Distributed Hypothesis Testing Under Privacy Constraints

Sreejith Sreekumar[†], Deniz Gündüz[†] and Asaf Cohen[‡]
[†]Dept. of Electrical and Electronic Engineering, Imperial College London
[‡]Dept. of Communication Systems Engineering, Ben-Gurion University of the Negev
{s.sreekumar15,d.gunduz}@imperial.ac.uk, coasaf@bgu.ac.il

***Abstract*—A distributed binary hypothesis testing problem involving two parties, a remote observer and a detector, is studied. The remote observer has access to a discrete memoryless source, and communicates its observations to the detector via a rate-limited noiseless channel. The detector tests for the independence of its own observations with that of the observer, conditioned on some additional side information. While the goal is to maximize the type 2 error exponent of the test for a given type 1 error probability constraint, it is also desired to keep a private part, which is correlated with the observer's observations, as oblivious to the detector as possible. Considering equivocation and average distortion as the metrics of privacy at the detector, a tight single-letter characterization of the rate-error exponent-equivocation and rate-error exponent-distortion tradeoff is obtained.**

I. Introduction

Data inference and privacy are often contradicting objectives. In a distributed learning system, the performance of the learning algorithm depends critically on the communication between the agents involved. Typically, in multi-agent systems, each node provides information about their data to a remote decision maker, whose decisions determine the utility achieved by the system. On the other hand, privacy of the underlying data is increasingly becoming important due to the availability of powerful data-mining and machine learning algorithms. Thus, it is critical that the agents only reveal information relevant for obtaining the desired utility so that maximum possible privacy is retained for the sensitive information.

In distributed learning applications, the goal is typically to learn the joint probability distribution of the data available at different locations. Usually, there is some prior knowledge about the joint distribution, for example, that it belongs to a certain set of known probability distributions. In such a scenario, the detector, which tries to infer the joint distribution, applies hypothesis testing (HT) to decide on the joint distribution of the data based on its own observations and the data that it receives from other nodes. While the performance of the hypothesis test depends on the data transferred between the remote nodes and the detector, more data transferred may also result in reduced privacy. With the efficient data mining and machine learning algorithms available today, the detector can illegitimately infer some unintended private information from the data provided to it exclusively for HT purposes. Such threats are becoming increasingly imminent as large amounts of seemingly irrelevant sensitive data are released without proper anonymization, such as in medical research, social networks, online shopping, etc. Therefore, there is an inherent trade-off between the utility acquired by sharing personal data and the associated privacy leakage.

In this paper, we study a special case of the problem of distributed HT known as the *testing against conditional independence* (TACI) problem, under a privacy constraint. In TACI, the detector tests whether its own observation is conditionally independent of data at a remote observer, conditioned on an additional side information Z, available at the detector. Distributed HT without any privacy constraints has been studied extensively from an information theoretic perspective in the past, although many open problems remain. Testing against independence, e.g., no side information Z, is studied in [1] and [2], where the best achievable type 2 error exponent (T2EE) is established, in addition to other fundamental results for the general HT problem. The TACI is first studied in [3], where the optimality of a random binning based encoding scheme is shown. Various multi-terminal scenarios have been studied in [4], [5] and [6]. Recently, the optimal T2EE for TACI over a noisy channel is established in [7].

HT under a mutual information and maximal leakage privacy constraint has been studied in [8] and [9], respectively, where the encoder uses a *memoryless privacy mechanism* to convey a noisy version of its observed data to the detector. The detector performs HT on this noisy data and the optimal privacy mechanism that maximizes the T2EE is studied. Several other privacy measures have been considered in the literature, such as k-anonymity, differential privacy etc.; see [10] for a detailed survey. Among these, *equivocation* (or, equivalently, the mutual information between the data and the revealed information) is most commonly used to quantify the average privacy leakage from an information theoretic perspective [11]. A more general rate-distortion approach to privacy was first explored by Yamamoto for the case of a noiseless channel with a rate constraint, where, in addition to a distortion constraint for the legitimate data, a minimum distortion requirement is enforced for the private data [12].

