Low-Latency Rate-Distortion-Perception Trade-off: A Randomized Distributed Function Computation Application

Onur Günlü[1], Maciej Skorski[2], and H. Vincent Poor[3]
[1]Information Theory and Security Laboratory (ITSL), Linköping University, onur.gunlu@liu.se
[2]Faculty of Information Technology, Czech Technical University in Prague, maciej.skorski@uw.edu.pl
[3]Department of Electrical and Computer Engineering, Princeton University, poor@princeton.edu

***Abstract*—Semantic communication systems, which focus on transmitting the semantics of data rather than its exact reconstruction, redefine the design of communication networks for transformative efficiency in bandwidth-limited and latency-critical applications. Addressing these goals, we tackle the rate-distortion-perception (RDP) problem for image compression, a critical challenge in achieving perceptually realistic reconstructions under rate constraints. Formulated within the randomized distributed function computation (RDFC) framework, we establish an achievable non-asymptotic RDP region, providing finite blocklength trade-offs between rate, distortion, and perceptual quality, aligning with semantic communication objectives. We extend this region to also include a secrecy constraint, providing strong secrecy guarantees against eavesdroppers via physical-layer security methods, ensuring resilience against quantum attacks. Our contributions include (i) establishing achievable bounds for non-asymptotic RDP regions under realism and distortion constraints; (ii) extending these bounds to provide strong secrecy guarantees; (iii) characterizing the asymptotic secure RDP region under a perfect realism constraint; and (iv) illustrating significant reductions in rates and the effects of secrecy constraints and finite blocklengths. Our results provide actionable insights for designing low-latency, high-fidelity, and secure image compression systems with realistic outputs, advancing applications, e.g., in privacy-critical domains.**

I. Introduction

The rapid evolution of communication systems toward semantic communications marks an important shift in the design and objectives of modern communication networks. Unlike conventional approaches that prioritize exact signal reconstruction, semantic communication systems transmit the semantics of the data, aligning more closely with the requirements of real-world applications [1], [2]. This transition is particularly impactful in bandwidth-constrained and low-latency scenarios, such as immersive multimedia systems, autonomous vehicles, and augmented reality, where transmitting only the semantically significant features of the data can drastically reduce communication overhead while maintaining functional utility.

The semantic communications problem is shown to be an instance of remote source coding problems [3], [4, pp. 118], [5, pp. 78], where the receiver computes a hidden function of the data observed at the transmitter. We recently extended this idea to introduce the randomized distributed function computation (RDFC) framework [6]. This framework takes into account that numerous practical distributed function computation scenarios require a controlled randomization step. Example applications that require such a controlled randomization in the form of preserving a synthesized probability distribution include neural data compression methods with generative models [7], [8], federated learning with side information [9], and compression methods used as differential privacy mechanisms [6], [10], [11]. The synthesis of the randomization step by using coded encoding and decoding methods, as in RDFC,–rather adding additional random noise– allows the receiver to generate outputs that align statistically with a target distribution while maintaining the output utility. This flexibility enables substantial gains in communication efficiency, especially in latency-sensitive and privacy-critical environments. Moreover, the RDFC framework provides strong performance guarantees with limited or no common randomness shared between the transmitter and receiver, as the performance guarantees are for each function computation instance [6], [12], [13].

In this work, we consider the rate-distortion-perception (RDP) problem that considers, for instance, image compression applications [14, Section 17.4.2]. The RDP problem is a specific instance of RDFC and a crucial challenge in achieving perceptually realistic image reconstructions under rate constraints [15]–[19]. This problem extends the classical rate-distortion trade-off by incorporating a perceptual quality constraint, ensuring that reconstructed signals, such as images, not only minimize distortion but also align with human-perceived quality. There are powerful deep learning-based discriminators that are used to output images indistinguishable from the original images, such as in [20]–[22]. A common way to impose a realism constraint such that its mathematical analysis is tractable is to impose that the probability distribution of the reconstructed image is close to the probability distribution of the original image; see, for instance, [23], [24] for an extensive summary and a list of such existing methods. This set of realism constraints is a special case of the RDFC framework, as the synthesized randomization step should ensure for the former that the receiver outputs are distributed according to a probability distribution close to the transmitter inputs' probability distribution. In this work, by formulating the RDP problem within the RDFC framework, we adopt coded RDFC methods with strong function computation guarantees for limited or no common randomness to establish the non-asymptotic limits of

RDP in finite blocklength regimes.

We characterize achievable non-asymptotic trade-offs between rate, distortion, and perceptual quality for the RDP problem. These non-asymptotic limits will provide system designers with theoretical benchmarks against which practical strategies can be tested and with design guidelines, offering actionable insights for real-world applications where low-latency image compression is vital.

