Computation of Rate-Distortion-Perception Function under f-Divergence Perception Constraints

Giuseppe Serra, Photios A. Stavrou, Marios Kountouris
Communication Systems Department, EURECOM, Sophia-Antipolis, France
{giuseppe.serra,fotios.stavrou,marios.kountouris}@eurecom.fr

***Abstract*—In this paper, we study the computation of the rate-distortion-perception function (RDPF) for discrete memoryless sources subject to a single-letter average distortion constraint and a perception constraint that belongs to the family of f-divergences. For that, we leverage the fact that RDPF, assuming mild regularity conditions on the perception constraint, forms a convex programming problem. We first develop parametric characterizations of the optimal solution and utilize them in an alternating minimization approach for which we prove convergence guarantees. The resulting structure of the iterations of the alternating minimization approach renders the implementation of a generalized Blahut-Arimoto (BA) type of algorithm infeasible. To overcome this difficulty, we propose a relaxed formulation of the structure of the iterations in the alternating minimization approach, which allows for the implementation of an approximate iterative scheme. This approximation is shown, via the derivation of necessary and sufficient conditions, to guarantee convergence to a globally optimal solution. We also provide sufficient conditions on the distortion and the perception constraints which guarantee that our algorithm converges exponentially fast. We corroborate our theoretical results with numerical simulations, and we draw connections with existing results.**

I. Introduction

The theoretical framework of rate distortion theory subject to a single-letter average distortion constraint stems from the seminal work of Shannon in [1], [2]. Therein, Shannon described, for the first time, the fundamental trade-offs between the desired bit rate used for a compressed representation of the source messages and the associated achievable distortion criterion attained between the source message and its reconstructed representation. The mathematical representation in rate distortion theory is manifested by the rate-distortion function (RDF). Rate distortion theory has been the foundational basis to study and develop lossy compression algorithms for various multimedia applications.

In recent years, it has been shown through multiple works, spanning from machine learning and computer vision to multimedia applications [3]–[8], that focusing exclusively on distortion minimization does not necessarily imply a good perceptual quality of the reconstructed signal, where perceptual quality refers to the property of a sample to appear visually pleasing from a human perspective. Motivated by the need for a general characterization of RDF to enable encompassing the perceptual quality of the sample, Blau and Michaeli in [9] introduced a generalization of the single-letter RDF, coined rate-distortion-perception function (RDPF). The specific information theoretic characterization complements the classical single-letter distortion constraint between a source message and its reconstruction, which is inherent in RDF formulation, with a divergence constraint between the induced distributions of the source and its estimated value. The divergence constraint in RDPF is precisely used as a proxy of the human perception, measuring the degree of satisfaction in the consumption of the data from a human perspective. The divergence constraint can also be viewed as a semantic quality metric, which measures the degree of relevance of the reconstructed source from the perspective of the observer, a point first hinted at in [10]. Another relevant yet different setup is the recently introduced robust source coding framework, see, e.g., [11] (and references therein), in which in place of the perception quality, there exists an additional distortion criterion.

Since its conception, RDPF has received substantial interest from the information theory community. Theis and Wagner in [12] proved that the RDPF can be achieved via stochastic, variable-length codes by making use of a strong functional representation lemma [13], whereas Chen *et al.* in [14] studied the role of stochasticity in the encoder/decoder structure and proved the achievability of the RDPF by deterministic codes (except certain extreme cases) in the asymptotic regime. Wagner in [15] considered the case of perfect perceptual quality in the definition of RDPF and devised a coding theorem that allows for a specified amount of common randomness between the encoder and decoder.

It should be emphasized that similarly to the classical RDF, there are no known closed-form representations of the RDPF for general alphabet sources. In fact, there are only a few examples in the literature where analytical expressions are provided for the RDPF, such as the case of binary sources subject to Hamming distortion and total variation distance [9], [16] or the case of a Gaussian source subject to a mean-squared error distortion and a 2-Wasserstein distance [4].

In this paper, we focus on the computation of the RDPF for discrete memoryless sources subject to a single-letter average distortion constraint and a perception constraint that belongs to the class of f-divergences[1]. In particular, we leverage the fact that under mild regularity conditions on the perception constraint (i.e., convexity on the second argument) the RDPF forms a convex program. This enables us to derive a parametric characterization of the optimal solution of the RDPF (Lemma

[1]For a mathematical background on f-divergences and their properties we refer the reader to [17].

1), which is subsequently utilized to construct an alternating minimization procedure[2] for which we also establish convergence guarantees (Theorem 2). The resulting structure of the iterations in Theorem 2 prohibits the implementation of a generic Blahut-Arimoto (BA) algorithm similar to what is already known for the classical rate distortion theory for i. i. d sources and single-letter distortions [20]. To overcome this technical difficulty, we introduce a new relaxed formulation of the structure of the iterations in Theorem 2, which results in a variant of an approximate BA algorithm (Theorem 3). Additionally, in Theorem 3, we derive necessary and sufficient conditions to ensure that our approximation algorithm converges to a globally optimal solution. In light of our result in Theorem 3, we also provide sufficient conditions on the structure of the distortion and the perception constraints based on which our algorithm converges exponentially fast (Theorems 6, 7). We corroborate our theoretical findings with numerical simulations and draw comparisons with existing results in the literature (Section V).

