Denoising and Filtering Under the Probability of Excess Loss Criterion

Stephanie Pereira and Tsachy Weissman

Abstract

Subclasses of finite alphabet denoising and filtering (causal denoising) schemes are compared. Performance is measured by the normalized cumulative single letter loss (a.k.a. distortion). We aim to minimize the probability that the normalized cumulative loss exceeds a given threshold. We call this quantity the probability of excess loss. Specifically, we consider a scheme to be optimal if it attains the maximal exponential decay rate of the probability of excess loss. This provides another way of comparing schemes that complements and contrasts previous work which considered only the expected value of the normalized cumulative loss.

In particular, the question of whether the optimal denoiser is symbol-by-symbol for an independent, identically distributed source and a discrete memoryless channel is investigated. For Hamming loss, the optimal denoiser is proven to be symbol-by-symbol. Perhaps somewhat counterintuitively, for a general single letter loss function, the optimal scheme need not be symbol-by-symbol.

The optimal denoiser requires unbounded delay and unbounded look-ahead while symbol-by-symbol schemes mandate zero delay and look-ahead. It is natural to wonder about the effect of limited delay and limited look-ahead. Consequently, finite sliding window denoisers and finite block denoisers are defined. They are shown to perform no better than symbol-by-symbol denoisers.

Finally, the effect of causality is investigated. While it is difficult to characterize the performance of filters with unbounded memory explicitly, it is shown that finite memory filters perform no better than symbol-by-symbol filters.

Index Terms

Denoising, filtering, large deviations, single letter loss, probability of excess loss, symbol-by-symbol, delay, look-ahead, causality, time-invariant, sliding-block, memory.

I. INTRODUCTION

THE denoising and filtering problems have a long history focussed on the continuous alphabet case. Recently, there has been work on the discrete alphabet case (cf. [1], [2]). To our knowledge, only the problem of minimizing *expected* loss has been considered. We study the probability that the loss *exceeds* a particular threshold, first considered by Marton in [3] in the context of lossy source codes. This *excess loss* criterion enables us to design denoisers and filters that have loss less than some target rate with high probability. Further, even if a denoiser/filter has low expected loss, the spread of this loss may be high. The excess loss criterion provides a handle on the spread of the loss. Our work was partially inspired by results in lossy source coding (cf. [3], [4], [5]).

In particular, we analyze the asymptotic excess loss probability by establishing a large deviations principle (LDP) for denoisers and determining the corresponding rate function. Large deviations characterizations have been used as a performance metric both in the information theory and statistics literature (see [3], [6], [4] and [7] and [8], resp.).

The LDP for denoising is a special case of the lossy source coding LDP discussed in [4] and [7]. However, while [4] and [7] are concerned with characterizing the performance of the optimal scheme, the basic question we ask in this work is how different subclasses of schemes compare to the optimal scheme. In other words, how much, if anything, is lost by restricting the class of allowable schemes? There is a clear practical motivation to this question. The subclasses we consider are those that limit the amount of noisy observations that the denoiser ‘sees’. In practice, a denoiser may not have an unbounded horizon so it is important to ascertain whether/when such practical schemes are close to the optimal bound. Furthermore, while the derivations of [4] and [7] are information-theoretic, our

Fig. 1. Denoising/Filtering Setup

Fig. 2. Denoiser

results have more of a large deviations flavor. I.e., while the characterizations in [4] and [7] are given in terms of minimum Kullback-Leibler divergences, in this work we emphasize the Fenchel-Legendre transform representation of the exponents. This representation makes the comparison of the rate functions for different subclasses more transparent and helps us to establish cases of strict suboptimality of symbol-by-symbol and other classes of schemes.

II. SETUP

The setup (see Figure 1) is as follows: a source generates n i.i.d. symbols, $\{X_i\}_{i=1}^n$, that take values in a discrete alphabet $\mathcal{X}$ of finite cardinality. These source symbols X_i pass through a discrete memoryless channel (d.m.c.) $\Pi(z|x)$ to produce Z_i that take values in a discrete alphabet $\mathcal{Z}$ of finite cardinality. Denote the distribution of the Z_i by Q .

A. Denoising

The goal of a denoiser is to estimate $X^n = (X_1, X_2, \dots, X_n)$ from $Z^n = (Z_1, Z_2, \dots, Z_n)$. The vector Z^n is *denoised* to produce the symbol $\hat{X}_i = \hat{x}_i(Z^n)$, for each i , where the *denoising functions* $\hat{x}_i(\cdot)$ are general, deterministic functions of the random vector Z^n with range $\hat{\mathcal{X}} = \{1, 2, \dots, |\hat{\mathcal{X}}|\}$, $|\hat{\mathcal{X}}| < \infty$. We note that while $\hat{x}_i(\cdot)$ is a deterministic function, $\hat{X}_i$ is a random variable. The *denoiser* is the collection of *denoising functions* $(\hat{x}_1(\cdot), \hat{x}_2(\cdot), \dots, \hat{x}_n(\cdot))$ and is denoted by $\hat{x}^n(\cdot)$. We illustrate a general denoiser in Figure 2.

If the denoising functions satisfy $\hat{x}_i(Z^n) = \hat{x}(Z^n, Z_i)$, $1 \leq i \leq n$ for some deterministic function $\hat{x} : \mathcal{Z}^n \times \mathcal{Z} \rightarrow \hat{\mathcal{X}}$, we call the denoiser *time-invariant*.

We refer to a denoiser with denoising functions that depend only on Z_i as a *symbol-by-symbol* denoiser, so that $\hat{X}_i = \hat{x}_i(Z^n) = \hat{x}_i(Z_i)$ for a symbol-by-symbol denoiser. A symbol-by-symbol denoiser is shown in Figure 3. Note that the $\hat{x}_i(\cdot)$ may vary with time (hence the subscript i). Applying the above definition and the definition of a symbol-by-symbol denoiser, we can see that a symbol-by-symbol denoiser is time-invariant if $\hat{x}_i(z_i) = \hat{x}(z_i)$ for some function $x : \mathcal{Z} \rightarrow \hat{\mathcal{X}}$.

We define the *k-finite sliding window denoiser* to allow $\hat{X}_i$ to depend on Z_{i-k}^{i+k} (i.e. $\hat{X}_i = \hat{x}_i(Z_{i-k}^{i+k})$) (see Figure 4). As

Fig. 3. Symbol-by-Symbol Denoiser/Filter

Fig. 4. k -Finite Sliding Window Denoiser Fig. 5. k -Finite Block Denoiser

above, a time-invariant k -finite sliding window denoiser satisfies $\hat{x}_i(Z_{i-k}^{i+k}) = \hat{x}(Z_{i-k}^{i+k}, Z_i)$ for $\hat{x} : \mathcal{Z}^{2k+1} \times \mathcal{Z} \rightarrow \hat{\mathcal{X}}$. We can view k -blocks of symbols as *supersymbols* to be denoised. We thus define the k -finite block denoiser to divide the output sequence Z^n sequentially into $\lfloor \frac{n}{k} \rfloor$ blocks of k symbols and up to one remainder block of less than k symbols: $Z_{M(k-1)+1}^{Mk}$, $M = 1, 2, \dots, \lfloor \frac{n}{k} \rfloor$ and $Z_{\lfloor \frac{n}{k} \rfloor k+1}^n$. A block of reconstruction symbols is produced after observing each output block, so that $\hat{X}_{M(k-1)+1}^{Mk} = \hat{x}_{M(k-1)+1}^{Mk}(Z_{M(k-1)+1}^{Mk})$, $M = 1, 2, \dots, \lfloor \frac{n}{k} \rfloor$ and $\hat{X}_{\lfloor \frac{n}{k} \rfloor k+1}^n = \hat{x}_{\lfloor \frac{n}{k} \rfloor k+1}^n(Z_{\lfloor \frac{n}{k} \rfloor k+1}^n)$ (see Figure 5). It is straightforward to deduce the form of a time-invariant k -finite block filter from the above definition.

B. Filtering

The basic idea in filtering is to reconstruct X^n *causally*. That is, the filtering function $\hat{x}_i(\cdot)$ at time i may depend only on Z^i so that $\hat{X}_i = \hat{x}_i(Z^i)$ as in Figure 6. So, the most general filtering functions can make use of all of Z^i when deciding on output $\hat{x}_i$. The memory of a general filter is unbounded in that the number of observation symbols Z_i used to make the decision on $\hat{x}_i$ grows arbitrarily large with i . We call this most general class of filters the class of *infinite* filters. It turns out to be difficult to analyze such filters so we now define some classes of filters with *finite* memory (i.e. their output at time i depends on a fixed number of past output symbols) which are interesting in their own right. In particular, we consider the *symbol-by-symbol filter*, which is the same as the symbol-by-symbol denoiser (i.e. $\hat{X}_i = \hat{x}_i(Z_i)$) of Figure 3, and the k -finite memory filter, which allows $\hat{X}_i$ to depend on Z_{i-k}^i (i.e. $\hat{X}_i = \hat{x}_i(Z_{i-k}^i)$) as shown in Figure 7. Similar to the above, a time-invariant k -finite memory filter satisfies $\hat{x}_i(Z_{i-k}^i) = \hat{x}(Z_{i-k}^i, Z_i)$ for $\hat{x} : \mathcal{Z}^{k+1} \times \mathcal{Z} \rightarrow \hat{\mathcal{X}}$.

Fig. 6. Filter

Fig. 7. k -Finite Memory Filter

C. Criterion for Optimality

We assume a given single letter loss function $\Lambda(x, \hat{x}) : \mathcal{X} \times \hat{\mathcal{X}} \rightarrow \mathbb{R}^+$. Note that what we call loss is also referred to as distortion, particularly in the context of source coding. An example of such a single letter loss function is Hamming loss, where $\Lambda(x, \hat{x}) = 0$ if $x = \hat{x}$ and $\Lambda(x, \hat{x}) = 1$ if $x \neq \hat{x}$. We denote the maximum value of the single-letter loss function by $\Lambda_{\max}$.

We consider the normalized cumulative loss

$$L_n = \frac{1}{n} \sum_{i=1}^n \Lambda(X_i, \hat{X}_i). \quad (1)$$

For the general denoising setup,

$$L_n = \frac{1}{n} \sum_{i=1}^n \Lambda(X_i, \hat{x}_i(Z^n)). \quad (2)$$

The cumulative loss for the other denoisers and for the filters is defined analogously, with the appropriate restrictions on the functions $\hat{x}_i$. Note that L_n depends on the particular denoiser/filter as well as on X^n and Z^n . We omit these from the notation for readability.

