot) = D_\alpha(\cdot || \cdot)$, and the distortion measure is the MSE metric, i.e., $d(X, \hat{X}) = (X - \hat{X})^2$. Following [6], we can characterize an upper bound to the RDPF for a Gaussian source X by constraining the reconstruction distribution \hat{X} to be jointly Gaussian with X as follows.

Definition 2: (JG-RDPF) Let $D \geq 0$, $P \geq 0$. Then, assuming a Gaussian random variable X is jointly Gaussian

with a reconstruction variable $\hat{X} \sim q = \mathcal{N}(\nu, \rho^2)$, we define the joint Gaussian RDPF (JG-RDPF) $R^G(D, P)$ as follows:

$$\begin{aligned} R^G(D, P) &= \min_{p_{\hat{X}|X}, \nu, \rho^2} I(X, \hat{X}) \\ \text{s.t. } &\mathbb{E} \left[(X - \hat{X})^2 \right] \leq D, \\ &D_\alpha(p||q) \leq P. \end{aligned} \quad (6)$$

We remark that the definition $R^G(D, P)$ induces the following upper bound to $R(D, P)$.

Lemma 1: Let $D \geq 0$, $P \geq 0$. Then, $R(D, P) \leq R^G(D, P)$.

Proof: The proof follows by observing that the constraint set of (6) is a proper subset of the constraint set of (1). ■

We stress the following technical remark.

Remark 2: The assumption of \hat{X} being jointly Gaussian with X leads to an upper bound to the optimal solution of the problem (6). This is because the distribution minimizing the perception constraint $D_\alpha(\cdot||\cdot)$ cannot be guaranteed to be Gaussian. However, in the case of reverse KL divergence, i.e., ($\alpha \rightarrow 0$), where it is proven that the minimizing distribution is itself Gaussian [13], the above upper bound is exact.

We note that, without loss of optimality, in (6) we can assume that $\mu = \nu = 0$ since this does not affect the mutual information $I(X, \hat{X})$ while being optimal in terms of MSE and D_α metrics, as shown in the following lemma.

Lemma 2: Let $\Delta = \mu - \nu$. Then, $D_\alpha(p||q)$ has a global minimum at $\Delta = 0$.

Proof: For $\alpha \in (-\infty, 0) \cup (1, +\infty)$, we see that $D_\alpha(p||q) \approx s(\Delta) = e^{\Delta^2} - c$, where c denotes constant terms. The first and second derivative of $s(\Delta)$ are, respectively,

$$s'(\Delta) = 2\Delta e^{\Delta^2}, \quad s''(\Delta) = 2e^{\Delta^2} + 4\Delta^2 e^{\Delta^2} > 0.$$

We conclude that $s(\Delta)$ is strongly convex and therefore minimized by its stationary point $\Delta_0 = 0$. Hence, also $D_\alpha(p||q)$ has a global minimum at $\Delta_0 = 0$.

For $\alpha \in (0, 1)$, $D_\alpha(p, q) \approx t(\Delta) = c - e^{-\Delta^2}$. The first and second derivative of $t(\Delta)$ are, respectively,

$$t'(\Delta) = 2\Delta e^{-\Delta^2}, \quad t''(\Delta) = 2e^{-\Delta^2}(1 - 2\Delta^2),$$

from which we characterize stationary point $\Delta_0 = 0$. For $\Delta < 0$, $t'(\Delta)$ is negative which means that $t(\Delta)$ is decreasing where for $\Delta > 0$, $t'(\Delta)$ is positive meaning that $t(\Delta)$ is increasing. For $\Delta \in (-\infty, -\frac{1}{\sqrt{2}}) \cup (\frac{1}{\sqrt{2}}, +\infty)$, $t''(\Delta)$ is negative where for $\Delta \in (-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ is positive. Therefore, the function $t(\Delta)$ and thus $D_\alpha(p||q)$ has a global minimum at $\Delta_0 = 0$. This concludes the proof. ■

The solution of (6) can be greatly facilitated by an in-depth analysis of the perception constraint $D_\alpha(p||q) \leq P$, focusing in particular on the case where it holds with equality. The following lemma characterizes essential aspects of this case.