Moreover, distributed function computation problems are usually susceptible to information leakage through public communication channels [25]. Since we propose coding-theoretic methods for the RDP problem, these code constructions can be modified by using physical layer security (PLS) methods to minimize information leakage to an eavesdropper. Thus, we also consider image compression scenarios, where the transmitter output is observed by not only the receiver but also any eavesdropper. This provides an additional layer of security guarantee that is independent of the attackers' computational power, making them resistant to quantum attacks, unlike classical cryptographic methods [26]. Moreover, PLS is also known to be vital for joint source-channel coding-based image compression applications [27], which are direct extensions of our RDP models. We establish asymptotic and non-asymptotic secure RDP regions with strong secrecy guarantees, where the leakage is asymptotically negligible, unlike common PLS methods that limit the normalized leakage, providing weak secrecy guarantees.

A. Main Contributions and Paper Organization

A summary of the main contributions is as follows. We

- Establish an achievable non-asymptotic RDP region under realism and distortion constraints, providing inner bounds on required rates for finite blocklengths with strong guarantees on perceptual quality;
- Extend the non-asymptotic RDP region to include an information-leakage constraint, deriving achievable bounds that ensure strong secrecy guarantees while maintaining perceptual quality and distortion performance;
- Characterize the asymptotic secure RDP region under perfect realism constraint, demonstrating the relationship between near-perfect and perfect realism constraints; and
- Analyze the rate regions by (i) illustrating significant communication load reductions achieved by RDFC methods over classical data compression methods; (ii) evaluating the impact of incorporating secrecy constraints on achievable rates; and (iii) drawing analogies between non-asymptotic and asymptotic results.

In Section II, we define the model used for a point-to-point RDFC problem aiming to obtain high perceptual quality in image compression under distortion and communication constraints, as well as nonasymptotic RDP regions with and without secrecy constraint. In Section III, we establish three main results of this work by characterizing two achievable rate regions for finite lengths and the optimal rate region for secure RDP problems in the asymptotic regime. In Section IV, we compare the established rate regions with each other and with classical data compression methods, as well as their asymptotic counterparts. In Section V, we provide the proof sketches for the two main results of this work given in Section III. In Section VI, we conclude the paper and list the potential scientific and technological impact of our RDP results.

B. Notation

We represent random variables with upper-case letters X and their realizations with lower-case letters x. A random variable X has probability distribution P_X with support $\mathrm{supp}(P_X)$. Calligraphic letters, e.g., $\mathcal{X}$, denote sets with cardinality $|\mathcal{X}|$. We represent n-letter random variable sequences as $X^n = X_1, X_2, \dots, X_n$. Denote the probability distribution of a sequence of independent and identically distributed (i.i.d.) random variables as $P_X^n \triangleq \prod_{i=1}^n P_X$. Denote total variation distance as

$$\|P_Y - P_X\|_{\mathrm{TV}} \triangleq \frac{1}{2}\sum_{b\in\mathcal{B}} |P_Y(b) - P_X(b)|. \tag{1}$$

$\mathcal{O}(\cdot)$ denotes the Big O notation. Denote information density of a probability distribution P_{XY} as

$$\imath(X,Y) = \log\frac{P_{XY}(x,y)}{P_X(x)P_Y(y)}. \tag{2}$$

The variance of a random variable is denoted as $\mathrm{Var}[\cdot]$ and the inverse Q-function, i.e., the tail distribution function of a standard normal distribution, as $Q^{-1}(\cdot)$. Define for any $k \in \mathbb{R}$, $[k]^+ = \max\{k, 0\}$. Logarithms are in base 2. $[a:b]$ denotes the set of integers $\{a, a+1, \dots, b\}$. $\{\cdot\}^c$ denotes a complementary event. For any $\delta > 0$, a sequence x^n is defined to be δ-letter typical with respect to P_X, denoted as $x^n \in T_\delta^n(P_X)$ if the empirical probability distribution $N(.|x^n)/n$ satisfies

$$\left|\frac{N(a|x^n)}{n} - P_X(a)\right| \le \delta P_X(a) \qquad \text{for all } a \in \mathcal{X}. \tag{3}$$

II. Problem Definition

Consider the point-to-point RDFC problem that aims to achieve a high perceptual quality under communication-rate and reconstruction-distortion constraints, as depicted in Fig. 1. The encoder observes an image $X^n \in \mathcal{X}^n$, where $\mathcal{X}$ is finite, and has access to common randomness $C \in [1:2^{nR_0}]$ shared between the encoder and decoder, which is uniformly distributed and independent of X^n. One promising way to obtain such a uniformly-distributed and independent common randomness is the use of digital security primitives, called physical unclonable functions [28], from which random sequences can be extracted and distributed. The encoder outputs an index $S \in [1:2^{nR}]$, obtained as $S = \mathsf{Enc}(X^n, C)$, where $\mathsf{Enc}(\cdot)$ is the encoding function. Assume that the index S is observed noiselessly at the decoder, which can be achieved by using error-correcting codes for reliable communications. The aim of the decoder is to output an image $Y^n = y^n$, obtained as $Y^n = \mathsf{Dec}(S, C) \in \mathcal{X}^n$, where $\mathsf{Dec}(\cdot)$ is the decoding function, such that