The normalized cumulative loss is a random variable. The performance of a denoiser is usually characterized by the expected value of the normalized cumulative loss. We take a different approach and examine the probability that the normalized cumulative loss exceeds some threshold D . For D less than or equal to the minimum achievable expected loss, this probability goes to 1 when using the optimal scheme, by the law of large numbers. Thus, we consider values of D that exceed the minimum achievable expected loss. In the sequel, a denoiser (i.e. the collection of denoising functions $\hat{x}^n(\cdot) = (\hat{x}_1(\cdot), \dots, \hat{x}_n(\cdot))$) as described in II-A) will be said to be *optimal* if it achieves the best exponential rate of decay of $P(L_n > D)$. Similarly, the *optimal* symbol-by-symbol, k -finite block, k -finite sliding window and k -finite memory denoisers/filters maximize the exponential rate of decay of $P(L_n > D)$ among all symbol-by-symbol, k -finite sliding window, k -finite block and k -finite memory schemes, respectively (i.e. among the schemes where the denoising/filtering functions $\hat{x}_i(\cdot)$ are chosen so that the denoiser/filter $\hat{x}^n(\cdot)$ is symbol-by-symbol, k -finite sliding window, k -finite block or k -finite memory, respectively, as in II-A and II-B).

III. MAIN RESULTS

In section IV, we prove

Theorem 1:

$$R(\mathcal{C}, D) \triangleq \lim_{n \rightarrow \infty} -\frac{1}{n} \log \min_{\hat{x}^n(\cdot) \in \mathcal{C}} P(L_n > D), \quad (3)$$

exists for $\mathcal{C}$ equal to

- $\mathcal{C}_1$, the class of all denoisers,
- $\mathcal{C}_2$, the class of symbol-by-symbol denoisers,
- $\mathcal{C}_3$, the set of k -finite block denoisers,
- $\mathcal{C}_4$, the class of k -finite sliding window denoisers,
- $\mathcal{C}_5$, the class of k -finite memory filters.

We call $R(\mathcal{C}, D)$ the *optimal rate function* for class $\mathcal{C}$. Furthermore, $\lim_{n \rightarrow \infty} -\frac{1}{n} \log P(L_n > D)$ exists for any symbol-by-symbol denoiser (i.e. $\hat{x}^n \in \mathcal{C}$) and we call it the *rate function for symbol-by-symbol denoisers*.

Theorem 2: $R(\mathcal{C}_1, D) \geq R(\mathcal{C}_2, D)$, where the inequality is strict for some sources, channels and distortion criteria. That is, in general, symbol-by-symbol denoisers are suboptimal. The optimal denoiser and the optimal symbol-by-symbol denoiser are time-invariant.

Theorem 3: Under Hamming loss, $R(\mathcal{C}_1, D) = R(\mathcal{C}_2, D)$, i.e. symbol-by-symbol denoising is optimal.

Theorem 4: $R(\mathcal{C}_2, D) = R(\mathcal{C}_3, D) = R(\mathcal{C}_4, D) = R(\mathcal{C}_5, D)$. That is, in general, finite block denoisers, finite sliding window denoisers and finite memory filters do no better than symbol-by-symbol denoisers/filters. Since $\mathcal{C}_2 \subset \mathcal{C}_i$ for $i > 2$, Theorem 2 implies that the optimal rate functions for $\mathcal{C}_i$, $i > 2$ are achieved by time-invariant symbol-by-symbol denoisers.

Remark 1: We give concrete examples where the inequality of Theorem 2 is strict.

Remark 2: Theorem 2 seems somewhat counterintuitive since the source is i.i.d., the channel is memoryless, and the distortion is single-letter.

Remark 3: To obtain Theorem 4, we first compute $R(\mathcal{C}_3, D)$ and show that $R(\mathcal{C}_2, D) = R(\mathcal{C}_3, D)$. We use these two results to obtain $R(\mathcal{C}_2, D) = R(\mathcal{C}_4, D)$ by an approximation argument. Finally, we show that $R(\mathcal{C}_2, D) = R(\mathcal{C}_5, D)$ by observing that $\mathcal{C}_4 \supset \mathcal{C}_5$.

IV. OPTIMAL RATE FUNCTIONS AND OPTIMALITY OF DENOISERS/FILTERS

We prove the first two parts of Theorem 1 in Section IV-A. We establish the last part of Theorem 1 and the time-invariance of the optimal symbol-by-symbol denoiser of Theorem 2 in Section IV-B. In Section IV-C, we show that the optimal denoisers in $\mathcal{C}_1$ and $\mathcal{C}_2$ are time-invariant and find a set of examples where symbol-by-symbol denoising is strictly suboptimal thereby showing Theorem 2. Section IV-D characterizes the performance of k -finite block denoisers and shows it to be equivalent to the that of symbol-by-symbol denoisers, thus establishing the third part of Theorem 1 and part of Theorem 4. Section IV-E does the same for k -finite sliding window denoisers. Finally, Section IV-F discusses filtering, explaining why it is difficult to analyze the performance of filters with unbounded memory and then characterizing finite memory filters in order to establish the remainder of Theorems 1 and 4.

A. $R(\mathcal{C}_1, D)$ and $R(\mathcal{C}_2, D)$

1) *Optimal Denoisers:* We present upper and lower bounds on $P(L_n > D)$ for an optimal denoiser, preceded by some notation. The results use a modification of the large-deviations analysis for arbitrarily varying sources found in Appendix B of [5]. Since our result is somewhat different, we provide the complete proof in Appendix I. Consider a set of r probability distributions on the real line: $\{P_1, P_2, \dots, P_r\}$, $r < \infty$. Denote the support of P_a by $S(P_a)$, $a = 1, 2, \dots, r$. Suppose that for all a , every element of $S(P_a)$, is upper bounded by $\Delta_a \in \mathbb{R}$ and that $0 < P_a(\{\Delta_a\}) < 1$. Now, let W_i be independent random variables with distribution in $\{P_a\}_{a=1}^r$. Then, following the terminology of [5], we call W_i an arbitrarily varying source (AVS). Define

$$S_n = \frac{1}{n} \sum_{i=1}^n W_i.$$

Denote the fraction of W_i in W^n with distribution P_a by Q_a^n . Let $\Psi_a(\cdot)$ be the moment generating function of a random variable with distribution P_a , i.e.

$$\Psi_a(\lambda) = E_{P_a}(e^{\lambda W}) = \sum_{w \in S(P_a)} e^{\lambda w} P_a(w).$$

Let $\Psi_{(n)}(\lambda) = \sum_{a=1}^r Q_a^n \log(\Psi_a(\lambda))$. Let

$$\Psi_{(n)}^*(x) = \sup_{\lambda \geq 0} [\lambda x - \Psi_{(n)}(\lambda)] = \sup_{\lambda \geq 0} \left[\lambda x - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda) \right].$$

We also define the following quantities which are used in the proof and later in this work although they are not part of the statement of the lemma. We let $\bar{w}_n = \sum_{a=1}^r Q_a^n E_{P_a} W$, $\Delta_{\max} = \max_a \Delta_a$, $\Delta_{(n)} = \sum_{a=1}^r Q_a^n \Delta_a$ and $\Psi_a^*(x) = \sup_{\lambda \geq 0} [\lambda x - \Psi_a(\lambda)]$.

Then, we have

Lemma 1 (Large Deviations for AVS): For $D > 0$,

$$P(S_n > D) \leq e^{-n\Psi_{(n)}^*(D)},$$

and

$$P(S_n > D) \geq e^{-n[\Psi_{(n)}^*(D) + o(n)]}.$$

Proof: The proof and the precise expression of $o(n)$ can be found in Appendix I. ■

With Lemma 1, we can compute $P(L_n > D)$ for a general denoiser.

Consider an arbitrary denoiser $\hat{x}^n(\cdot) : \mathcal{Z}^n \rightarrow \mathcal{X}^n$. We examine the relation

$$\begin{aligned} P(L_n > D | Z^n = z^n) \\ = P\left(\frac{1}{n} \sum_{i=1}^n \Lambda(X_i, \hat{x}_i(Z^n)) > D | Z^n = z^n\right). \end{aligned}$$

Now, since $\hat{x}(\cdot)$ is deterministic, there is no randomness in $\hat{x}(Z^n)$ given $Z^n = z^n$. Consider the set $I_{\hat{x},z} = \{i : (\hat{x}_i, z_i) = (\hat{x}, z)\}$. Then for all $i \in I_{\hat{x},z}$, $\Lambda(X_i, \hat{x}_i)$ are conditionally independent given $Z^n = z^n$ and have the same distribution as $\Lambda(X, \hat{x})|_{Z=z}$, where X has the source distribution and Z has the distribution induced by the channel and by the distribution of X . So, $\Lambda(X_i, \hat{x}_i)$ given $Z^n = z^n$ is an arbitrarily varying source, where the finite set of possible distributions of this random variable is indexed by the possible values of $(\hat{x}_i, z_i)$, and has magnitude $|\hat{\mathcal{X}}||\mathcal{Z}|$. We can thus use Lemma 1.

Let Q_z^n be the fraction of occurrences of $z_i = z$ in z^n , and let $Q_{\hat{x}|z}^n$ be the fraction of occurrences of $(\hat{x}_i, z_i) = (\hat{x}, z)$ among the nQ_z^n pairs $(\hat{x}_i, z_i)$ with $z_i = z$. Given $Z^n = z^n$, the denoiser $\hat{x}^n(\cdot)$ induces a particular $Q_{\hat{x}|z}^n$. To simplify the notation, we will not show the dependence of $Q_{\hat{x}|z}^n$ on $\hat{x}(\cdot)$ and Z^n . We have

$$\begin{aligned} P(L_n > D | Z^n = z^n) \\ = P\left(\sum_{z \in \mathcal{Z}} \sum_{\hat{x} \in \hat{\mathcal{X}}} \sum_{i=1}^{nQ_z^n Q_{\hat{x}|z}^n} Y_i^{\hat{x},z} > nD\right), \end{aligned}$$

where $\{Y_i^{\hat{x},z}\}$ are independent with $Y_i^{\hat{x},z}$ distributed as $\Lambda(X, \hat{x})|_{Z=z}$.