In the sequel, we study TACI under privacy constraints with average distortion and equivocation as the metrics of privacy. To contrast with [8] and [9], we do not assume memoryless coding mechanisms at the encoder. More specifically, the output of the encoder is allowed to depend on the entire sequence of observed samples, rather than a single sample. Also, while [8] and [9] are concerned with general HT in a point to point

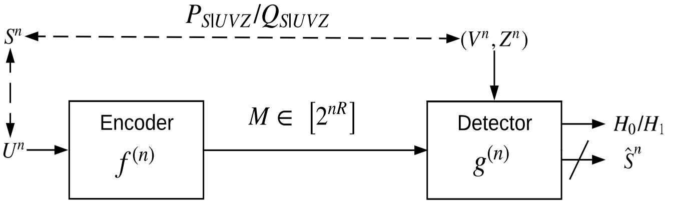

Fig. 1: HT under a privacy constraint.

setting, i.e., no side information at the detector, our focus is TACI in a distributed setting. Our main contribution is to establish a single-letter characterization of the complete rate-T2EE-distortion and rate-T2EE-equivocation trade-off.

A. Notations

We denote random variables (r.v.'s) and their realizations by upper and lower case letters (e.g., X and x), respectively. Sets are denoted by calligraphic letters, e.g., the alphabet of a r.v. X is denoted by $\mathcal{X}$. Sequence of r.v.'s $(X_1, \ldots, X_n)$ is denoted by X^n. Given distributions P_X and $P_{Y|X}$, the marginal distribution P_Y of Y (induced by P_X) is denoted by $P_X \circ P_{Y|X}$, where $P_Y(y) = (P_X \circ P_{Y|X})(y) = \sum_{x \in \mathcal{X}} P_X(x) P_{Y|X}(y|x)$. Following the notation in [13], $T^m_{[P_X]_\delta}$ (or $T^m_{[X]_\delta}$ or simply T^m_δ when there is no ambiguity) denotes the set of P_X-typical sequences of length m. $\mathbb{1}$ denotes the indicator function. $X - Y - Z$ denotes a Markov chain between r.v.'s X, Y and Z. $\xrightarrow{(n)}$ denotes asymptotic limit with respect to n, e.g., $a_n \xrightarrow{(n)} 0$ means the sequence a_n tends to zero asymptotically with n. $\mathbb{P}(\mathcal{E})$ denotes the probability of event $\mathcal{E}$. For positive real m, we define $[m] \triangleq \{1, \ldots, \lceil m \rceil\}$. For an arbitrary set $\mathcal{A}$, we denote its complement by $\mathcal{A}^c$ and for $\mathcal{A} \subseteq \mathbb{R}^n$, we denote its closure by $cl(\mathcal{A})$ (with respect to the Euclidean metric).

II. Problem formulation

Consider the HT setup illustrated in Fig. 1. The encoder (observer) observes a discrete memoryless source U^n independent and identically distributed (i.i.d.) according to $\prod_{i=1}^n P_U$, and sends an index $M \triangleq f^{(n)}(U^n)$ to the detector over an error-free channel using some encoding function (possibly stochastic) $f^{(n)} : \mathcal{U}^n \rightarrow [2^{nR}]$, where R is the rate of the error-free communication channel available from the encoder to the detector. Given its own i.i.d. observation V^n and side information Z^n, the detector performs TACI with null hypothesis

$$H_0 : P_{UVZ}$$

and alternate hypothesis

$$H_1 : Q_{UVZ} = P_{U|Z} P_{V|Z} P_Z$$

on the joint distribution of U, V and Z. Let $\mathcal{A} \subseteq [2^{nR}] \times \mathcal{Z}^n \times \mathcal{V}^n$ (resp. $\mathcal{A}^c$) denote the acceptance region for H_0 (resp. H_1). The decision rule of the detector is given by $g^{(n)}(m, z^n, v^n) = \mathbb{1}\left((m, z^n, v^n) \in \mathcal{A}^c\right)$, where 0 (resp. 1) denotes H_0 (resp. H_1). Let $\bar{\alpha}\left(f^{(n)}, g^{(n)}\right) \triangleq P_{MZ^nV^n}(\mathcal{A}^c)$ $\big($resp. $\bar{\beta}\left(f^{(n)}, g^{(n)}\right) \triangleq P_{MZ^n} \times P_{V^n|Z^n}(\mathcal{A})$ $\big)$ denote the type 1 (resp. type 2) error probability for an $(f^{(n)}, g^{(n)})$ pair. For a given type 1 error probability constraint ϵ, we define the minimum type 2 error probability over all possible decoders as

$$\beta\left(f^{(n)}, \epsilon\right) \triangleq \inf_{g^{(n)}} \bar{\beta}\left(f^{(n)},\ g^{(n)}\right), \tag{1}$$
$$\text{such that } \bar{\alpha}\left(f^{(n)},\ g^{(n)}\right) \leq \epsilon.$$

The performance of TACI is measured by the T2EE achieved by the test for a given constraint ϵ on the type 1 error probability, i.e., $\liminf_{n \rightarrow \infty} -\frac{1}{n} \log\left(\beta(f^{(n)}, \epsilon)\right)$.