Lemma 3: Assume $P \in [0, P_{\alpha, \max}]$ and let $p = \mathcal{N}(0, \sigma^2)$ and $q = \mathcal{N}(0, \rho^2)$. Then the constraint $D_\alpha(p||q) = P$ can be expressed as $f(x) = 0$, where

$$f(x) = x^\alpha - \alpha Cx - (1 - \alpha)C \quad (7)$$

with $x = \frac{\rho^2}{\sigma^2}$ and $C = (1 - \alpha(1 - \alpha)P)^2$. Furthermore, for all $\alpha \in (-\infty, 0) \cup (0, 1) \cup (1, +\infty)$, (7) has two roots, r_0 and r_1 , such that $r_0 \in [0, x_0]$ and $r_1 \in [x_0, y_0]$, where x_0 is the unique stationary point of (7) and $y_0 = (x_0 + \epsilon) - \frac{f(x_0 + \epsilon)}{f'(x_0 + \epsilon)}$ for $\epsilon > 0$.

Proof: The equivalence between $D_\alpha(p||q) = P$ and (7) easily follows from algebraic manipulations. We examine the properties of equation (7) through its derivatives, i.e.,

$$f'(x) = \alpha(x^{\alpha-1} - C), \quad f''(x) = \alpha(1 - \alpha)x^{\alpha-2}.$$

Considering $f'(x)$, we can characterize the stationary point x_0 of $f(x)$, i.e., $f'(x) = 0 \stackrel{\alpha \neq 0, 1}{\Leftrightarrow} x_0 = C^{\frac{1}{\alpha-1}}$. For $\alpha \in (-\infty, 0) \cup (1, +\infty)$, $f''(x) \geq 0$, implying that $f'(x)$ is increasing and $f(x)$ is concave up. Therefore, x_0 will be the global minimum of $f(x)$. For $\alpha \in (0, 1)$, it is $f''(x) \leq 0$ which means $f'(x)$ is decreasing and $f(x)$ is concave down. Hence, x_0 will be the global maximum of $f(x)$. We can now infer about the number of roots of $f(x)$ by determining the sign of $f(x_0)$, i.e.,

$$f(x_0) = (1 - \alpha)C \left(C^{\frac{1}{\alpha-1}} - 1 \right).$$

By definition, the condition $C \geq 0$ holds. We notice that for $C = 0$, $x_0 = 0$ is the unique root of the function. To evaluate the sign of $f(x_0)$, we can verify that the following holds

$$\begin{aligned} \alpha < 0 &\Rightarrow C^{\frac{1}{\alpha-1}} - 1 \leq 0 \Rightarrow f(x_0) \leq 0, \\ 0 < \alpha < 1 &\Rightarrow C^{\frac{1}{\alpha-1}} - 1 \geq 0 \Rightarrow f(x_0) \geq 0, \\ \alpha > 1 &\Rightarrow C^{\frac{1}{\alpha-1}} - 1 \geq 0 \Rightarrow f(x_0) \leq 0. \end{aligned}$$

This implies that when $f(x_0) \leq 0$ (resp. $f(x_0) \geq 0$) we have a positive (resp. negative) second derivative. Since $f'(x)$ is monotonic, $f(0) = -(1 - \alpha)C$ and

$$\lim_{x \rightarrow +\infty} f(x) = \begin{cases} +\infty & \text{if } \alpha \in (-\infty, 0) \cup (1, +\infty) \\ -\infty & \text{if } \alpha \in (0, 1) \end{cases}$$

we can deduce that $f(x)$ necessarily has two roots in \mathbb{R} .