(i) the induced output image probability distribution P_{Y^n}, where $y^n \sim P_{Y^n}$, is almost equal to the input image probability

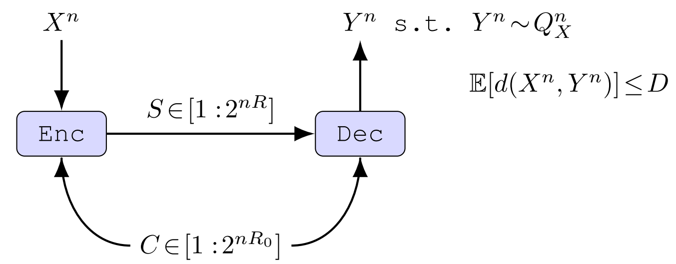

Fig. 1. An RDP model for deep learning-based image compression applications, where a realism constraint is imposed by ensuring that the distribution of the image reconstructed at the receiver is close to the distribution of the original image, i.e., $Y^n \sim Q_X^n$ given $X^n \sim Q_X^n$. This enables high perceptual quality for the reconstructed image Y^n. Moreover, the expected distortion between X^n and Y^n should be minimized to limit the image compression's effect on the image quality. We consider low-latency RDP by considering a finite blocklength n and minimizing the rate R for a given common randomness rate $R_0 \geq 0$.

distribution Q_X^n;
(ii) the rate R is minimized, given a common randomness rate R_0; and
(iii) the distortion between the input and output images is minimized.

Nonasymptotic performance limits of the RDP trade-off can be characterized by fixing the blocklength $n \geq 1$, which provides low-latency image compression. Thus, we next define two non-asymptotic regions for this RDFC problem with and without a secrecy constraint against an eavesdropper who might observe the index S sent through a public communication channel. In the following, assume any $\epsilon_{\rm r}, \epsilon_{\rm D}, \epsilon_{\rm sec} > 0$.

Definition 1. *An RDP tuple (R, R_0, D) is $(\epsilon_r, \epsilon_D, n)$-*achievable *for Q_X if there exist one encoder and one decoder such that*

$$||P_{Y^n} - Q_X^n||_{TV} \leq \epsilon_r \qquad (realism) \tag{4}$$
$$\mathbb{E}\big[d(X^n, Y^n)\big] \leq D + \epsilon_D \qquad (distortion) \tag{5}$$

where $d(x^n, y^n) = \dfrac{1}{n}\sum_{i=1}^{n} d(x_i, y_i)$ is any per-letter distortion metric bounded from above by a value $d_{max} > 0$.

The nonasymptotic RDP region $\mathcal{R}_{RDP}$ is the closure of the set of all achievable tuples. ◊

We next include a secrecy constraint to consider the information leakage about the reconstructed image Y^n to an eavesdropper who observes the index S. This is particularly relevant for applications such as generative artificial intelligence in the creation of artistic digital content, where the reconstructed output represents the final, valuable product.

Definition 2. *An RDP tuple (R, R_0, D) is $(\epsilon_r, \epsilon_D, \epsilon_{sec}, n)$-*achievable *for Q_X under a secrecy constraint if there exist one encoder and one decoder such that (4), (5), and*

$$||P_{SY^n} - P_S P_{Y^n}||_{TV} \leq \epsilon_{sec} \qquad (secrecy). \tag{6}$$

The nonasymptotic secure RDP region $\mathcal{R}_{SRDP}$ is the closure of the set of all achievable tuples under the secrecy constraint imposed. ◊

Note that the secrecy leakage constraint in (6) corresponds to a strong secrecy constraint, which measures unnormalized information leakage, unlike classical weak secrecy constraint, used, for instance, in [29].

We next provide non-asymptotic RDP tuples achievable for $\mathcal{R}_{\rm RDP}$ and $\mathcal{R}_{\rm SRDP}$.