The detector is also curious about a latent r.v. S^n that is correlated with the data observed by the encoder, U^n. S^n is referred to as the *private* part of U^n, and is i.i.d. according to the joint distribution $P_{SUVZ} = P_{S|UVZ} P_{UVZ}$ and $Q_{SUVZ} = Q_{S|UVZ} P_{U|Z} P_{V|Z} P_Z$ under the null and alternate hypothesis, respectively, where $P_{S|UVZ}$ and $Q_{S|UVZ}$ denotes two arbitrary conditional probability distributions. The observer desires to keep the private part as concealed as possible from the detector. We consider two metrics of privacy for S^n at the detector. The first metric is equivocation defined as $\frac{1}{n} H(S^n | M, V^n, Z^n)$. The second one is the average distortion between S^n and its reconstruction $\hat{S}^n$ at the detector, measured according to an arbitrary bounded additive distortion measure $d : \mathcal{S} \times \hat{\mathcal{S}} \rightarrow [0, D_m]$ with multi-letter distortion defined as

$$d(s^n, \hat{s}^n) \triangleq \frac{1}{n} \sum_{i=1}^{n} d(s_i, \hat{s}_i). \tag{2}$$

The goal is to ensure that the T2EE for HT is maximized, while satisfying the constraints on the type 1 error probability ϵ and the privacy of S^n. In the sequel, we state the results characterizing the trade-off between the rate, T2EE (henceforth referred to as the error exponent), and privacy achieved in the above setting for the case $\epsilon \rightarrow 0$.

Let $H, \hat{H} \in \{H_0, H_1\}$ denote the r.v.'s corresponding to the true hypothesis and the output of the HT, respectively.

Definition 1. *For a given type 1 error probability constraint ϵ, a rate - error exponent - distortion tuple $(R, \kappa, \Delta_0, \Delta_1)$ is* achievable, *if there exists a sequence of encoding and decoding functions $f^{(n)} : \mathcal{U}^n \rightarrow [2^{nR}]$ and $g^{(n)} : [2^{nR}] \times \mathcal{Z}^n \times \mathcal{V}^n \rightarrow \{0, 1\}$ such that*

$$\limsup_{n \rightarrow \infty} \frac{\log\left(\beta(f^{(n)}, \epsilon)\right)}{n} \leq -\kappa, \text{ and} \tag{3}$$

$$\liminf_{n \rightarrow \infty} \inf_{g_r^{(n)}} \mathbb{E}\left[d\left(S^n, \hat{S}^n\right) | H = H_i\right] \geq \Delta_i,\ i = 0, 1, \tag{4}$$

where $g_r^{(n)} : [2^{nR}] \times \mathcal{Z}^n \times \mathcal{V}^n \rightarrow \hat{\mathcal{S}}^n$ denotes an arbitrary (possibly stochastic) mapping. The rate - error exponent - distortion region $\mathcal{R}_d(\epsilon)$ is the closure of the set of all such achievable $(R, \kappa, \Delta_0, \Delta_1)$ tuples for a given ϵ.

Definition 2. *For a given type 1 error probability constraint ϵ, a rate-error exponent-equivocation $(R, \kappa, \Omega_0, \Omega_1)$ tuple is* achievable, *if there exists a sequence of encoding and decoding functions $f^{(n)} : \mathcal{U}^n \rightarrow [2^{nR}]$ and $g_h^{(n)} : [2^{nR}] \times \mathcal{Z}^n \times \mathcal{V}^n \rightarrow$*

$\{0,1\}$ *such that* (3) *is satisfied and*

$$\liminf_{n\to\infty}\frac{1}{n}H(S^n|M,V^n,Z^n,H=H_i)\geq\Omega_i,\ i=0,1. \quad (5)$$

The rate-error exponent-equivocation region $\mathcal{R}_e(\epsilon)$ is the closure of the set of all achievable $(R,\kappa,\Omega_0,\Omega_1)$ tuples for a given ϵ.

The goal of the paper is to provide a single-letter characterization for $\mathcal{R}_e(\epsilon)$ and $\mathcal{R}_d(\epsilon)$ in the regime of vanishing type 1 error probability, i.e., $\epsilon\to 0$, which we denote by $\mathcal{R}_e$ and $\mathcal{R}_d$, respectively.