We now identify the sets where the roots r_0 and r_1 reside. Starting with r_0 , we can notice that $f(x)$ is monotone on the set $[0, x_0]$ while changing the sign in its extremes. Hence, the root r_0 must belong to this set. In the case of r_1 , we notice that, for $\epsilon > 0$, the following inequalities hold

$$\begin{cases} f(y) \geq f_T(y) & \text{if } \alpha \in (-\infty, 0) \cup (1, +\infty) \\ f(y) \leq f_T(y) & \text{if } \alpha \in (0, 1) \end{cases}$$

$$f_T(y) = f(x_0 + \epsilon) + f'(x_0 + \epsilon)(y - (x_0 + \epsilon))$$

due to the convexity (resp. concavity) of f . This allows us to characterize the set where r_1 resides by taking as inferior x_0 and as superior the zero of the linear function $f_T(y)$, i.e., $y_0 = (x_0 + \epsilon) - f(x_0 + \epsilon)(f'(x_0 + \epsilon))^{-1}$. This concludes the proof. ■

We stress the following point regarding the evaluation of the roots r_0 and r_1 .

Remark 3: Lemma 3 proves that, in their respective interval, r_0 and r_1 are the unique roots of (7). This property facilitates the numerical estimation of the problem by allowing the

bisection method to be applied directly [14, Chapter 2.1] in each interval.

Equipped with the previous lemma, we are ready to develop a parametric solution for the JG-RDPF in (6).

Theorem 1: (Parametric JG-RDPF solution) Let X be a scalar Gaussian source $X \sim \mathcal{N}(0, \sigma^2)$. Then, the JG-RDPF $R^G(D, P)$ under squared error distortion and α -divergence perception is achieved by a jointly Gaussian reconstruction $\hat{X} \sim \mathcal{N}(0, \rho^2)$ and is given by

$$R^G(D, P) = \begin{cases} \max\left\{\frac{1}{2} \log \frac{\sigma^2}{D}, 0\right\} & \text{if } (D, P) \in \mathcal{S} \\ \frac{1}{2} \log \frac{2\rho^2\sigma^2}{\rho^2\sigma^2 - \left(\frac{\sigma^2 + \rho^2 - D}{2}\right)^2} & \text{if } (D, P) \notin \mathcal{S} \end{cases} \quad (8)$$

where

$$\begin{aligned} \rho^2 &= \begin{cases} \sigma^2 - D & \text{if } (D, P) \in \mathcal{S} \\ \min\{r_0, r_1\} & \text{if } (D, P) \notin \mathcal{S}. \end{cases} \\ \mathcal{S} &= \left\{ (D, P) \in \mathbb{R}_+^2 : P > g(D, \sigma) \wedge \right. \\ &\quad \left. (\alpha - 1) \left(\left| 1 - \frac{D}{\sigma^2} \right| - \left(1 - \frac{1}{\alpha} \right) \right) > 0 \right\} \\ g(D, \sigma) &= \frac{1}{\alpha(1-\alpha)} \left(1 - \frac{\sigma^{1-\alpha} |\sigma^2 - D|^{\alpha/2}}{\sqrt{\alpha |\sigma^2 - D| + (1-\alpha)\sigma^2}} \right), \end{aligned}$$

with r_0 and r_1 being the roots of (7).

Proof: The objective function and constraints of (6) can be rewritten in terms of parameters of the Gaussian source $X \sim \mathcal{N}(0, \sigma^2)$ and reconstruction $\hat{X} \sim \mathcal{N}(0, \rho^2)$ as follows

$$I(X, \hat{X}) = \frac{1}{2} \log \frac{\rho^2 \sigma^2}{\rho^2 \sigma^2 - \theta^2}, \quad (9)$$

$$\mathbb{E} \left[(X - \hat{X})^2 \right] = \sigma^2 + \rho^2 - 2\theta \leq D, \quad (10)$$

$$D_\alpha(p||q) = \frac{1}{\alpha(1-\alpha)} \left(1 - \frac{\rho^\alpha \sigma^{1-\alpha}}{\sqrt{\alpha \rho^2 + (1-\alpha)\sigma^2}} \right) \leq P \quad (11)$$

where $\theta = \mathbb{E}[X\hat{X}]$ is the covariance between X and \hat{X} .

We derive the closed-form solution of $R^G(D, P)$ by considering the different cases when only the distortion constraint is active (**Case I**), when only the perception constraint is active (**Case II**), or when both are active (**Case III**).