III. Nonasymptotic Limits of RDP Trade-off

Similar to [30]–[32], denote channel dispersions for the channels $P_{U|X}$ and $P_{U|Y}$, respectively, as

$$V_{U|X} = \mathbb{E}_{P_{UX}}\big[\mathrm{Var}[\imath(U,X)|U]\big], \tag{7}$$
$$V_{U|Y} = \mathbb{E}_{P_{UY}}\big[\mathrm{Var}[\imath(U,Y)|U]\big]. \tag{8}$$

We remark that $P_{U|X}$ and $P_{U|Y}$ are test channels connecting X and Y from the model in Fig. 1 to an auxiliary random variable U, which represents the designed codebook. Define

$$\mu_{xy} = \min_{(x,y)\in\mathrm{supp}(P_{XY})} P_{XY}(x,y). \tag{9}$$

Now, we provide an $(\epsilon_{\rm r}, \epsilon_{\rm D}, n)$-achievable nonasymptotic RDP region $\mathcal{R}_{\rm RDP}$. The proof sketch of Theorem 1 is given in Section V-A below.

Theorem 1. *An $(\epsilon_r, \epsilon_D, n)$-achievable nonasymptotic RDP region is the union over all joint distributions P_{XUY} of the rate tuples (R, R_0, D) satisfying*

$$R \geq \Big[I(U;X) + Q^{-1}\Big(\epsilon_r + \mathcal{O}\big(\tfrac{1}{\sqrt{n}}\big)\Big)\sqrt{\frac{V_{U|X}}{n}} + \mathcal{O}\Big(\frac{\log n}{n}\Big)\Big]^+, \tag{10}$$

$$R + R_0 \geq \Big[I(U;Y) + \mathcal{O}\Big(\frac{\log n}{n}\Big) + Q^{-1}\Big(\epsilon_r + \mathcal{O}\big(\tfrac{1}{\sqrt{n}}\big)\Big)\sqrt{\frac{V_{U|Y}}{n}}\Big]^+ \tag{11}$$

where $X - U - Y$ form a Markov chain, and

$$D \geq \mathbb{E}[d(X,Y))] - \delta_D \tag{12}$$

such that

$$\epsilon_D = \delta_D(1 + D + \delta_D) + 2|\mathcal{X}|^2 e^{-2n\delta_D^2\mu_{xy}^2} d_{max}. \tag{13}$$

It suffices to consider $|\mathcal{U}| \leq |\mathcal{X}|^2 + 1$.

The asymptotic counterparts of the bounds in (10)-(13) recover the RDP region established in [24, Theorem 6] that extends [33, Theorems 1 and 5], which follows by allowing $n \to \infty$ in (10)-(13).

Remark 1. *The asymptotic RDP region, obtained by allowing $n \to \infty$ in Theorem 1, shows that a rate of $R = I(U;X)$ is asymptotically achievable if enough common randomness C is available, i.e., if $R_0 \geq I(U;Y) - I(U;X)$. Since the minimum rate a lossless compression method can achieve is $H(X)$, the RDP methods can provide remarkable gains, measured as $H(X)/I(U;X)$, as compared to classical data compression methods. The rate gains from using common randomness for RDFC are illustrated in [6] to be more than* 214 *times when*

RDFC methods are used for differential privacy applications. Thus, depending on the target probability distortion Q_X and the choice of the distortion metric $d(\cdot,\cdot)$, such gains can also be illustrated for the RDP problem.

Next, we provide an $(\epsilon_{\text{r}}, \epsilon_{\text{D}}, \epsilon_{\text{sec}}, n)$-achievable nonasymptotic secure RDP region $\mathcal{R}_{\text{SRDP}}$. The proof sketch of Theorem 2 is given in Section V-B below.

Theorem 2. *An $(\epsilon_r, \epsilon_D, \epsilon_{sec}, n)$-achievable nonasymptotic secure RDP region is the union over all joint distributions P_{XUY} of the rate tuples (R, R_0, D) satisfying, for any $\theta \in [0,1]$, (12), (13), and*

$$R \geq \Big[I(U;X) + Q^{-1}\Big(\theta\big(\epsilon_r + \mathcal{O}\big(\frac{1}{\sqrt{n}}\big)\big)\Big)\sqrt{\frac{V_{U|X}}{n}} + \mathcal{O}\Big(\frac{\log n}{n}\Big)\Big]^+, \tag{14}$$

$$R_0 \geq \Big[I(U;Y) + \mathcal{O}\Big(\frac{\log n}{n}\Big) + Q^{-1}\Big((1-\theta)\big(\epsilon_{sec} + \mathcal{O}\big(\frac{1}{\sqrt{n}}\big)\big)\Big)\sqrt{\frac{V_{U|Y}}{n}}\Big]^+ \tag{15}$$

where $X - U - Y$ form a Markov chain. It suffices to consider $|\mathcal{U}| \leq |\mathcal{X}|^2 + 1$.

Note that the asymptotic counterparts of the bounds in (12)-(15) recover the secure RDP region established in [34], which follows by allowing $n \to \infty$ and removing multiplications with θ and $(1-\theta)$ in (14) and (15), respectively.