Notation: We denote with $p_x(x)$ the probability distribution over the source alphabet $\mathcal{X}$ evaluated on the symbol x, and by $q_{\hat{x}}(\hat{x})$ any arbitrary marginal probability on the output alphabet $\hat{\mathcal{X}}$ evaluated on the symbol $\hat{x}$. $Q_{\hat{x}|x}(\hat{x}|x)$ denotes the entry $(\hat{x},x)$ of the transition matrix $Q_{\hat{x}|x}$ while $Q_{\hat{x}}(\hat{x})$ denotes the marginal on $\hat{\mathcal{X}}$ induced by p_x and $Q_{\hat{x}|x}$, evaluated on the symbol $\hat{x}$. We denote by $\mathbb{E}[\cdot]$ the expectation operator, and by $\mathbb{E}_q[\cdot]$ the probability distribution q on which the expectation operator is applied. We denote by $p_1[p_2]$ the explicit form of the functional dependence of some probability distribution p_1 functionally dependent on another probability distribution p_2. We denote by $D(\cdot||\cdot)$ any divergence measure and by $D_f(\cdot||\cdot)$, $D_{KL}(\cdot||\cdot)$, $D_{\chi^2}(\cdot||\cdot)$, $TV(\cdot||\cdot)$, the class of f-divergence, the Kullback–Leibler divergence, the chi-squared divergence and the total variation distance, respectively. We denote by $f\in C^0$, the continuous function defining the f-divergence and ∂f denotes its sub-gradient [21, Definition 8.3]. For a continuous and twice differentiable function $f\in C^2$, we denote by $f''(\cdot)$ the second derivative with respect to its argument.

II. Preliminaries on the RDPF

In this section, we consider finite alphabets sources and deterministic encoders/decoders pairs and define the minimum achievable rates subject to average per-letter distortion and average per-letter perception constraints. Our analysis herein stems from recent results in [14].

We assume that we are given an i. i. d sequence of n-length random variables X^n that induce the probability distribution p_x. The source sequence is received by an encoder (e) that generates the index $e(X^n)\in\mathcal{Z}$, $\mathcal{Z}=\{1,2,\ldots,2^{nR}\}$, whereas at the decoder (g), the system reconstructs an estimate of $\hat{X}^n$. Formally, the encoder and decoder are deterministic mappings with $e:\mathcal{X}^n\rightarrow\mathcal{Z}$ and $g:\mathcal{Z}\rightarrow\hat{\mathcal{X}}^n$, respectively.

We let $d:\mathcal{X}\times\hat{\mathcal{X}}\rightarrow\mathcal{R}_0^+$ to denote a single-letter distortion function and $D:\mathcal{P}\times\mathcal{Q}\rightarrow\mathcal{R}_0^+$ to denote a divergence measure. Moreover, we define the fidelity criterion Δ as the expected per-symbol distortion and the fidelity criterion Φ as the expected per-symbol divergence[3] as follows:

$$\Delta\triangleq\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^{n}d(x_i,\hat{x}_i)\right]\qquad\Phi\triangleq\frac{1}{n}\sum_{i=1}^{n}D(p_x||q_{\hat{x}_i}).$$

We are now ready to introduce the definition of achievability and that of infimum of all achievable rates.

Definition 1. *(Achievability) Given a distortion level $D>0$ and a perception constraint $P>0$, a rate R is said to be (D,P)-achievable if for an arbitrary $\epsilon>0$, there exists, for large enough n, a deterministic lossy source code (n,M,Δ,Φ) with $M\leq 2^{n(R+\epsilon)}$ such that $\Delta\leq D+\epsilon$ and $\Phi\leq P+\epsilon$. Then, we define*

$$R_{nr}(D,P)\equiv\inf\{R:(R,D,P)\text{ is achievable}\}.\quad(1)$$

Next, we give the definition of the information theoretic characterization of the RDPF (see [9]) assuming that $D>0$ and $P>0$.

Definition 2. *(RDPF) For a given finite alphabet source distribution p_x, a single-letter distortion $d(\cdot,\cdot)$ and a divergence $D(\cdot||\cdot)$, the RDPF is characterized as follows:*

$$\begin{aligned}R(D,P)=&\min_{Q_{\hat{x}|x}}I(X,\hat{X})\\&\text{s.t.}\quad\mathbb{E}[d(x,\hat{x})]\leq D,\quad D(p_x||Q_{\hat{x}})\leq P\end{aligned}\quad(2)$$

where $D\in[D_{\min},D_{\max}]\subset(0,\infty)$, $P\in[P_{\min},P_{\max}]\subset(0,\infty)$ and

$$I(X,\hat{X})=D_{KL}(p_xQ_{\hat{x}|x},p_xQ_{\hat{x}})\triangleq I(p_x,Q_{\hat{x}|x})\quad(3)$$

where $I(p_x,Q_{\hat{x}|x})$ highlights the dependency on $\{p_x,Q_{\hat{x}|x}\}$.

We stress the following technical remarks on Definition 2.

Remark 1. *(On Definition 2) Following [9], it can be shown that* (2) *has some useful properties, under mild regularity conditions. In particular, [9, 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.*

In the sequel, we consider that in (2) *the perception constraint is an f-divergence*, i.e., $D(\cdot||\cdot)=D_f(\cdot||\cdot)$, which is known to be convex in both arguments [22, Lemma 4.1]. Hence in light of the discussion in Remark 1, $R(D,P)$ forms a convex programming problem.