We now apply Lemma 1 to the random variables $Y_i^{\hat{x},z}$. Δ_a corresponds to $\max_{x:p(x|z)>0} \Lambda(x, \hat{x}) \triangleq \Lambda_{\hat{x},z}$, $\bar{w}_n$ corresponds to $\bar{\Lambda}_n \triangleq \sum_z \sum_{\hat{x}} Q_z^n Q_{\hat{x}|z}^n EY^{\hat{x},z}$, $\Delta_{\max}$ corresponds to $\Lambda_{\max}$, $\Delta_{(n)}$ corresponds to $\Lambda_{(n)} \triangleq \sum_z \sum_{\hat{x}} Q_z^n Q_{\hat{x}|z}^n \Lambda_{\hat{x},z}$ and r corresponds to $|\hat{\mathcal{X}}||\mathcal{Z}|$. We will omit the explicit specification of the alphabets of $\hat{x}$ and z , for simplicity. Also, we use $\{Q_z^n\}$ to denote the empirical distribution of z , i.e. $\{Q_1^n, Q_2^n, \dots, Q_{|\mathcal{Z}|}^n\}$. Similarly, $\{Q_{\hat{x}|z}^n\}$ denotes the collection of conditional empirical distributions of $\hat{x}$ given $z \in \{1, 2, \dots, |\mathcal{Z}|\}$. Then, using Lemma 1, we have that:

$$P(L_n > D | Z^n = z^n) \leq e^{-nI(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\})}, \quad (4)$$

and

$$P(L_n > D | Z^n = z^n) \geq e^{-n(I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + o(n))}, \quad (5)$$

where $o(n)$ is independent of Z^n and where, for probability distributions (using similar notation to that of the empirical distributions) $\{Q_z\}$ on Z and $\{Q_{\hat{x}|z}\}$ on $\hat{\mathcal{X}}$,

$$I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) \quad (6)$$

$$= \sup_{\lambda \geq 0} \left[\lambda D - \sum_{z, \hat{x}} Q_z Q_{\hat{x}|z} \log \left(\sum_x e^{\lambda \Lambda(x, \hat{x})} p(x|z) \right) \right], \quad (7)$$

where $p(x|z)$ is the given conditional distribution of the channel input given the channel output. We now restrict our attention to the schemes that maximize the exponential rate of decay of

$$P(L_n > D) = \sum_{z^n} P(L_n > D | Z^n = z^n) P(Z^n = z^n), \quad (8)$$

i.e. those that achieve (3). Notice that the probabilities in (4) and (5) depend on z^n and the denoiser only through the joint empirical type of $(\hat{x}^n, z^n)$. We claim that the the best (in the sense of maximizing the exponential rate of

decay of $P(L_n > D)$), joint empirical type, $(\hat{x}^n, z^n)$, is constant for z^n of the same type. The reason is that the set of possible joint types of $(\hat{x}^n, z^n)$ is identical for z^n of the same type. This is easily seen by considering $z^{n,1}$ and $z^{n,2}$ to be of the same type and noting that, because they are of the same type, there is a bijection $\pi(\cdot)$ from $\mathcal{Z}^n \rightarrow \mathcal{Z}^n$ between $z^{n,2}$ and $z^{n,1}$. So, if a particular denoiser $\hat{x}^n(\cdot)$ produces joint type $\{Q_{\hat{x},z}^n\}$ when used on $z^{n,1}$, the denoiser resulting from the composition of $\hat{x}^n(\cdot)$ and $\pi(\cdot)$ on $z^{n,2}$ produces the same joint type $\{Q_{\hat{x},z}^n\}$. The opposite is clearly true since $\pi(\cdot)$ is a bijection.

Since the set of joint types is the same, the best exponent is the same (since it depends only on the joint type). So, rather than summing over all z^n in (8), a set with magnitude exponential in n , we can group the z^n according to their type and sum over the different types, $\{Q_z^n\}$ (a set with magnitude polynomial in n) and, we may restrict our attention to denoising functions that depend only on $\{Q_z^n\}$. These facts will help us to express the desired probability as asymptotically equal to an exponential in n .

We will omit explicit dependence of the notation of the denoiser on $\{Q_z^n\}$ for brevity. Further, we can use the classical typical sequence bounds on the probability that Z^n has type $\{Q_z^n\}$ (c.f. [9], [10], [11]). Thus, for a particular choice of the denoising functions chosen among the set of optimal denoising functions, i.e., where $Q_{\hat{x}|z}^n$ induced by the denoiser depends on z^n only through its type, Q_z^n ,

$$\begin{aligned} P(L_n > D) &= \sum_{z^n} P(L_n > D | Z^n = z^n) P(Z^n = z^n) \\ &\leq \sum_{Z^n = z^n} e^{-nI(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\})} P(Z^n = z^n) \\ &= \sum_{\{Q_z^n\}} e^{-nI(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\})} P(Z^n \text{ has type } \{Q_z^n\}) \\ &\leq \sum_{\{Q_z^n\}} e^{-n(I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + D(\{Q_z^n\} \| Q))}, \end{aligned} \quad (9)$$

where Q denotes the distribution of the channel output, and

$$\begin{aligned} P(L_n > D) &\geq \sum_{\{Q_z^n\}} e^{-n((D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + o(n))} P(Z^n \text{ has type } \{Q_z^n\}) \\ &\geq \sum_{\{Q_z^n\}} e^{-n(I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + o(n) + D(\{Q_z^n\} \| Q) + \frac{|\mathcal{X}| \log(n+1)}{n})}. \end{aligned} \quad (10)$$

Note that the summations in (9) and (10) are over the set of possible empirical distributions, which is of polynomial size.

The optimal denoiser chooses the best denoising functions given the type of z^n . Denoting the loss of the optimal denoiser by L_n^{opt} , we have

$$P(L_n^{opt} > D | Z^n = z^n) \leq e^{-n \max_{Q_{\hat{x}|z}^n} I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\})}, \quad (11)$$

and

$$P(L_n^{opt} > D | Z^n = z^n) \geq e^{-n(\max_{\{Q_{\hat{x}|z}^n\}} I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + o(n))}. \quad (12)$$

Therefore,

$$P(L_n^{opt} > D) \doteq e^{-n \min_{\{Q_z^n\}} \max_{\{Q_{\hat{x}|z}^n\}} (I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + D(\{Q_z^n\} \| Q))}, \quad (13)$$

where we use the notation $a_n \doteq b_n$ to denote that $\lim_{n \rightarrow \infty} \frac{1}{n} \log \frac{a_n}{b_n} = 0$.

2) *Optimal Symbol-by-Symbol Denoisers:* Now we derive the best performance among the class of symbol-by-symbol denoisers. The optimal symbol-by-symbol denoiser must choose the denoising functions before observing the realized type so that $\{Q_{\hat{x}|z}^n\}$ is a deterministic mapping (i.e. it is the same for all types $\{Q_z^n\}$). It thus picks a set of denoising functions that maximize the exponent over all types (see equation (14)). Denoting the loss of the symbol-by-symbol scheme by L_n^{ss} , we have

$$P(L_n^{ss} > D) \doteq e^{-n \max_{\hat{x}_1(z_1), \dots, \hat{x}_n(z_n)} \frac{1}{n} \log P(L_n > D)} \quad (14)$$

$$\doteq e^{-n \max_{\{Q_{\hat{x}|z}^n\}} \min_{\{Q_z^n\}} (I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + D(\{Q_z^n\} \| Q))}, \quad (15)$$

where (15) follows from (9) and (10). We present an alternate derivation of this optimal rate function in Section IV-B.

3) *Optimal Rate Functions:* We compute (3) by showing that we can move the limit inside the optimizations and then optimizing over a continuum of distributions instead of over the discrete sets $\{Q_z^n\}$ and $\{Q_{\hat{x}|z}^n\}$. We thus define this new domain of optimization variables.

Definition 1: For $z \in \mathcal{Z}$ and $\hat{x} \in \hat{\mathcal{X}}$, let $Q_z \in [0, 1]$ and $Q_{\hat{x}|z} \in [0, 1]$, such that $\sum_z Q_z = 1$, and $\sum_{\hat{x}|z} Q_{\hat{x}|z} = 1$ for each z . We can think of Q_z as the frequency of $Z = z$ in Z^n as $n \rightarrow \infty$ and likewise for $Q_{\hat{x}|z}$. We denote the collection of such frequencies by: $\{Q_z\} = \{Q_1, Q_2, \dots, Q_{|\mathcal{Z}|}\}$ and $\{Q_{\hat{x}|z}\} = \{Q_{1|1}, Q_{2|1}, \dots, Q_{|\hat{\mathcal{X}}||1}, \dots, Q_{|\hat{\mathcal{X}}|||\mathcal{Z}|}\}$.

We now claim that we can move the limit inside the optimization. We let $\bar{\Lambda} = \sum_z \sum_{\hat{x}} Q_z Q_{\hat{x}|z} EY^{\hat{x},z}$ and $\Lambda = \sum_z \sum_{\hat{x}} Q_z Q_{\hat{x}|z} \Lambda_{\hat{x},z}$.

Then,

Lemma 2: For $\bar{\Lambda} \leq D < \Lambda$,

$$\begin{aligned} & \lim_{n \rightarrow \infty} \max_{\{Q_z^n\}} \min_{\{Q_{\hat{x}|z}^n\}} (I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + D(\{Q_z^n\} \| Q)) \\ &= \max_{\{Q_{\hat{x}|z}\}} \min_{\{Q_z\}} (I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\} \| Q)), \end{aligned} \quad (16)$$

and

$$\begin{aligned} & \lim_{n \rightarrow \infty} \min_{\{Q_z^n\}} \max_{\{Q_{\hat{x}|z}^n\}} (I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + D(\{Q_z^n\} \| Q)) \\ &= \min_{\{Q_z\}} \max_{\{Q_{\hat{x}|z}\}} (I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\} \| Q)). \end{aligned} \quad (17)$$

We first need the following claims:

Claim 1: $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\})$ is convex in $\{Q_z\}$ and also in $\{Q_{\hat{x}|z}\}$.