The realism constraint imposed in (4) is sometimes called a *near-perfect* realism constraint, since for (4) one can achieve an ϵ_r such that $\epsilon_r \to 0$ when $n \to \infty$, but ϵ_r is not equal to zero. The *perfect realism* constraint imposes

$$\|P_{Y^n} - Q_X^n\|_{\text{TV}} = 0 \qquad \text{(perfect realism)} \tag{16}$$

for which we have the following asymptotic result.

Theorem 3. *The asymptotic secure RDP region with perfect realism constraint is the union over all joint distributions P_{XUY} of the rate tuples (R, R_0, D) satisfying*

$$R \geq I(U;X), \tag{17}$$

$$R_0 \geq I(U;Y), \tag{18}$$

$$D \geq \mathbb{E}[d(X,Y)] \tag{19}$$

where $X - U - Y$ forms a Markov chain. It suffices to consider $|\mathcal{U}| \leq |\mathcal{X}|^2 + 1$.

Proof Sketch. The proof follows by combining the asymptotic counterparts of the bounds in (12)-(15) by allowing $n \to \infty$, which recover the region established in [34], with [23, Theorem 1]. The latter proves that if $(d(\cdot,\cdot), Q_X)$ is uniformly integrable, defined in [23, Definition 3], then the rate tuples (R, R_0, D) are achievable with near-perfect realism if and only if they are achievable with perfect realism constraint. Note that it is shown in [23] that any $(d(\cdot,\cdot), Q_X)$ pair is uniformly integrable for finite $\mathcal{X}$, thus combining the results in [34] with [23, Theorem 1] proves Theorem 3. □

Remark 2. *Although the asymptotic secure RDP regions with near-perfect and perfect realism constraints have the same bounds, the optimal code constructions are not necessarily the same.*

IV. Rate Region Comparisons and Analysis

A. Comparisons Between The Rate Regions With and Without Secrecy Constraint

The achievable rate regions with and without secrecy constraints differ not only in their bounds but also in their structural forms. In the achievable secure RDP region, each rate component is separately bounded to meet the requirements of realism and secrecy. Moreover, the realism ϵ_{r} and secrecy ϵ_{sec} parameters are scaled by θ and $(1-\theta)$, respectively, which causes larger additive terms in the bounds as compared to the achievable non-secure RDP region. Moreover, the absence of secrecy constraints in the achievable non-secure RDP region results in a sum-rate constraint, which expands the achievable rate region as compared to the achievable secure RDP region. These structural differences underscore the additional cost of providing strong secrecy guarantees, as achieving them requires higher communication or common randomness rates, whose effect is increased for finite blocklengths. These differences are crucial for applications where the decision between secure and non-secure compression depends on specific trade-offs among security, latency, distortion, and realism performance.

B. Comparisons Between The Rate Regions and Their Asymptotic Counterparts

Comparisons between the nonasymptotic achievable RDP regions and their asymptotic counterparts elucidate the impact of finite blocklength constraints on the RDP trade-offs. In the non-asymptotic regimes, the achievable bounds include terms such as $Q^{-1}(\epsilon)\sqrt{V/n}$, reflecting finite-length penalties due to the statistical fluctuations inherent in finite blocklengths. The additional terms have a similar form to the terms in the finite-length channel and source coding results [30], [35]. These penalties vanish as $n \to \infty$, recovering the asymptotic rate regions, as mentioned above. However, their presence in finite blocklength settings imposes stricter constraints on rates and distortion as compared to asymptotic results. The insights gained from finite blocklength analysis is crucial for establishing reference performance baselines for image compression system designs that must operate efficiently within stringent latency, realism, and distortion constraints.

V. Proof Sketches of Theorem 1 and Theorem 2

A. Proof Sketch of Theorem 1

Proof Sketch. The achievability proof follows by applying nonasymptotic binning methods proposed and used, e.g., in [32], [36], [37], which basically develop finite-length code constructions as in [30], [35] for the output statistics of random binning (OSRB) method [38], [39]; see also [40]. We next present the main adaptations and nuanced details of the proof techniques employed in our work, emphasizing the differences from their conventional usage in the literature.