We conclude this section by providing a theorem that connects $R_{nr}(D,P)$ with $R(D,P)$ for finite alphabets, as long as $D>0$ and $P>0$.

[2]For details on the alternating minimization procedure, see e.g., [18], [19].

[3]Following for instance [14, Remark 3], one can choose the perception fidelity criterion to be $D(p_x||q_{\hat{x}_i})$, $i=1,\ldots,n$, or $D(p_x||\sum_{i=1}^{n}q_{\hat{x}_i})$. Both types of fidelity criteria result in the same operational quantity for finite alphabets in the asymptotic regime.

Theorem 1. *For $|\mathcal{X}| < \infty$, $D > 0$, $P > 0$, we obtain*

$$R_{nr}(D,P) = R(D,P). \tag{4}$$

Proof: This is a consequence of [14, Assumption 2] that results into [14, Theorem 2]. ∎

It should be noted that in the previous analysis, we excluded the extreme case where $P = 0$. This is because, in that scenario, deterministic encoders and decoders do not achieve $R(D,0)$ and one instead requires common [12, Theorem 3] or private [14, Theorems 4] randomness to achieve it.

III. Main Results

In this section, we present our main results. We start by reformulating (2) into a double minimization problem.

Lemma 1. *(Double minimization) Let $s_1 \geq 0$, $s_2 \geq 0$, and define $s = (s_1, s_2)$. Moreover, let $D > 0$, $P > 0$ and let $D(\cdot||\cdot) = D_f(\cdot||\cdot)$. Then* (2) *can be expressed as a double minimum as follows:*

$$\begin{aligned} R(D,P) =& - s_1 D - s_2 P + \min_{Q_{\hat{x}|x}} \min_{q_{\hat{x}}} \Big[s_1 \mathbb{E}[d(x,\hat{x})] \\ &+ s_2 D_f(p_x||Q_{\hat{x}}) + D_{KL}(p_x Q_{\hat{x}|x}||p_x q_{\hat{x}}) \Big] \end{aligned} \tag{5}$$

where $D = \mathbb{E}_{Q^*_{\hat{x}|x}}[d(x,\hat{x})]$, $P = D_f(p_x||Q^*_{\hat{x}})$, with $Q^*_{\hat{x}|x}$ achieving the minimum. Furthermore, for fixed $Q_{\hat{x}|x}(\hat{x}|x)$, the right side of (5) *is minimized by*

$$q_{\hat{x}}(\hat{x}) = \sum_{x\in\mathcal{X}} p_x(x) Q_{\hat{x}|x}(\hat{x}|x). \tag{6}$$

For fixed $q_{\hat{x}}$, the right side of (5) *is minimized by*

$$Q_{\hat{x}|x}(\hat{x}|x) = q_{\hat{x}}(\hat{x}) \cdot \frac{A(x,\hat{x},s)}{\sum_{i\in\mathcal{X}} q_{\hat{x}}(i) A(x,i,s)} \tag{7}$$

where

$$\begin{aligned} A(x,\hat{x},s) &= \exp\{-s_1 d(x,\hat{x}) - s_2 g(p_x, Q_{\hat{x}}, \hat{x})\} \\ g(p_x, Q_{\hat{x}}, \hat{x}) &= f\left(\frac{p_x(\hat{x})}{Q_{\hat{x}}(\hat{x})}\right) - \frac{p_x(\hat{x})}{Q_{\hat{x}}(\hat{x})} \partial f\left(\frac{p_x(\hat{x})}{Q_{\hat{x}}(\hat{x})}\right). \end{aligned} \tag{8}$$

We note that Lemma 1 differs from [23, Theorem 6.3.3] in (7). Specifically, the presence of the additional perception constraint $D_f(\cdot||\cdot)$ in (2) changes the functional properties of the parametric solution (7), as it requires an additional exponential term, i.e., $s_2 g(p_x, Q_{\hat{x}}, \hat{x})$, with $g(\cdot)$ depending on the induced marginal $Q_{\hat{x}}$.

The next corollary is a consequence of Lemma 1.

Corollary 1. *$R(D,P)$ in* (5) *can be reformulated as follows:*

$$\begin{aligned} R(D_s, P_s) =& - s_1 D_s - s_2 P_s \\ &+ \min_{q_{\hat{x}}} s_2 \sum_{\hat{x}\in\hat{\mathcal{X}}} p_x(\hat{x}) \partial f\left(\frac{p_x(\hat{x})}{Q_{\hat{x}}[q_{\hat{x}}](\hat{x})}\right) \\ &- \sum_{x\in\mathcal{X}} p_x(x) \log\left(\sum_{\hat{x}\in\hat{\mathcal{X}}} q_{\hat{x}}(\hat{x}) A(x,\hat{x},s)\right) \end{aligned} \tag{9}$$

where $q^*_{\hat{x}}$ achieves the minimum and

$$\begin{aligned} D_s &= \sum_{(x,\hat{x})\in\mathcal{X}\times\hat{\mathcal{X}}} p_x(x) \frac{q^*_{\hat{x}}(\hat{x}) A(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} q^*_{\hat{x}}(i) A(x,i,s)} d(x,\hat{x}) \\ P_s &= D_f(p_x||Q_{\hat{x}}[q^*_{\hat{x}}]). \end{aligned}$$

Next, we proceed to construct an alternating minimization procedure and show its convergence to a point of $R(D,P)$.