Proof: Consider the linear combination of $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}^{(1)})$ and $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}^{(2)})$.

$$\begin{aligned} & \alpha \left[\sup_{\lambda \geq 0} \lambda D - \sum_{z, \hat{x}} Q_z Q_{\hat{x}|z}^{(1)} \log \left(\sum_x e^{\lambda \Lambda(x, \hat{x}(z))} p(x|z) \right) \right] \\ &+ \bar{\alpha} \left[\sup_{\lambda \geq 0} \lambda D - \sum_{z, \hat{x}} Q_z Q_{\hat{x}|z}^{(2)} \log \left(\sum_x e^{\lambda \Lambda(x, \hat{x}(z))} p(x|z) \right) \right] \\ &= \sup_{\lambda \geq 0} \alpha \lambda D - \sum_{z, \hat{x}} Q_z \alpha Q_{\hat{x}|z}^{(1)} \log \left(\sum_x e^{\lambda \Lambda(x, \hat{x}(z))} p(x|z) \right) \\ &+ \sup_{\lambda \geq 0} \bar{\alpha} \lambda D - \sum_{z, \hat{x}} Q_z \bar{\alpha} Q_{\hat{x}|z}^{(2)} \log \left(\sum_x e^{\lambda \Lambda(x, \hat{x}(z))} p(x|z) \right) \\ &\geq \alpha \lambda' D - \sum_{z, \hat{x}} Q_z \alpha Q_{\hat{x}|z}^{(1)} \log \left(\sum_x e^{\lambda' \Lambda(x, \hat{x}(z))} p(x|z) \right) \\ &\bar{\alpha} \lambda'' D - \sum_{z, \hat{x}} Q_z \bar{\alpha} Q_{\hat{x}|z}^{(2)} \log \left(\sum_x e^{\lambda'' \Lambda(x, \hat{x}(z))} p(x|z) \right) \end{aligned}$$

$\forall \lambda', \lambda''$.

Choose $\lambda' = \lambda''$ so that this equals

$$\sup_{\lambda \geq 0} \lambda D - \sum_{z, \hat{x}} Q_z \left(\alpha Q_{\hat{x}|z}^{(1)} + \bar{\alpha} Q_{\hat{x}|z}^{(2)} \right) \log \left(\sum_x e^{\lambda \Lambda(x, \hat{x}(z))} p(x|z) \right).$$

A similar argument holds for $\{Q_z\}$. ■

Claim 2: $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)$ is convex in $\{Q_z\}$ and also in $\{Q_{\hat{x}|z}\}$.

Proof: The claim follows from the previous claim and the fact that $D(\{Q_z\}||Q)$ is convex in $\{Q_z\}$ and independent of $\{Q_{\hat{x}|z}\}$. ■

Claim 3: For $\bar{\Lambda} \leq D < \Lambda$, and M satisfying $0 < M < \infty$, $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)$ is uniformly continuous in $\{Q_{\hat{x}|z}\}$ and $\{\{Q_z\} : D(\{Q_z\}||Q) \leq M\}$.

Proof: The set of allowable $\{Q_{\hat{x}|z}\}$ is closed and bounded and thus compact. Since $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)$ is convex in $\{Q_{\hat{x}|z}\}$, it is uniformly continuous in $\{Q_{\hat{x}|z}\}$. Since $D(\{Q_z\}||Q)$ is continuous where finite, the set $\{\{Q_z\} : D(\{Q_z\}||Q) \leq M\}$ is closed. Clearly this set is also bounded since the range of $\{Q_z\}$ is bounded. Uniform continuity follows from the compactness of this set and the convexity of $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)$.

It is also clear that the types $\{Q_z\}$ such that $D(\{Q_z\}||Q) = \infty$ cannot minimize the rate function, so we may assume the existence of some $M > 0$ such that the optimization is equivalent to optimizing over $\{\{Q_z\} : D(\{Q_z\}||Q) \leq M\}$. To keep the notation simple, we will not explicitly state this restricted range of values of $\{Q_z\}$ in the following. ■

We are now ready to prove Lemma 2.

Proof: For all $\eta_1 > 0$, there exists an N such that $n > N$ implies

$$\max(\min_{\{Q_z^n\}} \|\{Q_z^n\} - \{Q_z\}\|_2, \min_{\{Q_{\hat{x}|z}^n\}} \|\{Q_{\hat{x}|z}^n\} - \{Q_{\hat{x}|z}\}\|_2) < \eta_1$$

for all $\{Q_z\}, \{Q_{\hat{x}|z}\}$.

Since $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)$ is uniformly continuous in $\{Q_{\hat{x}|z}\}$ and in the allowed $\{Q_z\}$, for all $\eta_2 > 0$, there is an $\eta_3 > 0$ such that

$$\max(\min_{\{Q_z^n\}} \|\{Q_z^n\} - \{Q_z\}\|_2, \min_{\{Q_{\hat{x}|z}^n\}} \|\{Q_{\hat{x}|z}^n\} - \{Q_{\hat{x}|z}\}\|_2) < \eta_3$$

implies

$$|(I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)) - (I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}) + D(\{Q_z^n\}||Q))| < \eta_2.$$

Thus, for fixed $\eta_2 > 0$, we can find the N such that for all $n > N$,

$$\max(\min_{\{Q_z^n\}} \|\{Q_z^n\} - \{Q_z\}\|_2, \min_{\{Q_{\hat{x}|z}^n\}} \|\{Q_{\hat{x}|z}^n\} - \{Q_{\hat{x}|z}\}\|_2) < \eta_3.$$

Then, for any fixed $\{Q_{\hat{x}|z}\}$ and $n > N$, we have

$$\begin{aligned} \min_{\{Q_z\}} I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q) &\leq \min_{\{Q_z^n\}} I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z^n\}||Q) \\ &\leq \min_{\{Q_z\}} I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q) + \eta_2. \end{aligned}$$

If we now optimize over $\{Q_{\hat{x}|z}\}$, we have

$$\begin{aligned} \max_{\{Q_{\hat{x}|z}\}} \min_{\{Q_z\}} I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q) - \eta_2 \\ \leq \max_{\{Q_{\hat{x}|z}\}} \min_{\{Q_z^n\}} I(D, \{Q_z^n\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z^n\}||Q) \\ \leq \max_{\{Q_{\hat{x}|z}\}} \min_{\{Q_z\}} I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q) + \eta_2. \end{aligned}$$

Since η_2 is arbitrary, we have the first part of the lemma. The second part follows analogously. ■

Thus, combining Lemma 2 with (13) and (15) shows that $R(\mathcal{C}_1, D)$ and $R(\mathcal{C}_2, D)$ are well defined and that

$$R(\mathcal{C}_1, D) = \min_{\{Q_z\}} \max_{\{Q_{\hat{x}|z}\}} (I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)) \quad (18)$$

and

$$R(\mathcal{C}_2, D) = \max_{\{Q_{\hat{x}|z}\}} \min_{\{Q_z\}} (I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\}||Q)). \quad (19)$$

B. Alternate Derivation of $R(\mathcal{C}_2, D)$ and Rate Function for Symbol-by-Symbol Denoisers

We can find the rate function for a symbol-by-symbol denoiser by noting that $\Lambda(X_i, \hat{x}_i(Z_i))$ is an A.V.S. with distribution depending only on the denoising function, $\hat{x}_i(\cdot)$. There are a finite number, $|\hat{\mathcal{X}}|^{|Z|}$, of different denoising functions, $\hat{x}(\cdot)$ which we will now label $\hat{x}^{(j)}(\cdot)$, $j = 1, \dots, |\hat{\mathcal{X}}|^{|Z|}$. Let the fraction of times function $\hat{x}^{(j)}(\cdot)$ appears be $\theta_j^{(n)}$ (there is a dependence on n since there are n total observations being denoised). Then, we can apply Lemma 1 to conclude that, for any symbol-by-symbol denoiser,

$$P(L_n > D) \doteq e^{-n \sup_{\lambda > 0} \left[\lambda D - \sum_{j=1}^{|\hat{\mathcal{X}}|^{|Z|}} \theta_j^{(n)} \log E_{X,Z} \exp(\lambda \Lambda(X, \hat{x}^{(j)}(Z))) \right]}, \quad (20)$$

thus establishing the last part of Theorem 1.

We get an alternate derivation of the optimal rate function by optimizing over the $\{\theta_j^{(n)}\}$ to get

$$\begin{aligned} P(L_n^{ss} > D) &\doteq \min_{\{\theta_j^{(n)}\}} e^{-n \sup_{\lambda > 0} \left[\lambda D - \sum_{j=1}^{|\hat{\mathcal{X}}|^{|Z|}} \theta_j^{(n)} \log E_{X,Z} \exp(\lambda \Lambda(X, \hat{x}^{(j)}(Z))) \right]} \\ &\doteq e^{-n \max_{\{\theta_j^{(n)}\}} \sup_{\lambda > 0} \left[\sum_{j=1}^{|\hat{\mathcal{X}}|^{|Z|}} \theta_j^{(n)} (\lambda D - \log E_{X,Z} \exp(\lambda \Lambda(X, \hat{x}^{(j)}(Z)))) \right]} \\ &\doteq e^{-n \max_{\{\theta_j^{(n)}\}} \sum_{j=1}^{|\hat{\mathcal{X}}|^{|Z|}} \theta_j^{(n)} \sup_{\lambda > 0} [\lambda D - \log E_{X,Z} \exp(\lambda \Lambda(X, \hat{x}^{(j)}(Z)))]} \\ &\doteq e^{-n \max_{\hat{x}(\cdot)} \sup_{\lambda > 0} [\lambda D - \log E_{X,Z} \exp(\lambda \Lambda(X, \hat{x}(Z)))]}. \end{aligned} \quad (21)$$

(21) follows from the fact that the optimal denoising function, $\hat{x}_j(\cdot)$, depends only on the distribution of (X, Z) for all j and is therefore the same function for all j . In other words, the optimal symbol-by-symbol scheme is time-invariant, establishing part of Theorem 2. We now have another expression for $R(\mathcal{C}_2, D)$, namely

$$R(\mathcal{C}_2, D) = \max_{\hat{x}(\cdot)} \sup_{\lambda > 0} [\lambda D - \log E_{X,Z} \exp(\lambda \Lambda(X, \hat{x}(Z)))]. \quad (22)$$

C. Optimal Denoising and Theorem 2

1) *Theory*: Our goal is to investigate whether $P(L_n^{opt} > D) \doteq P(L_n^{ss} > D)$, i.e. whether $R(\mathcal{C}_1, D) = R(\mathcal{C}_2, D)$. We can immediately see from (18) that

Lemma 3: For Hamming loss, symbol-by-symbol denoising is optimal.