Let $\mathcal{R}_d^{(n)}$ denote the set of $(R,\kappa,\Delta_0,\Delta_1)$-tuples such that there exists $f^{(n)}:\mathcal{U}^n\to\left[2^{nR}\right]$ satisfying[1]

$$\kappa\leq\frac{1}{n}I(M;V^n|Z^n), \quad (6)$$

$$\inf_{g_r^{(n)}}\mathbb{E}\left[d\left(S^n,\hat{S}^n\right)|H=H_i\right]\geq\Delta_i,\ i=0,1. \quad (7)$$

$$(V^n,Z^n,S^n)-U^n-M,\ M=f^{(n)}(U^n) \quad (8)$$

Similarly, let $\mathcal{R}_e^{(n)}=\{(R,\kappa,\Omega_0,\Omega_1)\}$ such that there exists $f^{(n)}:\mathcal{U}^n\to\left[2^{nR}\right]$ satisfying (6), (8) and

$$\frac{1}{n}H(S^n|M,V^n,Z^n,H=H_i)\geq\Omega_i,\ i=0,1. \quad (9)$$

The next theorem provides an $n-$letter characterization of $\mathcal{R}_d$ and $\mathcal{R}_e$ in terms of $\mathcal{R}_d^{(n)}$ and $\mathcal{R}_e^{(n)}$, respectively.

Theorem 3.

$$\mathcal{R}_d=cl\left(\cup_n\mathcal{R}_d^{(n)}\right), \quad (10)$$

$$\mathcal{R}_e=cl\left(\cup_n\mathcal{R}_e^{(n)}\right). \quad (11)$$

Due to space constraints, the proof of the theorem is omitted here. The details can be found in an extended version of the paper. In the next section, we introduce the one-helper lossless source coding problem under a privacy constraint, which will be instrumental in obtaining a single-letter characterization of $\mathcal{R}_d$ and $\mathcal{R}_e$.

III. One-helper lossless source coding Problem under a Privacy constraint

Consider the setup shown in Fig. 2, which we refer to as the one-helper lossless source coding problem under a privacy constraint. In this problem, the main encoder $f_v^{(n)}$ (resp. helper encoder $f^{(n)}$) sends the message $\tilde{M}$ (resp. M) based on its observation V^n (resp. U^n) to the legitimate decoder $g_v^{(n)}$ through a noiseless channel with rate constraint R_v (resp. R). The goal of the legitimate decoder g_v^n is to reconstruct V^n losslessly using the received indices M and $\tilde{M}$ as well as its side information Z^n. This is a source coding with coded side information problem, studied in [14]. However, in our case, there is an additional sequence S^n and an adversary decoder $g_r^{(n)}$ which has access to (M,V^n,Z^n). The goal is to keep S^n private from the adversary decoder such that (4) (resp. (5)) is satisfied when average distortion (resp. equivocation) is the measure of privacy. Note that the adversary decoder has access to all the information that the legitimate decoder has. Hence, protecting S^n cannot depend on coding techniques that are based on the adversary partially missing data (or having a noisier channel), as is common in many physical layer security related works. We measure the privacy of S^n at the adversary decoder under two cases, namely, when the joint distribution of the r.v.'s (S^n,U^n,V^n,Z^n) is (i) $\prod_{i=1}^n P_{SUVZ}$ and (ii) $\prod_{i=1}^n Q_{SUVZ}$. The pair of equivocation and average distortion tuples simultaneously achievable in these two cases are of interest.

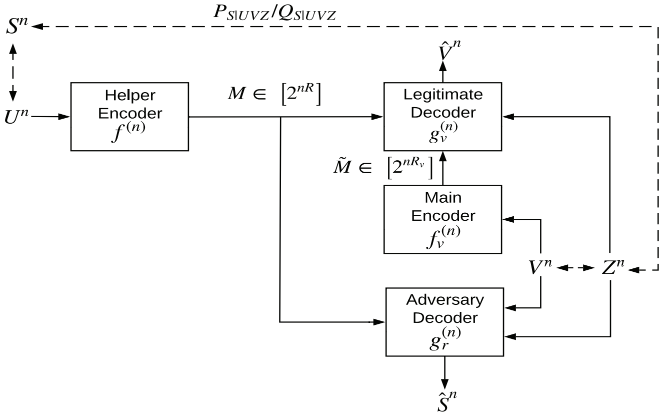

Fig. 2: Source coding problem in the presence of a helper under a privacy constraint.