Case I: When the perception constraint is not active (i.e., $P = \infty$), the solution to (6) reduces to the rate-distortion function of a Gaussian source [11, Theorem 10.3.2]

$$R^G(D, \infty) = \begin{cases} \frac{1}{2} \log \frac{\sigma^2}{D} & 0 \leq D \leq \sigma^2 \\ 0 & D > \sigma^2 \end{cases},$$

which is attained by some $p_{\hat{X}|X}$ with marginal distribution $\hat{X} \sim \mathcal{N}(0, \sigma^2 - D)$. Furthermore, since the distortion constraint is active and holds with equality while the perception constraint is inactive, $\rho = \sqrt{\sigma^2 - D} < \sigma$ and (11) must hold with strict inequality. However, we note that for $\alpha > 1$ the argument of the root in (11) may be negative, hence the

constraint can become ill-posed. To avoid this case, we enforce the following additional constraint

$$\left| 1 - \frac{D}{\sigma^2} \right| > \left(1 - \frac{1}{\alpha} \right).$$

Case II: When only the perception constraint is active, we can characterize the variance of the reconstruction ρ^2 by solving the generalized equation (7), following the results of Lemma 3. Furthermore, the inactive distortion constraint implies that (10) holds with strict inequality for all $\theta \in [0, \sigma\rho]$. We can therefore set $\theta = 0$, which in turns implies that $R^G(\infty, P) = 0$.

Case III: When both constraints hold with equality, the resulting system of equations can be used to characterize both the variance ρ^2 and the covariance θ . Using Lemma 3, we can find the two roots of (7), i.e., ρ_0^2 and ρ_1^2 . Despite both roots being optimal solutions for (7), we select $\rho^2 = \min\{\rho_0^2, \rho_1^2\}$ under the argument that a lower variance of the reconstruction is optimal in terms of MSE distortion. From the equality of the distortion constraint, we can instead derive $\theta = \frac{1}{2}(\sigma^2 + \rho^2 - D)$, which, once substituted in (9), recovers the closed-form expression for $R^G(D, P)$ in (8). This concludes the proof. ■

JG-RDPF Special Cases: We shall notice that for specific instances of D_α we can directly recover previously characterized analytical JG-RDPF [6], i.e., Hellinger distance ($\alpha = \frac{1}{2}$), direct and reverse Kullback-Leibler divergence ($\alpha \rightarrow 1$ and $\alpha \rightarrow 0$, respectively). In the remainder of this section, we derive a new characterization for the Pearson divergence ($\alpha = 2$). We should mention that in each of these cases, the roots of the generalized polynomial (7) can be easily found analytically, allowing for the closed-form characterization of $\rho(\sigma, P)$, based on which we recover the closed-form expression of $R^G(D, P)$ once substituted in (8).

Pearson Divergence ($\alpha = 2$):

$$(7) \stackrel{\alpha=2}{\Leftrightarrow} \rho^4 - 2(1+2P)^2 \rho^2 \sigma^2 + (1+2P)^2 \sigma^4 = 0$$

$$\rho^2 = \begin{cases} \sigma^2 (1+2P) \left(1 + 2P + 2\sqrt{P+P^2} \right) & \\ \quad \text{if } \rho^2 - (1+2P)^2 \sigma^2 > 0, & \\ \sigma^2 (1+2P) \left(1 + 2P - 2\sqrt{P+P^2} \right) & \\ \quad \text{if } \rho^2 - (1+2P)^2 \sigma^2 \leq 0. & \end{cases}$$

IV. NUMERICAL RESULTS

In this section, we provide plots that visually illustrate the main results and provide insights into the behavior and functionality of the RDPF. In all the presented examples, we consider a Gaussian source $X \sim \mathcal{N}(0, 1)$.

a) *Polynomial representation of perception constraint:*

Fig. 1 plots the polynomial (7) for different values of α and fixed perception constraint $P = 0.2$, taking into account the constraint of the variances $\alpha\rho^2 + (1-\alpha)\sigma^2 > 0$. We see that for every value of α , the polynomial has two roots in the disjoint sets defined based on the global extrema.