Theorem 2. *(Alternating minimization) Let $s_1 \geq 0$, $s_2 \geq 0$ be given, with $s = (s_1, s_2)$, and let $A[q_{\hat{x}}]$ be such that*

$$\begin{aligned} A[q_{\hat{x}}](x,\hat{x},s) &= \exp\{-s_1 d(x,\hat{x}) - s_2 g(p_x, Q_{\hat{x}}[q_{\hat{x}}], \hat{x})\} \\ g(p_x, Q_{\hat{x}}[q_{\hat{x}}], \hat{x}) &= f\left(\frac{p_x(\hat{x})}{Q_{\hat{x}}[q_{\hat{x}}](\hat{x})}\right) - \frac{p_x(\hat{x})}{Q_{\hat{x}}[q_{\hat{x}}](\hat{x})} \partial f\left(\frac{p_x(\hat{x})}{Q_{\hat{x}}[q_{\hat{x}}](\hat{x})}\right). \end{aligned}$$

Let $q^{(0)}_{\hat{x}}$ denote any probability vector with nonzero components and let $q^{(n+1)}_{\hat{x}} \equiv Q_{\hat{x}}[q^{(n)}_{\hat{x}}]$ and $Q^{(n+1)}_{\hat{x}|x} \equiv Q_{\hat{x}|x}[q^{(n)}_{\hat{x}}]$ be defined as functions of the previous iteration $q^{(n)}_{\hat{x}}$ as follows:

$$\begin{aligned} Q^{(n+1)}_{\hat{x}|x}(\hat{x}|x) &= q^{(n)}_{\hat{x}}(\hat{x}) \frac{A^{(n)}(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} q^{(n)}_{\hat{x}}(i) A^{(n)}(x,i,s)} \\ q^{(n+1)}_{\hat{x}}(\hat{x}) &= q^{(n)}_{\hat{x}}(\hat{x}) \sum_{x\in\mathcal{X}} \frac{p_x(x) A^{(n)}(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} q^{(n)}_{\hat{x}}(i) A^{(n)}(x,i,s)} \end{aligned} \tag{10}$$

where $A^{(n)}(x,\hat{x},s) = A[q^{(n)}_{\hat{x}}](x,\hat{x},s)$. Then, as $n \to \infty$, we obtain

$$D(Q^{(n)}_{\hat{x}|x}) \to D_s,\ P(Q^{(n)}_{\hat{x}|x}) \to P_s,\ I(p_x, Q^{(n)}_{\hat{x}|x}) \to R(D_s, P_s).$$

Despite being optimal, the alternating minimization scheme of Theorem 2 does not allow the implementation of a BA type of algorithmic embodiment. Due to the structure of the iterations in (10), a non-reversible dependency of $q^{(n+1)}_{\hat{x}}$ on itself appears once $A^{(n)}(x,\hat{x},s)$ is substituted therein, thus requiring the knowledge of $q^{(n+1)}_{\hat{x}}$ to evaluate $q^{(n+1)}_{\hat{x}}$ itself. To circumvent this technical difficulty, we introduce a relaxation in the structure of the iterations of Theorem 2. This results in a variant of an approximate alternating minimization scheme, for which we derive necessary and sufficient conditions that ensure its convergence to a globally optimal point. The aforementioned relaxation and the necessary and sufficient conditions that lead to a globally optimal solution are stated in the following theorem, which is a major result of this paper.

Theorem 3. *(Approximate alternating minimization) Let $s_1 \geq 0$, $s_2 \in [0, s_{2,\max}]$ be given with $s = (s_1, s_2)$ and let $\tilde{A}[q_{\hat{x}}](x,\hat{x},s)$ be*

$$\begin{aligned} \tilde{A}[q_{\hat{x}}](x,\hat{x},s) &= \exp\{-s_1 d(x,\hat{x}) - s_2 \tilde{g}(p_x, v[q_{\hat{x}}], \hat{x})\} \\ \tilde{g}(p_x, v[q_{\hat{x}}], \hat{x}) &= f\left(\frac{p_x(\hat{x})}{v[q_{\hat{x}}](\hat{x})}\right) - \frac{p_x(\hat{x})}{v[q_{\hat{x}}](\hat{x})} \partial f\left(\frac{p_x(\hat{x})}{v[q_{\hat{x}}](\hat{x})}\right) \end{aligned}$$

where $v[q_{\hat{x}}]$ is any probability vector. Let $\tilde{q}^{(0)}_{\hat{x}}$ be any probability vector with nonzero components and let $\tilde{q}^{(n+1)}_{\hat{x}} = Q_{\hat{x}}[\tilde{q}^{(n)}_{\hat{x}}]$ and $\tilde{Q}^{(n+1)}_{\hat{x}|x} = Q_{\hat{x}|x}[\tilde{q}^{(n)}_{\hat{x}}]$ be defined as functions of the past iteration $\tilde{q}^{(n)}_{\hat{x}}$ as follows:

$$\begin{aligned} \tilde{Q}^{(n+1)}_{\hat{x}|x} &= \tilde{q}^{(n)}_{\hat{x}}(\hat{x}) \frac{\tilde{A}^{(n)}(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} \tilde{q}^{(n)}_{\hat{x}}(i) \tilde{A}^{(n)}(x,i,s)} \\ \tilde{q}^{(n+1)}_{\hat{x}} &= \tilde{q}^{(n)}_{\hat{x}}(\hat{x}) \sum_{x\in\mathcal{X}} \frac{p_x(x) \tilde{A}^{(n)}(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} \tilde{q}^{(n)}_{\hat{x}}(i) \tilde{A}^{(n)}(x,i,s)} \end{aligned}$$

where $\tilde{A}^{(n)}(x,\hat{x},s) = \tilde{A}[\tilde{q}_{\hat{x}}^{(n)}](x,\hat{x},s)$. Then, as $n \to \infty$, we obtain

$$D(\tilde{Q}_{\hat{x}|x}^{(n)}) \to D_s,\ P(\tilde{Q}_{\hat{x}|x}^{(n)}) \to P_s,\ I(p_x,\tilde{Q}_{\hat{x}|x}^{(n)}) \to R(D_s,P_s)$$

if and only if $\lim_{n\to\infty} \frac{\tilde{q}_{\hat{x}}^{(n+1)}}{v[\tilde{q}_{\hat{x}}^{(n)}]} \to 1$ with at least linear rate of convergence.

Theorem 3 enables the implementation of the alternating minimization scheme by introducing an auxiliary variable $v[q_{\hat{x}}^{(n)}]$ (approximation), designed as a function of only the current iteration of $q_{\hat{x}}^{(n)}$. Nevertheless, depending on v, this approximation may incur restrictions on the domain of the Lagrangian multiplier s_2 in order to have convergence guarantees. The implementation of Theorem 3 is illustrated in Algorithm 1.

We conclude this section, with a lemma where we provide necessary and sufficient conditions for the optimal solution of the studied problem.

Lemma 2. *Let $D_f(\cdot||\cdot)$ be such that $f \in C^1(0,\infty)$ continuous and differentiable on $(0,\infty)$. Then, a probability vector $q_{\hat{x}}$ yields a point on the $R(D,P)$ curve via the transition matrix*

$$Q_{\hat{x}|x}[q_{\hat{x}}](\hat{x}) = q_{\hat{x}}(\hat{x})\frac{p_x(x)\tilde{A}(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} q_{\hat{x}}(i)\tilde{A}(x,i,s)}$$

if and only if, $\forall \hat{x} \in \hat{\mathcal{X}}$,

$$c(\hat{x}) = \sum_{x\in\mathcal{X}} \frac{p_x(x)\tilde{A}(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} \tilde{q}_{\hat{x}}(i)\tilde{A}(x,i,s)} \leq 1,$$

holding with equality for any $\hat{x}$ for which $q_{\hat{x}}(\hat{x})$ is nonzero.

IV. Analysis of Algorithm 1

In this section, we study a stopping criterion and the convergence rate of Algorithm 1.

Before obtaining stopping conditions for Algorithm 1, we first derive a useful lemma which is needed to obtain the stopping conditions.

Lemma 3. *Let $s_1 \geq 0$, $s_2 \in [0, s_{2,\max})$ be given with $s = (s_1,s_2)$ and let $Q_{\hat{x}|x}$ be a transition matrix included in the set $\mathcal{L}_{(D,P)}$ defined as follows:*

$$\mathcal{L}_{(D,P)} = \{Q_{\hat{x}|x} : \mathbb{E}_{Q_{\hat{x}|x}}[d(x,\hat{x})] \leq D \wedge D_f(p_x||Q_{\hat{x}}) \leq P\}.$$

Then, $\forall \lambda \in \Lambda_{s,v[Q_{\hat{x}}]}$, with $\Lambda_{s,v[Q_{\hat{x}}]} = \big\{\lambda \in \mathcal{R}^{|\mathcal{X}|} : \forall x \in \mathcal{X}, \lambda(x) \geq 0 \wedge \forall \hat{x} \in \hat{\mathcal{X}},\ \sum_{x\in\mathcal{X}} p_x(x)\lambda(x)\tilde{A}(x,\hat{x},s) \leq 1\big\}$, we obtain

$$R(D,P) \geq -\sum_{x\in\mathcal{X}} p_x(x)\log\Big(\frac{1}{\lambda(x)}\Big) - s_1 D - s_2 \sum_{(x,\hat{x})\in\mathcal{X}\times\hat{\mathcal{X}}} p_x(x)Q_{\hat{x}|x}(x,\hat{x})\tilde{g}(\hat{x}).$$

Theorem 4. *(Stopping criterion) Let $\tilde{Q}_{\hat{x}|x}$ and $\tilde{q}_{\hat{x}}$ be defined as in Theorem 3, $c(\hat{x})$ be as defined as in Lemma 2 and $c_{\max} = \max_{\hat{x}\in\hat{\mathcal{X}}} c(\hat{x})$. Then, at the point $D = \mathbb{E}_{\tilde{Q}_{\hat{x}|x}}[d(x,\hat{x})]$, and $P = D_f(p_x||\tilde{Q}_{\hat{x}}))$, the following bounds hold*

$$R(D,P) \geq -s_1 D - s_2 P + \tilde{W}[\tilde{q}_{\hat{x}}] - \log(c_{\max}) \quad (11)$$