Definition 4. *Given a distortion measure $d:\mathcal{S}^n\times\hat{\mathcal{S}}^n\to[0,D_m]$, a rate-distortion tuple $(R,R_v,\Delta_0,\Delta_1)$ is achievable if there exists a sequence of encoding functions $f^{(n)}:\mathcal{U}^n\to[2^{nR}]$, $f_v^{(n)}:\mathcal{V}^n\to[2^{nR_v}]$ and decoding functions $g_v^{(n)}:[2^{nR}]\times[2^{nR_v}]\times\mathcal{Z}^n\to\hat{V}^n$ such that*

$$\limsup_{n\to\infty}\mathbb{P}(V^n\neq\hat{V}^n)=0 \quad (12)$$

and (4) *are satisfied. Let $\hat{\mathcal{R}}_d$ denote the closure of all achievable $(R,R_v,\Delta_0,\Delta_1)$ tuples.*

Definition 5. *A rate-equivocation tuple $(R,R_v,\Omega_0,\Omega_1)$ is achievable if there exists a sequence of encoding functions $f^{(n)}:\mathcal{U}^n\to[2^{nR}]$, $f_v^{(n)}:\mathcal{V}^n\to[2^{nR_v}]$ and decoding functions $g_v^{(n)}:[2^{nR}]\times[2^{nR_v}]\times\mathcal{Z}^n\to\hat{V}^n$ such that* (5) *and* (12) *are satisfied. Let $\hat{\mathcal{R}}_e$ denote the closure of all achievable $(R,R_v,\Omega_0,\Omega_1)$ tuples.*

The next theorem provides a multi-letter characterization of $\hat{\mathcal{R}}_d$ and $\hat{\mathcal{R}}_e$.

Theorem 6. *Let $\hat{\mathcal{R}}_d^{(n)}$ denote the set of $(R,R_v,\Delta_0,\Delta_1)$ tuples such that* (7), (8) *and*

$$R_v\geq\frac{1}{n}H(V^n|M,Z^n). \quad (13)$$

are satisfied. Similarly, let $\hat{\mathcal{R}}_e^{(n)}$ denote the set of

[1] The mutual information in (6) is computed with respect to the joint distribution induced under H_0, and this will also be the case in the rest of the paper unless specified otherwise.

$(R, R_v, \Omega_0, \Omega_1)$ *tuples such that* (8), (9) *and* (13) *are satisfied. Then,*

$$\hat{\mathcal{R}}_d = cl\left(\cup_n \hat{\mathcal{R}}_d^{(n)}\right), \tag{14}$$

$$\hat{\mathcal{R}}_e = cl\left(\cup_n \hat{\mathcal{R}}_e^{(n)}\right). \tag{15}$$

The proof of this theorem is omitted here due to space constraints and is available in an extended version of the paper.

Noting that $I(M; V^n|Z^n) = nH(V|Z) - H(V^n|Z^n, M)$, it follows from Theorems 3 and 6 that,

$$\begin{aligned}(R, \kappa, \Delta_0, \Delta_1) \in \mathcal{R}_d \Leftrightarrow \\ (R, H(V|Z) - \kappa, \Delta_0, \Delta_1) \in \hat{\mathcal{R}}_d\end{aligned} \tag{16}$$

An equivalence similar to (16) also holds between $\mathcal{R}_e$ and $\hat{\mathcal{R}}_e$ with Δ_i replaced by Ω_i, $i = 0, 1$. In Section IV, we obtain a single-letter characterization of $\mathcal{R}_d$ and $\mathcal{R}_e$ by exploiting these equivalences.

IV. Main Results

The main results of the paper are presented in this section.

Theorem 7. $(R, \kappa, \Omega_0, \Omega_1) \in \mathcal{R}_e$ *if and only if there exists an auxiliary r.v.* W *such that the Markov chain* $(Z, V, S) - U - W$ *is satisfied and*

$$R \geq I(W; U|Z) \tag{17}$$

$$\kappa \leq I(W; V|Z) \tag{18}$$

$$\Omega_i \leq H_{P^{(i)}}(S|W, Z, V), \ i = 0, 1 \tag{19}$$

where the joint distribution $P^{(0)} = P_{SUVZ} P_{W|U}$ *and* $P^{(1)} = Q_{SUVZ} P_{W|U}$.

Proof: We show that $(R, R_v, \Omega_0, \Omega_1) \in \hat{\mathcal{R}}_e$ if and only if there exists an auxiliary r.v. W such that (17), (19) and

$$R_v \geq H(V|W, Z). \tag{20}$$

The result then follows from the equivalence mentioned above.

Achievability: Fix a conditional probability distribution $P_{W|U}$.