Consider i.i.d. random variables (X^n, U^n, Y^n) such that

$$\mathbb{E}[d(X,Y)] \leq D + \delta_D \tag{20}$$

for some $\delta_D \geq 0$ satisfying (13). Define the error event that the sequences (X^n, Y^n) are not δ_{D}-letter typical as

$$\mathcal{E} = \{(X^n, Y^n) \notin T^n_{\delta_{\text{D}}}(P_{XY})\}. \tag{21}$$

Similar to [41], we have (13), given (5), since

$$\begin{aligned}
&\mathbb{E}[d(X^n, Y^n)] \\
&= \Pr[\mathcal{E}]\ \mathbb{E}[d(X^n, Y^n)|\mathcal{E}] + \Pr[\mathcal{E}^c]\ \mathbb{E}[d(X^n, Y^n)|\mathcal{E}^c] \\
&\overset{(a)}{\leq} \Pr[\mathcal{E}]\ d_{\max} + \Pr[\mathcal{E}^c]\ (1+\delta_{\text{D}})\ \mathbb{E}[d(X,Y))] \\
&\overset{(b)}{\leq} 2|\mathcal{X}|^2 e^{-2n\delta_{\text{D}}^2\mu_{xy}^2}\ d_{\max} + (1+\delta_{\text{D}})\ (D+\delta_{\text{D}})
\end{aligned} \tag{22}$$

where (a) follows from the typical average lemma [42, pp. 26] and since the distortion metric is per-letter with bound $d_{\max}$, and (b) follows from the bound on the probability of the error event $\mathcal{E}$ given in [43, Eq. (6.34)], applied since a per-letter estimator is used, and by (20).

We demonstrate the existence of non-asymptotic random binning schemes simultaneously meeting the realism and distortion constraints. Following the structure of the OSRB method, we first analyze a source coding problem, called Protocol A, closely related to our problem. In this problem, the encoder observes (U^n, X^n) and then maps U^n independently and uniformly to three random bin indices: $F \in [1:2^{n\widetilde{R}}]$, $S \in [1:2^{nR}]$, and $C \in [1:2^{nR_0}]$. In Protocol A, the index F represent the public choice of encoder-decoder pairs. Using a mismatch stochastic likelihood coder (SLC) as the decoder, as in [32, Eq. (12)] and [37, pp. 3], we bound the expected error probability averaged over the random binning ensemble.

Broadly speaking, the rate constraints are imposed to ensure that the encoder-decoder pair aims to satisfy the following, with the penalties for finite blocklengths: i) (C, F) are almost independent of X^n; ii) (C, F, S) almost recover U^n; and iii) F is almost independent of Y^n. To impose constraints ensuring near independence, we apply [37, Theorem 1]. Similarly, to impose reliable sequence reconstruction constraints, we apply [37, Theorem 2]. These steps yield rate constraints on $(\widetilde{R}, R, R_0)$, derived by applying Berry-Esseen Theorem such that total variation distances between the target and observed probability distributions are bounded by a fixed value. This analysis corresponds to Protocol B, a channel coding problem dual to our problem with extra randomness F. Moreover, the proof of the realism constraint (4) follows by applying the soft covering lemma [12, Lemma IV.1] as in the achievability proof of [23, Theorem 2].

To eliminate the extra randomness F such that $\widetilde{R}$ is also eliminated from the rate constraints, we show that a fixed realization $F = f$ can be agreed upon publicly by the encoder and decoder, which follows by applying arguments similar to in [32], [38]. Finally, by selecting the free parameters similar to the choices in [32, Eq. (36)], we obtain the results in (10) and (11).

The cardinality bound on the auxiliary random variable U follows by using the support lemma [4, Lemma 15.4]. We preserve the probability distribution P_{XY} by using $(|\mathcal{X}|^2 - 1)$ continuous real-valued functions, since we have $\mathcal{Y} = \mathcal{X}$. We must preserve two more expressions that are the lower bounds in (10) and (11). Therefore, we can limit the cardinality $|\mathcal{U}|$ of U to $|\mathcal{U}| \leq |\mathcal{X}|^2 + 1$. □

B. Proof Sketch of Theorem 2

Proof Sketch. The achievability proof follows similarly to the achievability proof of Theorem 1 given in Section V-A. However, we here need to demonstrate that there exist non-asymptotic random binning schemes that simultaneously meet the following constraints

$$||P_{Y^n} - Q_X^n||_{\text{TV}} \leq \theta\epsilon_{\text{r}}, \tag{23}$$

$$||P_{SY^n} - P_S P_{Y^n}||_{\text{TV}} \leq (1-\theta)\epsilon_{\text{sec}} \tag{24}$$

for any $\theta \in [0,1]$; see also [37, Theorem 4]. Broadly speaking, we use a similar random code construction with the additional strong secrecy constraint imposed such that, rather than imposing that F is almost independent of Y^n, we impose that (S, F) are almost independent of Y^n. Applying steps similar to those in Section V-A, we obtain (14) and (15).