$$R(D,P) \leq -s_1 D - s_2 P + \tilde{W}[\tilde{q}_{\hat{x}}] - \sum_{\hat{x}\in\hat{\mathcal{X}}} \tilde{q}_{\hat{x}} c(\hat{x})\log(c(\hat{x})) \quad (12)$$

where $\tilde{W}[\tilde{q}_{\hat{x}}]$ is given by

$$\tilde{W}[\tilde{q}_{\hat{x}}^{(n)}] = -\sum_{x\in\mathcal{X}} p_x(x)\log\Bigg(\sum_{\hat{x}\in\hat{\mathcal{X}}} \tilde{q}_{\hat{x}}^{(n)}(\hat{x})\tilde{A}^{(n)}(x,\hat{x},s)\Bigg) + s_2 \sum_{\hat{x}\in\hat{\mathcal{X}}} p_x(\hat{x})\frac{p_x(\hat{x})}{v^{(n)}(\hat{x})}\partial f\Big(\frac{p_x(\hat{x})}{v^{(n)}(\hat{x})}\Big) + s_2\Bigg[\sum_{\hat{x}\in\hat{\mathcal{X}}} \tilde{q}_{\hat{x}}^{(n+1)}\Bigg(f\Big(\frac{p_x(\hat{x})}{q_{\hat{x}}^{(n+1)}(\hat{x})}\Big) - f\Big(\frac{p_x(\hat{x})}{v^{(n)}(\hat{x})}\Big)\Bigg)\Bigg]. \quad (13)$$

Algorithm 1 Implementation of Theorem 3

Require: source distribution p_x; Lagrangian multipliers $s = (s_1,s_2)$ with $s_1 \geq 0$ and $s_2 \in [0, s_{2,\max}]$; error tolerance ϵ; divergence measure $D_f(\cdot||\cdot)$; distortion measure $d(\cdot,\cdot)$; initial assignment $q_{\hat{x}}^{(0)}$.

- $n \leftarrow 0$
- flag $\leftarrow 0$
- **while** flag $== 0$ **do**
  - $g(\hat{x}) \leftarrow f\Big(\frac{p_x(\hat{x})}{v^{(n)}(\hat{x})}\Big) - \frac{p_x(\hat{x})}{v^{(n)}(\hat{x})}\partial f\Big(\frac{p_x(\hat{x})}{v^{(n)}(\hat{x})}\Big)$
  - $\tilde{A}^{(n)}(x,\hat{x},s) \leftarrow \exp[-s_1 d(x,\hat{x}) + s_2 g(p_x, v[q_{\hat{x}}],\hat{x})]$
  - $c^{(n)}(\hat{x}) \leftarrow \sum_{x\in\mathcal{X}} \frac{p_x(x)\tilde{A}^{(n)}(x,\hat{x},s)}{\sum_{i\in\mathcal{X}} \tilde{q}_{\hat{x}}^{(n)}(i)\tilde{A}^{(n)}(x,i,s)}$
  - $q_{\hat{x}}^{(n+1)} \leftarrow q_{\hat{x}}^{(n)} \cdot c^{(n)}$
  - $\omega \leftarrow \log c_{\max}^{(n)}(\hat{x}) - \sum_{\hat{x}\in\hat{\mathcal{X}}} q_{\hat{x}}^{(n)} c^{(n)}(\hat{x})\log(c^{(n)}(\hat{x}))$
  - **if** $\omega \leq \epsilon$ **then**
    - flag $\leftarrow 1$
  - **end if**
  - $n \leftarrow n+1$
- **end while**

Ensure: $Q_{\hat{x}|x} = q_{\hat{x}}^{(n)}(\hat{x})\frac{\tilde{A}^{(n)}(x,\hat{x},s)}{\sum_{i\in\hat{\mathcal{X}}} q_{\hat{x}}^{(n)}(i)\tilde{A}^{(n)}(x,i,s)}$, $D_s = \mathbf{E}_{p_x Q_{\hat{x}|x}}[d(x,\hat{x})]$, $P_s = D_f(p_x, q_{\hat{x}}^{(n)})$, $R(D_s,P_s) = \tilde{W}[q_{\hat{x}}^{(n)}] - s_1 D_s - s_2 P_s - \sum_{\hat{x}\in\hat{\mathcal{X}}} q_{\hat{x}}^{(n)} c(\hat{x})\log(c(\hat{x}))$, $\tilde{W}[\cdot] =$ (13).

Next, we study the convergence rate of Algorithm 1. This is done by first studying the convergence rate of the alternating minimization procedure of Theorem 2 and then, using it as a reference to analyze the convergence rate of the approximate alternating minimization procedure of Theorem 3.