Codebook of the helper encoder: Generate codewords $W^n(m, m')$, $m \in [2^{nR}]$, $m' \in [2^{nR'}]$ drawn independently according to distribution $\prod_{i=1}^n P_W$, where $P_W = P_U \circ P_{W|U}$. Denote this codebook by $\mathcal{C}_u^n$.

Codebook of the main encoder: This codebook is generated by performing uniform random binning of the V^n sequences, i.e., an index $\tilde{M}$ is assigned to each sequence v^n uniformly at random from the set $[2^{nR_v}]$. Denote this codebook by $\mathcal{C}_v^n$. The two codebooks are revealed to all the encoders and decoders.

Encoding: The helper encoder uses joint-typicality[2] encoding, i.e., it first looks for a unique (M, M') pair such that $(u^n, W^n(M, M')) \in T^n_{[UW]_\delta}$, $\delta > 0$, where T^n_δ denotes the $\delta-$ typical set as defined in [13]. If successful, M is transmitted to the decoder; otherwise, it transmits a message chosen uniformly at random from the set $[2^{nR}]$. The message M' is not transmitted, but is intended to be recovered by the decoder using its side information Z^n. The encoder of source V^n sends the bin-index $\tilde{M} = f_v^{(n)}(V^n)$ via its channel.

Decoding: The decoder first looks for a unique index $\hat{M}'$ such that $(W^n(M, \hat{M}'), Z^n) \in T^n_{[WZ]_{\delta'}}$, $\delta' > \delta > 0$. If successful, it then checks for the unique sequence $\tilde{V}^n$ in bin $\tilde{M}$ such that $(\tilde{V}^n, W^n(M, \hat{M}'), Z^n) \in T^n_{[VWZ]_{\delta''}}$, $\delta'' > \delta' > 0$. If this is also successful, it sets the estimate as $\hat{V}^n = \tilde{V}^n$; otherwise, a random sequence from set $\mathcal{V}^n$ is chosen as the estimate. It can be shown that $\mathbb{P}(\mathcal{E}) \triangleq \mathbb{P}(V^n \neq \hat{V}^n)$ can be made to decay to zero with n, provided that (17) and (20) are satisfied. It can also be proved that an equivocation of Ω_i is achievable for S^n at the adversary under hypothesis H_i, provided (19) is satisfied. The details are available in an extended version of the paper.

Converse: Let Q be a r.v. uniformly distributed over $[n]$ and independent of all the other r.v.'s (U^n, V^n, Z^n, M). Define $U = U_Q$, $Z = Z_Q$, $V = V_Q$, $S = S_Q$ and auxiliary r.v. $W \triangleq (W_Q, Q)$, where $W_i \triangleq (M, V^{i-1}, Z^{i-1}, Z^n_{i+1})$, $i \in [n]$. Note that $(Z, V) - U - W$. Then, for any $\epsilon' > 0$ and sufficiently large n, we have

$$\begin{aligned}
& n(R + \epsilon') \geq H(M) \geq H(M|Z^n) \geq I(M; U^n|Z^n) \\
& = \sum_{i=1}^n I(M; U_i|U^{i-1}, Z^n) \\
& = \sum_{i=1}^n I(M, U^{i-1}, Z^{i-1}, Z^n_{i+1}; U_i|Z_i) && (21) \\
& = \sum_{i=1}^n I(M, U^{i-1}, Z^{i-1}, Z^n_{i+1}, V^{i-1}; U_i|Z_i) && (22) \\
& \geq \sum_{i=1}^n I(M, Z^{i-1}, Z^n_{i+1}, V^{i-1}; U_i|Z_i) \\
& = \sum_{i=1}^n I(W_i; U_i|Z_i) = nI(W_Q; U_Q|Z_Q, Q) \\
& = nI(W_Q, Q; U_Q|Z_Q) && (23) \\
& = nI(W; U|Z). && (24)
\end{aligned}$$

Here, (21) follows since the sequences (U^n, Z^n) are memoryless; (22) follows from the Markov chain $V^{i-1} - (M, U^{i-1}, Z^{i-1}, Z^n_{i+1}) - U_i$; (23) follows from the fact that Q is independent of all the other r.v.'s.