Fig. 1. Polynomial of (7) for $\alpha = -1.2$, $\alpha = 0.1$, $\alpha = 0.5$ and $\alpha = 2$, with the perception constraint fixed to $P = 0.2$.




Fig. 2. Rate-distortion curves for $\alpha = -5$, $\alpha = -0.5$, $\alpha = 1.5$ and $\alpha = 3$. The first figure represents $P = 0$ whereas the second $P = 0.7$.

b) Comparison between α -divergences: Fig. 2 plots the RD curve for varying α values. As expected, we see that for $P = 0$ (perfect perceptual quality), all curves overlap. On the other hand, when $P = 0.7$, the curves diverge from each other, since different α values lead to different perception measures. We notice that as α increases, the curves start to differ more from each other indicating that the perceptual constraint becomes more stringent, leading to higher distortion for the same rate compared to lower α values. Fig. 3 evinces the matching between our analytical closed-form solutions and our numerical results for the special cases of $\alpha = 0.5$, $\alpha \rightarrow 1$ and $\alpha = 2$. In the first two cases, we confirm our results with [6], theorems 4 and 2, which provide a closed-form solution for JG-RDPF with Hellinger and KL divergence as perception constraints, respectively.

ACKNOWLEDGMENT

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




Fig. 3. Rate-distortion curves for the special cases of $\alpha = 0.5$, $\alpha \rightarrow 1$ and $\alpha = 2$. The perception constraint is fixed $P = 0.5$.

G. Serra is supported. The work of M. Sourla was supported by the ERASMUS+ programme and the work of P. A. Stavrou is supported by the SNS JU project 6G-GOALS [15] under the EU’s Horizon programme Grant Agreement No. 101139232.

REFERENCES

- [1] Y. Blau and T. Michaeli, “The Perception-Distortion tradeoff,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6228–6237.
- [2] —, “Rethinking lossy compression: The rate-distortion-perception tradeoff,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 675–685.
- [3] R. Matsumoto, “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] M. Kountouris and N. Pappas, “Semantics-empowered communication for networked intelligent systems,” *IEEE Commun. Mag.*, vol. 59, no. 6, pp. 96–102, 2021.
- [5] P. A. Stavrou and M. Kountouris, “The role of fidelity in goal-oriented semantic communication: A rate distortion approach,” *IEEE Transactions on Communications*, vol. 71, no. 7, pp. 3918–3931, 2023.
- [6] 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*, vol. 5, pp. 314–330, 2024.
- [7] —, “Copula-based estimation of continuous sources for a class of constrained rate-distortion-functions,” in *IEEE International Symposium on Information Theory (ISIT)*, 2024, (to appear).
- [8] H. Chernoff, “A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations,” *Annals of Mathematical Statistics*, vol. 23, no. 4, pp. 493–507, 1952.
- [9] S. Amari, “Differential geometry of curved exponential families—curvatures and information loss,” *Annals of Mathematical Statistics*, vol. 10, no. 2, p. 357–385, 1982.
- [10] —, *Differential geometric methods in statistics*. New York: Springer Verlag, 1985.
- [11] T. Cover and J. Thomas, *Elements of Information Theory*. John Wiley & Sons, 2006.
- [12] N. Cressie and T. Read, *Goodness-of-Fit Statistics for Discrete Multivariate Data*. Springer: New York, NY, USA, 1988.
- [13] S. Fang and Q. Zhu, “Independent Gaussian distributions minimize the Kullback-Leibler (KL) divergence from independent Gaussian distributions,” *arXiv:2011.02560*, 2020.
- [14] R. L. Burden, J. D. Faires, and A. M. Burden, *Numerical analysis*. Cengage learning, 2015.
- [15] E. C. Strinati et al., “Goal-oriented and semantic communication in 6G AI-native networks: the 6G-GOALS approach,” in *European Conference on Networks and Communications & 6G Summit*, June 2024.