The cardinality bound follows similarly to Section V-A by preserving P_{XY} and the lower bounds in (14) and (15). □

VI. Conclusion and Impact

This paper addressed the RDP problem, extending classical rate-distortion trade-offs by incorporating perceptual quality and secrecy constraints within the RDFC framework. Achievable non-asymptotic regions were established, and their asymptotic counterparts were analyzed. These results provided foundational insights into low-latency, high-fidelity, and secure image compression with realistic outputs. Although the i.i.d. assumption for input sequences was idealistic, it offered a basis for insightful theoretical benchmarks. Extending these results to non-i.i.d. cases using methods such as information spectrum techniques, as used in [44], is possible, but it results in cumbersome expressions that limit their usefulness. Thus, we focused on i.i.d. models to achieve a balance between theoretical rigor and actionable insights.

The technological impact of these findings spans multiple domains. The derived bounds and trade-offs inform the design of systems capable of secure and efficient image compression, addressing critical needs in privacy-sensitive applications. By ensuring security against threats from powerful quantum computers, these results align with the ongoing transition to quantum-safe communications. Future work could extend the current framework to address noisy transmissions. Such efforts would deepen theoretical understanding of joint RDP-channel coding methods, while enhancing the practical applicability of RDP methods in emerging deep learning-based systems.

ACKNOWLEDGMENT

This work of O. Günlü was supported by the ZENITH Research and Leadership Career Development Fund under Grant ID23.01, ELLIIT funding endowed by the Swedish government, and the Swedish Foundation for Strategic Research (SSF) under Grant ID24-0087. The work of H. V. Poor was supported in part by the U.S National Science Foundation under Grant ECCS-2335876.

REFERENCES

[1] D. Gündüz *et al.*, "Timely and massive communication in 6G: Pragmatics, learning, and inference," *IEEE BITS Inf. Theory Mag.*, Oct. 2023.

[2] ——, "Beyond transmitting bits: Context, semantics, and task-oriented communications," *IEEE J. Sel. Areas Commun. (JSAC)*, Nov. 2022.

[3] R. Dobrushin and B. Tsybakov, "Information transmission with additional noise," *IRE Trans. Inf. Theory (T-IT)*, vol. 8, no. 5, pp. 293–304, Sep. 1962.

[4] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2011.

[5] T. Berger, "Rate-distortion theory," *Wiley Encyclopedia of Telecommunications*, 2003.

[6] O. Günlü, "Randomized distributed function computation with semantic communications: Applications to privacy," in *Proc. IEEE Int. Workshop Inf. Forensics Security (WIFS)*, Dec. 2024, pp. 1–6.

[7] M. Havasi, R. Peharz, and J. M. Hernández-Lobato, "Minimal random code learning: Getting bits back from compressed model parameters," in *Proc. Int. Conf. Learn. Representations (ICLR)*, May 2019.

[8] G. Flamich, M. Havasi, and J. M. Hernández-Lobato, "Compressing images by encoding their latent representations with relative entropy coding," *Adv. Neural Inf. Process. Sys. (NeurIPS)*, Dec. 2020.

[9] B. Isik, F. Pase, D. Gündüz, S. Koyejo, T. Weissman, and M. Zorzi, "Adaptive compression in federated learning via side information," in *Proc. Int. Conf. Artificial Intelligence Statistics (AISTATS)*, May 2024.

[10] A. Shah, W.-N. Chen, J. Balle, P. Kairouz, and L. Theis, "Optimal compression of locally differentially private mechanisms," in *Proc. Int. Conf. Artificial Intelligence Statistics (AISTATS)*, Mar. 2022.

[11] M. Hegazy, R. Leluc, C. T. Li, and A. Dieuleveut, "Compression with exact error distribution for federated learning," in *Proc. Int. Conf. Artificial Intelligence Statistics (AISTATS)*, May 2024.

[12] P. Cuff, "Distributed channel synthesis," *IEEE Trans. Inf. Theory (T-IT)*, Aug. 2013.

[13] D. Bergström and O. Günlü, "Deep randomized distributed function computation (DeepRDFC): Neural distributed channel simulation," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Michigan, June 2025, submitted.

[14] K. Sayood, *Introduction to Data Compression*, 4th ed. Morgan Kaufmann, 2017.

[15] Y. Blau and T. Michaeli, "Rethinking lossy compression: The rate-distortion-perception tradeoff," in *Proc. Int. Conf. Mach. Learn. (ICML)*, June 2019, pp. 675–685.

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

[17] J. Chen, L. Yu, J. Wang, W. Shi, Y. Ge, and W. Tong, "On the rate-distortion-perception function," *IEEE J. Sel. Areas Inf. Theory (JSAIT)*, vol. 3, no. 4, pp. 664–673, Dec. 2022.

[18] G. Zhang, J. Qian, J. Chen, and A. Khisti, "Universal rate-distortion-perception representations for lossy compression," *Adv. Neural Inf. Process. Sys. (NeurIPS)*, vol. 34, pp. 11 517–11 529, 2021.