Proof: For all types $\{Q_z\}$, $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\})$ (see (6)) is maximized over $\{Q_{\hat{x}|z}\}$ by $Q_{\hat{x}|z} = \mathbb{I}(\{\hat{x} = \arg\max_x p(x|z)\})$, i.e. the deterministic conditional distribution that sets $\hat{x}$ as the the most likely x given z . ■ We thus have Theorem 3. Using (18) and (19), we now show that

Lemma 4: The best denoisers and the best symbol-by-symbol denoisers are time invariant.

Proof: The best denoiser picks a conditional distribution $Q_{\hat{x}|z}$ based on Q_z . It is easy to extend Claim 1 to continuous distributions. Thus, for a fixed $\{Q_z\}$, $I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\} \| Q)$ is convex in $Q_{\hat{x}|z}$. Thus, for each $z \in \mathcal{Z}$, the best choice of $Q_{\hat{x}|z}$ sets 1 for some $\hat{x}$ and 0 else, i.e. the best denoiser is time-invariant. The best symbol-by-symbol denoiser chooses $\{Q_{\hat{x}|z}\}$ to maximize $\min_{\{Q_z\}} (I(D, \{Q_z\}, \{Q_{\hat{x}|z}\}) + D(\{Q_z\} \| Q))$. It is easy to see that this expression is convex in $\{Q_{\hat{x}|z}\}$. So, for each $z \in \mathcal{Z}$, the best symbol-by-symbol denoiser has $Q_{\hat{x}|z}$ is 1 for some $\hat{x}$ and 0 else. Thus, it is time-invariant. ■

2) *Concrete Examples of Suboptimality*: We now show that there are cases for which symbol-by-symbol denoising is strictly suboptimal, i.e. the inequality between (18) and (19) is strict.

Consider a Bernoulli(p), $p = 0.5$, source that passes through a BSC(δ), $\delta < 0.5$, with an asymmetric loss function. Specifically, a loss of 1 is incurred when we decode 0 as 1 and a loss of 2 is incurred when we decode a 1 as a 0. This setup is illustrated in Figure 8. By Lemma 4, it is clear that the best denoiser has $Q_{0|z} = 0$ or 1 for $z = 0, 1$. So, for a given $\{Q_z\}$, there is no need to timeshare; the best denoiser makes the same decision at each time for the same output symbol z . Thus, there are only four possible optimal denoising schemes: say-what-you see (SWYS), say-the-opposite ($\overline{\text{SWYS}}$), decode all ones (ONES) and decode all zeros (ZEROS). We represent the denoising decision by $\hat{x}(\cdot)$ which takes a single symbol z as an argument. So, for the SWYS denoiser, $\hat{x}(z) = z$ and for the ONES denoiser, $\hat{x}(z) = 1$.

Since we are in the binary setting, we can simplify the notation for the frequency/distribution vectors. Instead of

Fig. 8. Setup for Symbol-by-Symbol Suboptimality Examples

writing $\{Q_z\}$ or Q , we specify the frequency/distribution, respectively, by Q_0 or $Q(0)$. In our setting, the objective function, $I\left(D, \{Q_z^n\}, \{Q_{\hat{x}|z}^n\}\right) + D(\{Q_z^n\} \| Q)$, for a particular denoiser, $\hat{x}(\cdot)$, is:

$$\begin{aligned} & \sup_{\lambda \geq 0} \lambda D - Q_0 \log \left(e^{\lambda \Lambda(0, \hat{x}(0))} \frac{p(0, 0)}{Q(0)} + e^{\lambda \Lambda(1, \hat{x}(0))} \frac{p(1, 0)}{Q(0)} \right) \\ & - Q_1 \log \left(e^{\lambda \Lambda(0, \hat{x}(1))} \frac{p(0, 1)}{Q(1)} + e^{\lambda \Lambda(1, \hat{x}(1))} \frac{p(1, 1)}{Q(1)} \right) \end{aligned} \quad (23)$$

$$\begin{aligned} & + Q_0 \log \frac{Q_0}{Q(0)} + Q_1 \log \frac{Q_1}{Q(1)} \\ & = \sup_{\lambda \geq 0} \lambda D - \log E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))} \\ & + D_b \left(Q_0 \left\| \frac{\sum_x p(x, 0) e^{\lambda \Lambda(X, \hat{x}(0))}}{E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))}} \right\| \right), \end{aligned} \quad (24)$$

where D_b denotes binary divergence. That is, $D_b(p(0) \| q(0)) = p(0) \log \frac{p(0)}{q(0)} + p(1) \log \frac{p(1)}{q(1)}$, where p and q are probability distributions on $\{0, 1\}$.

We can now show that

Claim 4: $\overline{\text{SWYS}}$ and ZEROS are suboptimal.

Proof: For SWYS, the quantity inside the supremum of (23) is:

$$\lambda D - Q_0 \log(\bar{\delta} + e^{2\lambda} \delta) - Q_1 \log(e^\lambda \delta + \bar{\delta}). \quad (25)$$

For $\overline{\text{SWYS}}$, it is:

$$\lambda D - Q_1 \log(\delta + e^{2\lambda} \bar{\delta}) - Q_0 \log(e^\lambda \bar{\delta} + \delta). \quad (26)$$

Since $\delta < 0.5$, (25) for $Q_0 = Q_0^1$ is greater than or equal to (26) for $Q_0 = 1 - Q_0^1$, for each $\lambda \geq 0$. So, for any fixed denoiser and $Q_0 = Q_0^1$, (23) for SWYS is better than for $\overline{\text{SWYS}}$ for that same denoiser and $Q_0 = 1 - Q_0^1$. Thus, the performance of SWYS is better than that of $\overline{\text{SWYS}}$.

Since $p = 0.5$, we are as likely to incorrectly decode a 1 as we are to incorrectly decode a 0. Since it is more costly to mistake a 1, the ONES denoiser is better than the ZEROS denoiser. ■

We also have that

Claim 5:

$$\lambda D - \log E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))} + D_b \left(Q_0 \left\| \frac{\sum_x p(x, 0) e^{\lambda \Lambda(X, \hat{x}(0))}}{E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))}} \right\| \right) \quad (27)$$

is concave in λ .

Proof: Since log convexity is preserved under sums and $e^{\lambda a}$, $a \in \mathbb{R}$ is log convex in λ , the terms inside the logarithm of (23) are log convex. Hence, the log terms are convex and so (27) is concave in λ . ■

Fig. 9. Region of symbol-by-symbol Suboptimality

Now, (16) can be expressed as follows:

$$\max_{\hat{x}(\cdot)} \min_{Q_0} \sup_{\lambda \geq 0} \lambda D - \log E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))} + D_b \left(Q_0 \left\| \frac{\sum_x p(x, 0) e^{\lambda \Lambda(X, \hat{x}(0))}}{E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))}} \right\| \right) \quad (28)$$

$$= \max_{\hat{x}(\cdot)} \sup_{\lambda \geq 0} \min_{Q_0} \lambda D - \log E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))} + D_b \left(Q_0 \left\| \frac{\sum_x p(x, 0) e^{\lambda \Lambda(X, \hat{x}(0))}}{E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))}} \right\| \right) \quad (29)$$

$$= \max_{\hat{x}(\cdot)} \sup_{\lambda \geq 0} \lambda D - \log E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))}, \quad (30)$$

where we can switch the min and sup in (28) to get (29) since the objective is convex in the minimization variable Q_0 and concave in λ , the variable over which the supremum is taken. Equality (30) follows by setting $Q_0 = \frac{\sum_x p(x, 0) e^{\lambda \Lambda(X, \hat{x}(0))}}{E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))}}$ so that the binary divergence term is minimized. We could have obtained (30) directly from (20) but we re-derived it here because we use the form of (29) in the following.

Since only the SWYS and ONES denoisers can be optimal, the problem reduces to comparing the exponents of these two denoisers. That is, we use (30) and substitute either the SWYS or ONES function for $\hat{x}(\cdot)$. We use a Matlab program to search the space of channels in terms of δ and the range of thresholds, D , to determine for which channels and threshold values symbol-by-symbol denoising is strictly suboptimal. Although the region of such (δ, D) pairs is computed numerically, each point in the region can be verified analytically to show the suboptimality of symbol-by-symbol denoising. The details of our method and an explanation of why each point in the region can be verified analytically can be found in Appendix 2. Figure 9 shows a plot of the region of (δ, D) for which symbol-by-symbol denoising is suboptimal. This concludes the proof of Theorem 2.

D. $R(\mathcal{C}_3, D)$

We derive an expression for the exponent of the probability of excess loss for the k -finite block denoiser and show that the best k -finite block denoiser does no better than the best symbol-by-symbol denoiser.

For simplicity, assume that $n = mk$, for integer m . Since we take $\lim_{n \rightarrow \infty} -\frac{1}{n} \log P(L_n > D)$, it will be clear that this does not affect the validity of the derivation. We index the set of denoising functions, $\hat{x}^k(\cdot) : \mathcal{Z}^k \rightarrow \hat{\mathcal{X}}^k$ by $i \in \mathcal{I}_k = \left\{ 1, 2, \dots, |\hat{\mathcal{X}}|^{k|\mathcal{Z}|^k} \right\}$ to get $\hat{x}^{k,(i)}(\cdot)$. Note that for fixed k , this set is finite. Now, given $n = mk$, the most general deterministic scheme will use a certain fraction of each type of denoising function. Denote the fraction of time denoiser i is used by $\theta_i^{(m)} \in \left\{ 0, \frac{1}{m}, \frac{2}{m}, \dots, 1 \right\}$. Since the $\left(X_{(j-1)k+1}^{jk}, Z_{(j-1)k+1}^{jk} \right)$ are i.i.d., for a particular i ,