Note that based on the structure of the parametric solution of $q_{\hat{x}}^{(n+1)}$ at the previous iteration n, we can similarly define a vector function $S : \mathcal{R}^{|\mathcal{X}|} \to \mathcal{R}^{|\mathcal{X}|}$ as $S[q_{\hat{x}}](i) = q_{\hat{x}}(i)\cdot c[q_{\hat{x}}](i)$. Using Lemma 2, a distribution $q_{\hat{x}}^*$ that achieves the RDPF is a fixed point of $S(q_{\hat{x}})$. Following [24], we can analyze the convergence rate of the alternating minimization procedure in both Theorems 2 and 3. Using the first order Taylor expansion of $S[q_{\hat{x}}]$ around a fixed point $q^*_{\hat{x}}$, we obtain $S[q_{\hat{x}}] = S[q^*_{\hat{x}}] + J(q^*_{\hat{x}}) \cdot (q_{\hat{x}} - q^*_{\hat{x}}) + o(||q_{\hat{x}} - q^*_{\hat{x}}||)$, where $J(q_{\hat{x}})$ is the Jacobian matrix of $S(q_{\hat{x}})$ with entries $J_{i,j}(q_{\hat{x}}) \triangleq \frac{\partial S[q_{\hat{x}}](i)}{\partial q_{\hat{x}}(j)}, (i,j) \in \mathcal{X} \times \hat{\mathcal{X}}$. The next theorem provides the functional form of the Jacobian for the case of Theorem 2.

Theorem 5. *(Jacobian form) The Jacobian $J(q_{\hat{x}})$ computed at the fixed point $q^*_{\hat{x}}$ is given as*

$$J(q^*_{\hat{x}}) = (I - M)(I - \Gamma J(q^*_{\hat{x}})) \tag{14}$$

where

$$M \triangleq \left[q^*_{\hat{x}}(i) \sum_{x \in \mathcal{X}} p_x(x) \frac{A(x,i,s)A(x,j,s)}{(\sum_{k \in \hat{\mathcal{X}}} q^*_{\hat{x}}(k) A(x,k,s))^2} \right]_{(i,j) \in \mathcal{X} \times \hat{\mathcal{X}}} \tag{15}$$

$$\Gamma \triangleq s_2 \cdot \operatorname{diag} \left[q^*_{\hat{x}}(i) \cdot \frac{\partial^2}{\partial q(i)^2} D_f(p_x||q) \Big|_{q^*_{\hat{x}}} \right]_{i \in \mathcal{X}}. \tag{16}$$

Next, we introduce two lemmas, in which we use the structure of (14) to identify properties of matrix M.

Lemma 4. *Let $\{\lambda_i\}_{i=1:|\mathcal{X}|}$ be the set of eigenvalues of M. Then, given a distortion function $d : \mathcal{X} \times \hat{\mathcal{X}} \to \mathcal{R}^+_0$ that induces a full-rank matrix $D = [e^{-s_1 d(i,j)}]_{(i,j) \in \mathcal{X} \times \hat{\mathcal{X}}}$, then $\lambda_i > 0, \forall i \in [1 : |\mathcal{X}|]$, i.e., M has only positive eigenvalues.*

Remark 2. *(On Lemma 4) We note that a popular example that satisfies the assumptions imposed on Lemma 4 is the Hamming distortion denoted hereinafter as d_H [25].*

Lemma 5. *Let $\{\lambda_i\}_{i=1:|\mathcal{X}|}$ be the set of eigenvalues of M. Then, at fixed point $q^*_{\hat{x}}$, we have that $\lambda_i \leq 1, \forall i \in 1 : |\mathcal{X}|$.*

Using Lemmas 4 and 5, we can now characterize the interval that contains the set of eigenvalues of $J(q^*_{\hat{x}})$ and subsequently the convergence rate of Theorem 2.

Theorem 6. *(Convergence rate of Theorem 2) Let $\{\theta_i\}_{i \in \mathcal{X}}$ be the eigenvalues of $J(q^*_{\hat{x}})$. Then,*

$$0 \leq \{\theta_i\}_{i \in \mathcal{X}} < 1.$$

Moreover, let $\gamma \in [\theta_{\max}, 1)$. Then, there exists $\delta > 0$ and $K > 0$ such that if $q^0_{\hat{x}} \in \{q_{\hat{x}} : ||q_{\hat{x}} - q^*_{\hat{x}}|| \leq \delta\}$, we obtain

$$||q^{(n)}_{\hat{x}} - q^*_{\hat{x}}|| < K \cdot ||q^{(0)}_{\hat{x}} - q^*_{\hat{x}}|| \cdot \gamma^n \tag{17}$$

i.e., the iterations converge exponentially.

Following similar steps that led to Theorem 6, we can now write the Jacobian $J_a(q^*_{\hat{x}})$ for Theorem 3 considering a specific form of the auxiliary variable $v[q^{(n)}_{\hat{x}}] = q^{(n)}_{\hat{x}}$. This results into the following structure of the Jacobian matrix

$$J_a(q^*_{\hat{x}}) = (I - M)(I - \Gamma)$$

where M and Γ are given by (15) and (16), respectively. Unlike Theorem 2, where the structure of (14) bounds its own eigenvalues, in this case we need to bound the Lagrangian multiplier s_2, hence matrix Γ, to guarantee exponential convergence of the algorithm. This is proved in the following theorem.

Theorem 7. *For a given $s_1 \geq 0$, let $I_{s_2} = [0, s_{2,\max}]$ be the domain of s_2, $\{\theta_{a,i}\}_{i \in \mathcal{X}}$ the set of eigenvalues of $J_a(q^*_{\hat{x}})$ and $\theta_{\max}$ the maximum eigenvalue of $J(q^*_{\hat{x}})$ in (14). Define the set $I^{\epsilon}_{s_2} = [0, s_{2,\max} - \epsilon]$ for $0 < \epsilon < s_{2,\max}$. Then, there exists an ϵ' such that if $s_2 \in I^{\epsilon'}_{s_2}$ then $0 \leq \{\theta_{a,i}\}_{i \in \mathcal{X}} < 1$.*

Theorem 7 guarantees exponential convergence for Theorem 3 only for $s_2 \in I^{\epsilon}_{s_2}$ which means that we are able to consider $P \in [P_{\min}(\epsilon), P_{\max}]$, depending on the characteristics of the specific problem.