[19] L. Theis and A. B. Wagner, "A coding theorem for the rate-distortion-perception function," in *Neural Compression: From Information Theory to Applications–Workshop@ ICLR 2021*, May.

[20] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Adv. Neural Inf. Process. Sys. (NeurIPS)*, Dec. 2014.

[21] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Int. Conf. Mach. Learning (ICML)*, Aug. 2017, pp. 214–223.

[22] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of Wasserstein GANs," *Adv. Neural Inf. Process. Sys. (NeurIPS)*, vol. 30, Dec. 2017.

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

[24] Y. Hamdi, A. B. Wagner, and D. Gündüz, "The rate-distortion-perception trade-off: The role of private randomness," *arXiv preprint arXiv:2404.01111*, Apr. 2024.

[25] O. Günlü, M. Bloch, and R. F. Schaefer, "Secure multi-function computation with private remote sources," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Melbourne, Australia, July 2021, pp. 1403–1408.

[26] "Post-Quantum Cryptography," [Online]. Available: https://csrc.nist.gov/projects/post-quantum-cryptography, Accessed: Dec. 21, 2024.

[27] S. A. Ameli Kalkhoran, M. Letafati, E. Erdemir, B. H. Khalaj, H. Behroozi, and D. Gündüz, "Secure Deep-JSCC against multiple eavesdroppers," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2023, pp. 3433–3438.

[28] O. Günlü and R. F. Schaefer, "Low-complexity and reliable transforms for physical unclonable functions," in *Proc. IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)*, Barcelona, Spain, May 2020, pp. 2807–2811.

[29] A. D. Wyner, "The wire-tap channel," *Bell Labs Tech. J.*, vol. 54, no. 8, pp. 1355–1387, Oct. 1975.

[30] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Trans. Inf. Theory (T-IT)*, vol. 56, no. 5, pp. 2307–2359, May 2010.

[31] O. Günlü, M. Bloch, R. F. Schaefer, and A. Yener, "Nonasymptotic performance limits of low-latency secure integrated sensing and communication systems," in *Proc. IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)*, Seoul, Republic of Korea, April 2024, pp. 12 971–12 975.

[32] G. Cervia, T. J. Oechtering, and M. Skoglund, "(ϵ, n) fixed-length strong coordination capacity," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Kanazawa, Japan, Oct. 2021, pp. 1–6.

[33] N. Saldi, T. Linder, and S. Yüksel, "Output constrained lossy source coding with limited common randomness," *IEEE Trans. Inf. Theory (T-IT)*, vol. 61, no. 9, pp. 4984–4998, Sep. 2015.

[34] G. Åhlgren and O. Günlü, "Secure rate-distortion-perception trade-off over channels: A randomized distributed function computation (RDFC) application," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Michigan, June 2025, submitted.

[35] V. Kostina and S. Verdú, "Fixed-length lossy compression in the finite blocklength regime," *IEEE Trans. Inf. Theory*, vol. 58, no. 6, pp. 3309–3338, June 2012.

[36] M. H. Yassaee, M. R. Aref, and A. Gohari, "A technique for deriving one-shot achievability results in network information theory," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Istanbul, Turkey, July 2013, pp. 1287–1291.

[37] ——, "Non-asymptotic output statistics of random binning and its applications," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Istanbul, Turkey, July 2013, pp. 1849–1853.

[38] ——, "Achievability proof via output statistics of random binning," *IEEE Trans. Inf. Theory (T-IT)*, vol. 60, no. 11, pp. 6760–6786, Nov. 2014.

[39] J. M. Renes and R. Renner, "Noisy channel coding via privacy amplification and information reconciliation," *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7377–7385, Nov. 2011.

[40] V. Y. F. Tan, "Achievable second-order coding rates for the wiretap channel," in *Proc. IEEE Int. Conf. Commun. Syst. (ICCS)*, Nov. 2012, pp. 65–69.

[41] T. Welling, O. Günlü, and A. Yener, "Low-latency secure integrated sensing and communication with transmitter actions," in *Proc. IEEE Int. Workshop Sign. Process. Adv. Wireless Commun. (SPAWC)*, Lucca, Italy, September 2024, pp. 351–355.

[42] A. E. Gamal and Y.-H. Kim, *Network Information Theory*. Cambridge, U.K.: Cambridge University Press, 2011.

[43] G. Kramer, *Multi-User Information Theory*. Munich, Germany: Techn. Univ. Munich, 2018.

[44] S. Verdú and T. S. Han, "A general formula for channel capacity," *IEEE Trans. Inf. Theory (T-IT)*, vol. 40, no. 4, pp. 1147–1157, July 1994.