$\Lambda(X_{(j-1)k+1}^{jk}, \hat{x}^{k,(i)}(Z_{(j-1)k+1}^{jk}))$ are i.i.d.. So, $\Lambda(X^k, \hat{x}^{k,(i)}(Z^k))$ is an AVS. Thus,

$$\begin{aligned} P(L_n > D) &= P\left(\sum_{j=1}^n \Lambda(X_j, \hat{X}_j) > nD\right) \\ &= P\left(\sum_{j=1}^m \Lambda(X_{(j-1)k+1}^{jk}, \hat{X}_{(j-1)k+1}^{jk}) > nD\right) \\ &= P\left(\frac{1}{m} \sum_{j=1}^m \Lambda(X_{(j-1)k+1}^{jk}, \hat{X}_{(j-1)k+1}^{jk}) > kD\right) \\ &= P\left(\frac{1}{m} \sum_{i=1}^{|\mathcal{I}_k|} \sum_{j=1}^{m\theta_i^{(m)}} \Lambda(X^{k,(i,j)}, \hat{x}^{k,(i)}(Z^{k,(i,j)})) > kD\right), \end{aligned}$$

where for each (i, j) , $(X^{k,(i,j)}, Z^{k,(i,j)})$ has the same distribution as (X^k, Z^k) . We can therefore use Lemma 1 on the AVS $\Lambda(X^k, \hat{x}^{k,(i)}(Z^k))$ to compute

$$\begin{aligned} & -\frac{1}{n} \log P(L_n > D) \\ &= -\frac{1}{mk} \log P\left(\frac{1}{m} \sum_{i=1}^{|\mathcal{I}_k|} \sum_{j=1}^{m\theta_i^{(m)}} \Lambda(X^{k,(i,j)}, \hat{x}^{k,(i)}(Z^{k,(i,j)})) > kD\right) \\ &= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda kD - \sum_{i=1}^{|\mathcal{I}_k|} \theta_i^{(m)} \log E_{X^k, Z^k} e^{\lambda \Lambda(X^k, \hat{x}^{k,(i)}(Z^k))} \right] + o(m) \end{aligned}$$

where $o(m) \rightarrow 0$ as $m \rightarrow \infty$. Since we are concerned about the behavior for fixed k as $n \rightarrow \infty$, we have $m \rightarrow \infty$. So, we can neglect the $o(m)$ term and optimize over $\{\theta_i\}$, $\theta_i \in [0, 1]$ instead of over the $\{\theta_i^{(m)}\}$. We can rewrite this as

$$\begin{aligned} & \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda kD - \sum_{i=1}^{|\mathcal{I}_k|} \theta_i \log \left(\sum_{x^k, z^k} (p(x^k, z^k) e^{\lambda \Lambda(x^k, \hat{x}^{k,(i)}(z^k))}) \right) \right] \\ &= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda kD - \sum_{i=1}^{|\mathcal{I}_k|} \theta_i \log \left(\sum_{x^k, z^k} (p(z^k) p(x^k | z^k) e^{\lambda \Lambda(x^k, \hat{x}^{k,(i)}(z^k))}) \right) \right] \\ &= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda kD - \sum_{i=1}^{|\mathcal{I}_k|} \theta_i \log \left(\sum_{x^k, z^k} \left(p(z^k) \prod_{j=1}^k (p(x_j | z_j)) e^{\sum_{j=1}^k \lambda \Lambda(x_j, \hat{x}_j^{(i)}(z^k))} \right) \right) \right] \\ &= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda kD - \sum_{i=1}^{|\mathcal{I}_k|} \theta_i \log \left(\sum_{x^k, z^k} \left(p(z^k) \prod_{j=1}^k (p(x_j | z_j) e^{\lambda \Lambda(x_j, \hat{x}_j^{(i)}(z^k))}) \right) \right) \right] \\ &= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda kD - \sum_{i=1}^{|\mathcal{I}_k|} \theta_i \log \left(\sum_{z^k} \left(p(z^k) \prod_{j=1}^k \sum_{x_j} (p(x_j | z_j) e^{\lambda \Lambda(x_j, \hat{x}_j^{(i)}(z^k))}) \right) \right) \right] \\ &= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda kD - \sum_{i=1}^{|\mathcal{I}_k|} \theta_i \log \left(\sum_{z^k} \left(p(z^k) \prod_{j=1}^k \sum_x (p(x | z_j) e^{\lambda \Lambda(x, \hat{x}_j^{(i)}(z^k))}) \right) \right) \right]. \end{aligned}$$

To maximize this expression, we should set $\theta_i = 1$ for $\hat{x}_j^{(i)}(z^k) = \operatorname{argmin}_{\hat{x}} \sum_x p(x | z_j) e^{\lambda \Lambda(x, \hat{x})}$ and 0 else since this minimizes $\sum_x p(x | z_j) e^{\lambda \Lambda(x, \hat{x}_j^{(i)}(z^k))}$ for each j . So, the best denoiser uses the time invariant, symbol-by-symbol function $\hat{x}(z) = \operatorname{argmin}_{\hat{x}} \sum_x p(x | z) e^{\lambda \Lambda(x, \hat{x})}$, where λ achieves the above supremum. Since the function is

continuous in λ , starts at 0 for $\lambda = 0$ and tends to $-\infty$ as $\lambda \rightarrow \infty$, the supremum is achieved and so our definition of the optimal function is valid. Letting $\hat{x}(\cdot)$ denote this choice of denoising function, we thus have

$$\begin{aligned}
& \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda k D - \log \left(\sum_{x^k, z^k} p(x^k, z^k) \prod_{j=1}^k (e^{\lambda \Lambda(x_j, \hat{x}(z_j))}) \right) \right] \\
&= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda k D - \log \left(\sum_{x^{k-1}, z^{k-1}} p(x^{k-1}, z^{k-1}) \prod_{j=1}^{k-1} (e^{\lambda \Lambda(x_j, \hat{x}(z_j))}) \sum_{x_k, z_k} p(x_k, z_k) e^{\lambda \Lambda(x_k, \hat{x}(z_k))} \right) \right] \\
&= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda k D - \log \left(\prod_{j=1}^k \sum_{x_j, z_j} p(x_j, z_j) e^{\lambda \Lambda(x_j, \hat{x}(z_j))} \right) \right] \\
&= \frac{1}{k} \sup_{\lambda \geq 0} \left[\lambda k D - \log \left(\sum_{x, z} p(x, z) e^{\lambda \Lambda(x, \hat{x}(z))} \right)^k \right] \\
&= \sup_{\lambda \geq 0} \left[\lambda D - \log \left(\sum_{x, z} p(x, z) e^{\lambda \Lambda(x, \hat{x}(z))} \right) \right] \\
&= R(\mathcal{C}_3, D).
\end{aligned}$$

Clearly, the value of λ that maximizes this expression coincides with the value of λ that maximizes the symbol-by-symbol denoiser rate function (22). Therefore, finite block denoisers have the same performance as symbol-by-symbol denoisers, i.e. $R(\mathcal{C}_2, D) = R(\mathcal{C}_3, D)$, giving the third part of Theorem 1 and the first part of Theorem 4.

E. $R(\mathcal{C}_4, D)$

We show that, given a k -finite sliding window denoiser, we can find a sequence of finite block denoisers of increasing order whose performance is a lower bound on $P(L_n > D)$ for the k -finite sliding window denoiser. Consider a p -finite block denoiser, where $n = mp$, $p = lk$ and $l \gg 1$ and m, l are integers. It is straightforward to extend the argument to general m, l . We now show that

$$\begin{aligned}
& P(L_n > D) \\
&= P \left(\sum_{i=1}^n \Lambda(X_i, \hat{x}_i(Z_{i-k}^{i+k})) > nD \right) \\
&\geq P \left(\sum_{j=1}^m \Lambda \left(X_{(j-1)p+1}^{jp}, \hat{x}_{(j-1)p+1}^{jp}(Z_{(j-1)p+1}^{jp}) \right) > n \left(D + \frac{2\Lambda_{\max} k}{p} \right) \right),
\end{aligned}$$

This latter expression is the probability of excess loss of a p -finite block denoiser with threshold $D + \frac{2\Lambda_{\max} k}{p}$ and uses the notation for block denoisers defined in Section II-A. Figure 10 illustrates the reason for the inequality. The inequality follows from the fact that the p -finite block denoiser has more information than the k -finite sliding window denoiser for all indices except those of the form $jp - k + 1, jp - k + 2, \dots, jp + k$. So the best block denoiser does at least as well as the sliding window denoiser for indices that are not of this form. Furthermore, the loss for indices that are of this form cannot exceed $\Lambda_{\max}$. Since there are $2km$ such indices, increasing nD by $2km\Lambda_{\max}$ to get $n(D + \frac{2\Lambda_{\max} k}{p})$ gives the lower bound. We know the p -finite block denoiser can do no better than the optimal symbol-by-symbol denoiser with the same threshold, i.e. $D + \frac{2\Lambda_{\max} k}{p}$. Since the exponent is continuous in the threshold parameter and p was arbitrary, taking $p \rightarrow \infty$ gives us a tighter lower bound on the exponent associated with the probability of excess loss of a sliding window denoiser. This lower bound is the probability of excess loss for an optimal symbol-by-symbol denoiser with parameter D .

It is obvious that the best finite sliding window denoiser is no worse than a symbol-by-symbol denoiser of the same threshold. Thus, the performance of the best finite sliding window denoiser is the same as the performance of the best symbol-by-symbol denoiser, i.e. $R(\mathcal{C}_4, D) = R(\mathcal{C}_2, D)$. So we have another part of Theorem 1 and of Theorem 4.

Fig. 10. Illustration for Sliding Window Denoiser Proof

F. Filtering

1) *Infinite Memory Filters*: Explicit characterization of the performance of the infinite memory filter appears to be difficult. The problem shares some intricacies with the characterization of the exponent of zero-delay, infinite memory source codes, which is mentioned but left open in [5]. It is not clear how to use the AVS Lemma (Lemma 1) since the single-letter losses at different times are dependent for the infinite memory filter. This was also the case with the finite sliding window denoiser, but because the memory and look-ahead were finite, we could get a handle on the rate function by using a series of finite block denoisers to upper bound it. We are, however, able to characterize the performance of the finite memory filter, by sandwiching its performance between schemes whose performance we already know.

2) $R(\mathcal{C}_5, D)$: The analysis of finite memory filters is greatly simplified by the preceding results for denoisers. We observe that the set of k -finite sliding window denoising functions includes the set of k -finite memory filtering functions which includes the set of symbol-by-symbol denoising/filtering functions. The equivalence of the best k -finite sliding window filter and the best symbol-by-symbol denoiser/filter thus implies the performance of finite memory filters is the same as the performance of symbol-by-symbol filters, i.e. $R(\mathcal{C}_5, D) = R(\mathcal{C}_2, D)$. This gives the remaining parts of Theorems 1 and 4.

V. CONCLUSIONS AND FUTURE WORK

We have studied the effect of limiting the domain of denoising and filtering functions under the probability of excess loss criterion. We established Theorems 1-4, which we now rephrase.

Symbol-by-symbol denoising of a DMC-corrupted memoryless source is found to be suboptimal using a general single-letter loss function, under the probability of excess loss criterion. In the case of Hamming loss, symbol-by-symbol denoising is optimal. In general, the best denoising and symbol-by-symbol denoising schemes are time-invariant.