V. Numerical Results

In this section, we provide numerical results to demonstrate the utility of Algorithm 1.

Example 1. *Suppose that $\mathcal{X} = \hat{\mathcal{X}} = \{0, 1\}$ and let $p_x \sim Ber(0.15)$ with $d(\cdot,\cdot) = d_H(\cdot,\cdot)$ and perception constraint chosen to be either (a) $D_f(p_x||q_{\hat{x}}) = TV(p_x||q_{\hat{x}}) = \frac{1}{2}\sum_{i \in \mathcal{X}} |p_x(i) - q_{\hat{x}}(i)|$, or (b) $D_f(\cdot||\cdot) = D_{KL}(\cdot||\cdot)$, or (c) $D_f(\cdot||\cdot) = D_{\chi^2}(\cdot||\cdot)$.*

In Fig. 1a we compare the theoretical results of [9, Equation 6] with the numerical results obtained using Algorithm 1. We observe that Algorithm 1 achieves exactly the theoretical solution of [9, Equation 6] as long as $D \leq D_{\max} = 0.15$.

In Fig. 1b and Figure 1c we use Algorithm 1 to compute $R(D, P)$ for $D_f(\cdot||\cdot) = D_{KL}(\cdot||\cdot)$, and for $D_f(\cdot||\cdot) = D_{\chi^2}(\cdot||\cdot)$, respectively.

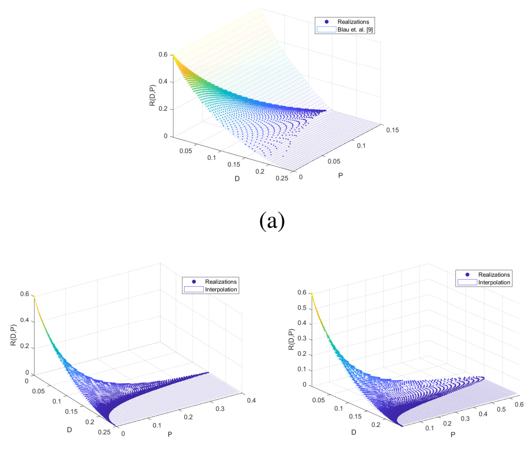

(a)

(b) (c)

Fig. 1: $R(D, P)$ for a Bernoulli source under a Hamming distortion and various perception constraints.

Acknowledgement

This work 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] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.

[2] ——, “Coding theorems for a discrete source with a fidelity criterion,” *Institute of Radio Engineers, National Convention Record*, vol. 4, pp. 142–163, 1993.

[3] 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.

[4] 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.

[5] 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.

[6] C. Ledig *et al.*, “Photo-realistic single image super-resolution using a generative adversarial network,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, jul 2017, pp. 105–114.

[7] 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.

[8] 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.

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

[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] P. A. Stavrou and M. Kountouris, “The role of fidelity in goal-oriented semantic communication: A rate distortion approach,” *IEEE Trans. Commun.*, to appear, 2023.

[12] 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.

[13] C. T. Li and A. E. Gamal, “Strong functional representation lemma and applications to coding theorems,” *IEEE Trans. Inf. Theory*, vol. 64, no. 11, pp. 6967–6978, 2018.

[14] 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.

[15] A. B. Wagner, “The rate-distortion-perception tradeoff: The role of common randomness,” 2022. [Online]. Available: https://arxiv.org/abs/2202.04147

[16] J. Qian, G. Zhang, J. Chen, and A. Khisti, “A rate-distortion-perception theory for binary sources,” in *International Zurich Seminar on Information and Communication*, 2022, pp. 34–38.

[17] I. Sason, “On f-divergences: Integral representations, local behavior, and inequalities,” *Entropy*, vol. 20, no. 5, 2018.

[18] I. Csiszar and G. Túsnady, “Information geometry and alternating minimization procedures,” *Statistics and Decisions, Dedewicz*, vol. 1, pp. 205–237, 1984.

[19] R. Yeung, *Information theory and network coding*. Springer, 2008.

[20] R. E. Blahut, “Computation of channel capacity and rate-distortion functions,” *IEEE Trans. Inf. Theory*, vol. 18, no. 4, pp. 460–473, 1972.

[21] R. T. Rockafellar and R. J. B. Wets, *Variational Analysis*. Springer Verlag, 1998.

[22] I. Csiszár and P. C. Shields, “Information theory and statistics: A tutorial,” *Foundations and Trends® in Communications and Information Theory*, vol. 1, no. 4, pp. 417–528, 2004.

[23] R. E. Blahut, *Principles and practice of Information Theory*. Addison-Wesley, 1987.

[24] K. Nakagawa, Y. Takei, S.-i. Hara, and K. Watabe, “Analysis of the convergence speed of the Arimoto-Blahut algorithm by the second-order recurrence formula,” *IEEE Trans. Inf. Theory*, vol. 67, no. 10, pp. 6810–6831, 2021.

[25] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. John Wiley & Sons, Inc., NJ, 2006.