The equivocation of source S^n can be bounded as follows.

$$\begin{aligned}
& H(S^n|M, V^n, Z^n, H = H_i) \\
& = \sum_{i=1}^n H(S_i|M, S^{i-1}, V^n, Z^n, H = H_i) \\
& \leq \sum_{i=1}^n H(S_i|M, V^{i-1}, V_i, Z^{i-1}, Z_i, Z^n_{i+1}, H = H_i) \\
& = \sum_{i=1}^n H(S_i|W_i, V_i, Z_i, H = H_i) \\
& = nH(S_Q|W_Q, V_Q, Z_Q, Q, H = H_i) \\
& = nH_{P^{(i)}}(S|W, V, Z). && (25)
\end{aligned}$$

Finally, we prove the bound on R_v. First, note that

$$\begin{aligned}
& n(R_v + \epsilon') \geq H(\tilde{M}|M, Z^n) \\
& = H(\tilde{M}|M, Z^n) + H(V^n|\tilde{M}, M, Z^n) - H(V^n|\tilde{M}, M, Z^n) \\
& \geq H(\tilde{M}, V^n|Z^n, M) - \gamma_n'', && (26)
\end{aligned}$$

where $\gamma_n'' \xrightarrow{(n)} 0$. Eqn. (26) follows from Fano's inequality.

[2] The typical sets defined in this paper are with respect to the joint distribution induced under H_0, i.e., $P^{(0)}$.

Defining $\epsilon_n'' \triangleq \epsilon' + \frac{\gamma_n''}{n} \xrightarrow{(n)} \epsilon'$, from (26), we get

$$\begin{aligned} n(R_v + \epsilon_n'') &\geq H(V^n|M, Z^n) + H(\tilde{M}|V^n, Z^n, M) \\ &\geq H(V^n|M, Z^n) = \sum\nolimits_{i=1}^{n} H(V_i|V^{i-1}, M, Z^n) \\ &= \sum\nolimits_{i=1}^{n} H(V_i|Z_i, W_i) = nH(V_Q|Z_Q, W_Q, Q) \\ &= nH(V|Z, W). \qquad (27) \end{aligned}$$

Eqns. (24), (25) and (27), along with the fact that $\hat{\mathcal{R}}_e$ (and $\mathcal{R}_e$) is closed completes the proof of the converse via the equivalence in (16). ■

Next, we state the result for the case when privacy is measured using an arbitrary distortion measure $d(\cdot, \cdot)$.

Theorem 8. $(R, \kappa, \Delta_0, \Delta_1) \in \mathcal{R}_d$ *if and only if there exist an auxiliary r.v.* W *such that*

$$R \geq I(W; U|Z) \qquad (28)$$

$$\kappa \leq I(W; V|Z) \qquad (29)$$

$$\Delta_i \leq \min_{\phi(\cdot,\cdot,\cdot)} \mathbb{E}_i \left[d\left(S, \phi(W, V, Z)\right)\right], \ i = 0, 1 \qquad (30)$$

for some deterministic function $\phi : \mathcal{W} \times \mathcal{V} \times \mathcal{Z} \to \hat{\mathcal{S}}$. *Here,* $\mathbb{E}_i$ *denotes expectation under* $P^{(i)}$ *defined in Theorem 7.*

Proof: Similarly to Theorem 7, we show that $(R, R_v, \Delta_0, \Delta_1) \in \hat{\mathcal{R}}_d$ if and only if there exists an auxiliary r.v. W such that (20), (28) and (30) are satisfied. The result then follows from the equivalence in (16). To prove achievability, the same codebook generation as in Theorem 7 is used.

Encoding: The encoder $f^{(n)}$ uses stochastic encoding to choose the messages $(M, M') \in [2^{nR}] \times [2^{nR'}]$ according to the following probability.

$$P_{E_u}(m, m'|u^n) = \frac{\prod_{i=1}^{n} P_{U|W}(u_i|W_i(m, m'))}{\sum_{m,m'} \prod_{i=1}^{n} P_{U|W}(u_i|W_i(m, m')))}.$$

The message M is transmitted over the noiseless link, whereas M' is not, but is intended to be recovered at the decoder using the side- information Z^n. The encoder $f_v^{(n)}$ transmits the bin-index $\tilde{M} \in [2^{nR_v}]$ over its own noiseless link.

Decoding: The decoder first uses maximum-likelihood (ML) decoding to retrieve $\hat{M}'$. It then looks for a unique sequence $\tilde{V}^n$ in the bin with index $\tilde{M}$ such that $(\tilde{V}^n, W^n(M, \hat{M}'), Z^n) \in T_{[VWZ]_{\delta''}}$. If such a sequence exists, it sets $\hat{V}^n = \tilde{V}^n$, else it declares an error.