A region of suboptimality for a $\text{Bern}(\frac{1}{2})$ source passing through a $\text{BSC}(\delta)$ under an asymmetric loss function was found numerically. Each point of the region can be verified analytically, but an analytical characterization of the region of suboptimality is yet to be found and may be of interest.

We have shown that finite memory filters, finite sliding window denoisers, and finite block denoisers all do no better than time-invariant symbol-by-symbol denoisers/filters.

We note that the case where the filter has unbounded memory is also of interest. An open question is how to characterize the performance of these infinite memory filters, or even to determine whether/when the performance is strictly better or worse than that of symbol-by-symbol filters and optimal denoisers, respectively.

APPENDIX I: PROOF OF THE AVS LEMMA

We use the notation given in the statement of Lemma 1. We start by showing that

Claim 6: $\lambda x - \log \Psi_a(\lambda)$ is concave in λ .

Proof: Since $e^{\lambda z}$ is log convex and log convexity is preserved under sums, $\Psi_a(\lambda)$ is a log convex function of λ . This implies that $\log \Psi_a(\lambda)$ is convex. Thus, $\Psi_{(n)}(\lambda)$ is also convex in λ and Claim 6 follows.

Clearly $\lambda x - \log \Psi_a(\lambda)$ is 0 when $\lambda = 0$. If $x < \Delta_a$,

$$\begin{aligned} \lim_{\lambda \rightarrow \infty} [\lambda x - \log \Psi_a(\lambda)] &= \lim_{\lambda \rightarrow \infty} \log \frac{e^{\lambda x}}{\sum_{z \in S(P_a)} P_a(z) e^{\lambda z}} \\ &= \log \lim_{\lambda \rightarrow \infty} \frac{e^{\lambda(x - \Delta_a)}}{P_a(\Delta_a) + \sum_{z \in S(P_a), z \neq \Delta_a} P_a(z) e^{\lambda(z - \Delta_a)}} \\ &= -\infty, \end{aligned}$$

since $P_a(\Delta_a) > 0$. Also,

$$\begin{aligned} \lim_{\lambda \rightarrow 0} \frac{d}{d\lambda} [\lambda x - \log \Psi_a(\lambda)] &= \lim_{\lambda \rightarrow 0} \left[x - \frac{\sum_{z \in S(P_a)} z P_a(z) e^{\lambda z}}{\sum_{z \in S(P_a)} P_a(z) e^{\lambda z}} \right] \\ &= x - E_{P_a} Z. \end{aligned}$$

So, if $E_{P_a} Z \leq x < \Delta_a$, the derivative is nonnegative near zero and the sup is achieved for λ in $[0, \infty)$. Conversely, $x < E_{P_a} Z$, then the supremum is achieved by $\lambda = 0$ and has value 0.

Thus, we have the following lemma:

Lemma 5: The $\sup_{\lambda \geq 0}$ of $\lambda x - \log e^{\lambda \Delta_a}$ is always achieved for $x < \Delta_a$ and is the solution of

$$\frac{d}{d\lambda} [\lambda x - \log \Psi_a(\lambda)] = 0$$

when $x > E_{P_a} Z$. It is 0 otherwise.

We note that we will be using the variable δ in what follows. This δ is *not* the same as the parameter of the BSC mentioned in the body of this paper. We reuse the variable here to simplify the notation. There should be no ambiguity since this appendix is self-contained.

Lemma 5 implies the existence of $\lambda_a(\delta)$, the achiever of $\Psi_a^*(\Delta_a - \delta)$, for $0 < \delta < \Delta_a$.

Define $\lambda_0(\delta) = \max_a \lambda_a(\delta)$. Now, we have that

Claim 7: $\lambda x - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda)$ is concave.

Proof: This follows easily since $\lambda x - \log \Psi_a(\lambda)$ is concave. ■

Now, $\lambda x - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda)$ is 0 at $\lambda = 0$ and, if $x < \Delta_{(n)}$,

$$\begin{aligned} \lim_{\lambda \rightarrow \infty} \left[\lambda x - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda) \right] &= \lim_{\lambda \rightarrow \infty} \log \frac{e^{\lambda x}}{\prod_{a=1}^r \Psi_a(\lambda)^{Q_a^n}} \\ &= \lim_{\lambda \rightarrow \infty} \log \frac{e^{\lambda(x - \Delta_{(n)})}}{\prod_{a=1}^r (P_a(\Delta_a) + \sum_{z \in S(P_a), z \neq \Delta_a} P_a(z) e^{\lambda(z - \Delta_a)})^{Q_a^n}} \\ &= \log \lim_{\lambda \rightarrow \infty} \frac{e^{\lambda(x - \sum_{a=1}^r Q_a^n \Delta_a)}}{\prod_{a=1}^r (P_a(\Delta_a) + \sum_{z \in S(P_a), z \neq \Delta_a} P_a(z) e^{\lambda(z - \Delta_a)})^{Q_a^n}} \\ &= -\infty \end{aligned}$$

Also,

$$\begin{aligned}
& \lim_{\lambda \rightarrow 0} \frac{d}{d\lambda} \left[\lambda x - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda) \right] \\
&= \lim_{\lambda \rightarrow 0} \left[x - \sum_{a=1}^r Q_a^n \frac{\sum_{z \in S(P_a)} z P_a(z) e^{\lambda z}}{\sum_{z \in S(P_a)} P_a(z) e^{\lambda z}} \right] \\
&= x - \sum_{a=1}^r Q_a^n E_{P_a} Z \\
&= x - \bar{z}_n.
\end{aligned}$$

So, if $\bar{z}_n \leq x < \Delta_{(n)}$, the derivative is nonnegative near zero and the sup is achieved for λ in $[0, \infty)$. Conversely, if $x < \bar{z}_n$, the supremum is zero and is achieved by $\lambda = 0$.

Thus, we have the following lemma:

Lemma 6: The achiever always exists for $x < \Delta_{(n)}$ and is the solution of

$$\frac{d}{d\lambda} [\lambda x - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda)] = 0 \quad (31)$$

when $x > \bar{z}_n$. It is 0 otherwise.

We thus have the existence of $\lambda_{(n)}(\delta)$, the achiever of $\Psi_{(n)}^*(\Delta_{(n)} - \delta)$, for $0 < \delta < \Delta_{(n)}$. We are now ready to prove the upper bound of Lemma 1.

Proof: We first assume that $\bar{z}_n \leq D < \Delta_{(n)}$. Then, using $\mathbb{I}$ to denote the indicator function, for all $\lambda \geq 0$,

$$\begin{aligned}
P(S_n > D) &= E(\mathbb{I}(S_n - D \geq 0)) \\
&\leq E(e^{n\lambda(S_n - D)}) \\
&= e^{-n\lambda D} \prod_{a=1}^r \prod_{i=1}^{nQ_a^n} E_{P_a}(e^{\lambda Z}) \\
&= e^{-n\lambda D} \prod_{a=1}^r E_{P_a}(e^{\lambda Z})^{nQ_a^n} \\
&= e^{-n[\lambda D - \sum_{a=1}^r Q_a^n \log E_{P_a}(e^{\lambda Z})]} \\
&= e^{-n[\lambda D - \sum_{a=1}^r Q_a^n \Psi_a(\lambda)]}
\end{aligned}$$

Thus,

$$\begin{aligned}
P(S_n > D) &\leq e^{[-n \sup_{\lambda \geq 0} (\lambda D - \sum_{a=1}^r Q_a^n \Psi_a(\lambda))]} \\
&= e^{-n\Psi_{(n)}^*(D)}
\end{aligned}$$

If $D \geq \Delta_{(n)}$,

$$\begin{aligned}
P(S_n > D) &\leq \sum_{a=1}^r P\left(\frac{1}{nQ_a^n} \sum_{i=1}^{nQ_a^n} Z_i^a \geq \Delta_a\right) \\
&= 0
\end{aligned}$$

where we use the notation Z_i^p, Z_j^q , $i \neq j$ to denote independent random variables distributed according to P_p, P_q , respectively.

Since $\Psi_{(n)}(D) = \infty$ for $D \geq \Delta_{(n)}$, the upper bound holds.

Finally, for $0 < D < \bar{z}_n$, $\Psi_{(n)}(D) = 0$ and since $P(S_n > D) < 1$, the upper bound holds. This concludes the proof of the upper bound. $\blacksquare$

We prove the lower bound of Lemma 1 by first showing a restricted version of the lemma:

Lemma 7: Fix $\delta > 0$ such that $\delta < \frac{\Delta_{(n)} - \bar{z}_n}{2}$. Let $\bar{z}_n \leq D \leq \Delta_{(n)} - 2\delta$. Choose an ϵ such that $0 < \epsilon < \delta$. For $n > \frac{r^2 \Delta_{\max}^2 \log 8r}{\delta^2}$,

$$P(S_n > D) \geq \exp \left[-n \left(\Psi_{(n)}^*(D) + \frac{\lambda_0(\delta) \Delta_{\max} r \sqrt{2 \log(8r)}}{\sqrt{n}} - \frac{\log(\frac{3}{4})}{n} \right) \right].$$

Proof: Define the events

$$E = \left\{ \sum_{t=1}^n Z_t \geq nD \right\}$$

and

$$F = \left\{ \sum_{t=1}^n Z_t \leq n(D + 2\epsilon) \right\}.$$

Define a new set of probability measures, $\{P_{a,\lambda}^n\}_{a=1}^r$, by $P_{a,\lambda}^n(z) = P_a(z) \frac{e^{\lambda z}}{\Psi_a(\lambda)}$.