Due to space constraints, rest of the proof of achievability and converse is omitted and can be found in an extended version of the paper. ■

Remark 9. *It can be shown using standard arguments based on the Fenchel-Eggleston-Carathéodory's Theorem [15] that, considering auxiliary r.v.* W *such that* $|\mathcal{W}| \leq |\mathcal{U}| + 3$ *suffices in Theorems 7 and 8.*

Example. Here we provide an example in which maximum privacy (under alternate hypothesis) can be achieved together with a non-zero error exponent. Let Z be a constant, $\mathcal{S} = \mathcal{U} = \{0, 1, 2, 3\}$, $\mathcal{V} = \{0, 1\}$,

$$P_{SU} = 0.125 * \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \quad P_{V|U} = \begin{bmatrix} 0.4 & 0.6 \\ 0.2 & 0.8 \\ 0.3 & 0.7 \\ 0.1 & 0.9 \end{bmatrix},$$

$P_{SUV} = P_{SU} P_{V|U}$ and $Q_{SUV} = P_{SU} P_V$, where $P_V = P_U \circ P_{V|U}$. Then, $H(U) = H(S) = 2$ bits. If we set $W = U \ mod\ 2$, then we have $I(S; W) = 0$, $I(U; W) = 1$ bit, and $I(V; W) = 0.0393$ bits. Thus, by revealing only W to the detector, it is possible to achieve a positive T2EE while ensuring maximum privacy, i.e., the tuple $(1, 0.0393, I(S; W|V), 2)$ is achievable. Alternatively, $(2, I(V; U), I(S; U|V), 1)$ can be achieved by setting $W = U$.

V. Conclusions

We have studied the TACI problem over a rate-limited noiseless channel under privacy constraints. With equivocation and average distortion as the metrics of privacy, we establish single-letter characterizations of the rate-error exponent-equivocation and rate-error exponent-distortion trade-offs. Extending this problem to the case where general hypothesis testing is considered in place of TACI is an interesting avenue for future research.

References

[1] R. Ahlswede and I. Csiszár, "Hypothesis testing with communication constraints," *IEEE Trans. Inf. Theory*, vol. 32, no. 4, pp. 533–542, Jul. 1986.

[2] T. S. Han, "Hypothesis testing with multiterminal data compression," *IEEE Trans. Inf. Theory*, vol. 33, no. 6, pp. 759–772, Nov. 1987.

[3] M. S. Rahman and A. B. Wagner, "On the optimality of binning for distributed hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 58, no. 10, pp. 6282–6303, Oct. 2012.

[4] M. Wigger and R. Timo, "Testing against independence with multiple decision centers," in *IEEE Int. Conf. on Signal Proc. and Comm.*, Bengaluru, India, Jun. 2016.

[5] G. Katz, P. Piantanida, and M. Debbah, "Distributed binary detection with lossy data compression," *IEEE Trans. Inf. Theory*, vol. 63, no. 8, pp. 5207–5227, Aug. 2017.

[6] W. Zhao and L. Lai, "Dist. testing against independence with multiple terminals," in *52nd Annual Allerton Conference*, IL, USA, Oct. 2014.

[7] S. Sreekumar and D. Gündüz, "Distributed hypothesis testing over noisy channels," in *IEEE Int. Symp. Inf. Theory*, Aachen, Germany, Jun. 2017.

[8] J. Liao, L. Sankar, V. Tan, and F. Calmon, "Hypothesis testing under mutual information privacy constraints in the high privacy regime," *IEEE Trans. on Inf. Forensics and Security*, vol. 13, no. 4, pp. 1058 – 1071, Apr. 2018.

[9] J. Liao, L. Sankar, F. Calmon, and V. Tan, "Hypothesis testing under maximal leakage privacy constraints," in *IEEE Int. Symp. Information Theory (ISIT)*, Aachen, Germany, Jun. 2017.

[10] I. Wagner and D. Eckhoff, "Technical privacy metrics: a systematic survey," *arXiv:1512.00327v1 [cs.CR].*

[11] F. Calmon and N. Fawaz, "Privacy against statistical inference," in *50th Annual Allerton Conference*, IL, USA, Oct.2012.

[12] H. Yamamoto, "A rate-distortion problem for a communication system with a secondary decoder to be hindered," *IEEE Trans. Inf. Theory*, vol. 34, no. 4, pp. 835–842, Jul. 1988.

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

[14] R. Ahlswede and J. Körner, "Source coding with side information and a converse for degraded broadcast channels," *IEEE Trans. Inf. Theory*, vol. 21, no. 6, pp. 629–637, Nov. 1975.

[15] A. E. Gamal and Y.-H. Kim, *Network Information theory*. Cambridge University Press, 2011.