Since $\bar{z}_n < D + \epsilon < D + \delta \leq \Delta_{(n)} - \delta$, by Lemma 6, the achiever of $\Psi_{(n)}^*(D + \epsilon)$ exists. We denote it by λ_n^* . By Lemma 6, λ_n^* satisfies:

$$\frac{d}{d\lambda} [\lambda(D + \epsilon) - \Psi_{(n)}(\lambda)] = 0. \quad (32)$$

so that

$$\begin{aligned} D + \epsilon &= \frac{d}{d\lambda} \Psi_{(n)}(\lambda_n^*) \\ &= \sum_{a=r}^r Q_a^{(n)} \frac{\sum_{z \in S(P_a)} z P_a(z) e^{\lambda_n^* z}}{\sum_{z \in S(P_a)} P_a(z) e^{\lambda_n^* z}} \\ &= \sum_{a=r}^r Q_a^{(n)} \frac{\sum_{z \in S(P_a)} z P_a(z) e^{\lambda_n^* z}}{\Psi_a(\lambda_n^*)} \\ &= \sum_{a=r}^r Q_a^{(n)} E_{P_{a,\lambda_n^*}}(Z). \end{aligned}$$

Now,

$$\begin{aligned} P \left(\frac{1}{n} \sum_{t=1}^n Z_t > D \right) &\geq P(E \cap F) \\ &= \sum_{z^n \in E \cap F} \prod_{t=1}^n P(z_t) \\ &= \sum_{z^n \in E \cap F} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_a(z_n \sum_{i=1}^a Q_i^n + t). \end{aligned}$$

Letting $h(a, t, n) = n \sum_{i=1}^a Q_i^n + t$, and substituting in the modified probability measure, the above expression

becomes

$$\begin{aligned}
& \sum_{z^n \in E \cap F} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} \Psi_a(\lambda_n^*) \exp[-\lambda_n^* z_{h(a,t,n)}] P_{a,\lambda_n^*}(z_{h(a,t,n)}) \\
&= \sum_{z^n \in E \cap F} \prod_{a=1}^r \exp(nQ_a^n \log(\Psi_a(\lambda_n^*)) - \lambda_n^* \sum_{t=1}^{nQ_a^n} Q_a^n z_{h(a,t,n)}) \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a,\lambda_n^*}(z_{h(a,t,n)}) \\
&\geq \exp \left[n \left(\sum_{a=1}^r Q_a^n \log \Psi_a(\lambda_n^*) - \lambda_n^* (D + 2\epsilon) \right) \right] \sum_{z^n \in E \cap F} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a,\lambda_n^*}(z_{h(a,t,n)}) \tag{33}
\end{aligned}$$

$$= \exp[-n(\lambda_n^* (D + 2\epsilon) - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda_n^*))] \sum_{z^n \in E \cap F} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a,\lambda_n^*}(z_{h(a,t,n)}) \tag{34}$$

where (33) follows from the fact that $z^n \in F$.

Now,

$$\exp \left[-n \left(\lambda_n^* (D + 2\epsilon) - \sum_{a=1}^r Q_a^n \log \Psi_a(\lambda_n^*) \right) \right] \geq \exp \left[-n \left(2\lambda_n^* \epsilon + \Psi_{(n)}^*(D) \right) \right] \tag{35}$$

since $\lambda_n^* D - \Psi_{(n)}(\lambda_n^*) \leq \Psi_{(n)}^*(D)$. To proceed, we need

Lemma 8: For all $x, y > 0$, if λ_1 achieves $\Psi_{(n)}^*(x)$ and λ_2 achieves $\Psi_{(n)}^*(x + y)$, then $\lambda_1 \leq \lambda_2$.

Proof:

$$\Psi_{(n)}^*(x + y) = \lambda_2(x + y) - \Psi_n(\lambda_2) \tag{36}$$

$$\geq \lambda_1(x + y) - \Psi_n(\lambda_1) \tag{37}$$

$$= \lambda_1 y + \Psi_{(n)}^*(x) \tag{38}$$

$$\geq \lambda_1 y + \lambda_2 x - \Psi_n(\lambda_2) \tag{39}$$

and (36) $\geq$ (39) implies $\lambda_2 \geq \lambda_1$.

Similarly, if λ_1 achieves $\Psi_a^*(x)$ and λ_2 achieves $\Psi_a^*(x + y)$, then $\lambda_1 \leq \lambda_2$. ■

It is straightforward to see that $\lambda_a(\delta)$ is continuous, and that, as $\delta \rightarrow 0$, $\lambda_a(\delta) \rightarrow \infty$, for all a . For a fixed $\delta > 0$, $\lambda_a(\delta) < \infty$ and is independent of n . Thus $\lambda_0(\delta) = \max_a \lambda_a(\delta)$ achieves all $\Psi_a^*(\Delta_a - \phi_a)$ for some collection $\{\phi_a\}_{a=1}^r$ such that, for each a , $\phi_a \leq \delta$. We thus have that:

$$\sum_{a=1}^r Q_a^n \sup_{\lambda \geq 0} [\lambda(\Delta_a - \phi_a) - \Psi_a(\lambda)] = \sum_{a=1}^r Q_a^n [\lambda_0(\delta)(\Delta_a - \phi_a) - \Psi_a(\lambda_0(\delta))] \tag{40}$$

$$= \sup_{\lambda \geq 0} \left[\sum_{a=1}^r Q_a^n (\lambda(\Delta_a - \phi_a) - \Psi_a(\lambda)) \right] \tag{41}$$

where the last equality follows by the choice of the $\{\phi_a\}$, which implies that $\lambda_0(\delta)$ achieves (41).

Now,

$$\begin{aligned}
(41) &= \sup_{\lambda \geq 0} \left\{ \lambda \left[\sum_{a=1}^r Q_a^n (\Delta_a - \phi_a) \right] - \Psi_{(n)}(\lambda) \right\} \\
&= \sup_{\lambda \geq 0} \left[\lambda \left(\Delta_{(n)} - \sum_{a=1}^r Q_a^n \phi_a \right) - \Psi_{(n)}(\lambda) \right]
\end{aligned}$$

and $\sum_{a=1}^r Q_a^n \phi_a \leq \delta$. Thus,

$$\Delta_{(n)} - \sum_{a=1}^r Q_a^n \phi_a \geq \Delta_{(n)} - \delta.$$

And so, by Lemma 8, we have

$$\lambda_{(n)}(\delta) \leq \lambda_0(\delta) = \max_a \lambda_a(\delta) < \infty.$$

Also, $\epsilon < \delta$ and $D < \Delta_{(n)} - 2\delta$, so

$$D + \epsilon < \Delta_{(n)} - 2\delta + \epsilon < \Delta_{(n)} - \delta.$$

Thus,

$$\lambda_n^* \leq \lambda_{(n)}(\delta) \leq \lambda_0(\delta) < \infty.$$

We use this in (35) to get,

$$\exp \left[-n \left(2\lambda_n^* \epsilon + \Psi_{(n)}^*(D) \right) \right] \geq \exp \left[-n \left(2\lambda_0(\delta) \epsilon + \Psi_{(n)}^*(D) \right) \right], \quad (42)$$

which is a lower bound on the first part of (34). We now bound the second part of (34).

$$\begin{aligned} & \sum_{z^n \in E \cap F} \sum_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a, \lambda_n^*}(z_{h(a,t,n)}) \\ & \geq 1 - \sum_{z^n \in E^C} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a, \lambda_n^*}(z_{h(a,t,n)}) - \sum_{z^n \in F^C} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a, \lambda_n^*}(z_{h(a,t,n)}) \end{aligned} \quad (43)$$

Also, note that $n(D + \epsilon) = \sum_{a=1}^r Q_a^n E_{P_{a, \lambda_n^*}}(Z)$ from (33). Thus, we have

$$\sum_{z^n \in F^C} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a, \lambda_n^*}(z_{h(a,t,n)}) = P \left(\sum_{a=1}^r \sum_{t=1}^{nQ_a^n} z_{h(a,t,n)}^a - n(D + \epsilon) > n\epsilon \right) \quad (44)$$

$$= P \left(\sum_{a=1}^r \sum_{t=1}^{nQ_a^n} (z_{h(a,t,n)}^a - E_{P_{a, \lambda_n^*}} Z) > n\epsilon \right) \quad (45)$$

$$\leq P \left(\bigcup_{a=1}^r \left[\sum_{t=1}^{nQ_a^n} (z_{h(a,t,n)}^a - E_{P_{a, \lambda_n^*}} Z) > \frac{n\epsilon}{r} \right] \right) \quad (46)$$

$$\leq \sum_{a=1}^r P \left(\sum_{t=1}^{nQ_a^n} (z_{h(a,t,n)}^a - E_{P_{a, \lambda_n^*}} Z) > \frac{(nQ_a^n)\epsilon}{Q_a^n r} \right) \quad (47)$$

$$\leq \sum_{a=1}^r \exp \left[-\frac{2nQ_a^n \epsilon^2}{(Q_a^n)^2 r^2 \Delta_a^2} \right] \quad (48)$$

$$= \sum_{a=1}^r \exp \left[-\frac{2n\epsilon^2}{Q_a^n r^2 \Delta_a^2} \right] \quad (49)$$

$$\leq r \exp \left[-\frac{2n\epsilon^2}{r^2 \Delta_{\max}^2} \right], \quad (50)$$

where (47) follows from the union bound, (48) follows from the Hoeffding bound and (50) follows from the facts

that $Q_a^n \leq 1$ and $\Delta_a^2 \leq \Delta_{\max}^2$.

Similarly, we have

$$\begin{aligned}
\sum_{z^n \in E^C} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a,\lambda_n^*}(z_{h(a,t,n)}) &= P \left(\sum_{a=1}^r \sum_{t=1}^{nQ_a^n} z_{h(a,t,n)}^a - n(D + \epsilon) < -n\epsilon \right) \\
&= P \left(\sum_{a=1}^r \sum_{t=1}^{nQ_a^n} (-z_{h(a,t,n)}^a - (-E_{P_{a,\lambda_n^*}} Z)) > n\epsilon \right) \\
&\leq P \left(\bigcup_{a=1}^r \left[\sum_{t=1}^{nQ_a^n} (-z_{h(a,t,n)}^a - (-E_{P_{a,\lambda_n^*}} Z)) > \frac{n\epsilon}{r} \right] \right) \\
&\leq \sum_{a=1}^r P \left(\sum_{t=1}^{nQ_a^n} (-z_{h(a,t,n)}^a - (-E_{P_{a,\lambda_n^*}} Z)) > \frac{(nQ_a^n)\epsilon}{Q_a^n r} \right) \\
&\leq \sum_{a=1}^r \exp \left[-\frac{2nQ_a^n \epsilon^2}{(Q_a^n)^2 r^2 \Delta_a^2} \right] \\
&= \sum_{a=1}^r \exp \left[-\frac{2n\epsilon^2}{Q_a^n r^2 \Delta_a^2} \right] \\
&\leq r \exp \left[-\frac{2n\epsilon^2}{r^2 \Delta_{\max}^2} \right].
\end{aligned}$$

Thus,

$$\sum_{z^n \in E \cap F} \prod_{a=1}^r \prod_{t=1}^{nQ_a^n} P_{a,\lambda_n^*}(z_{h(a,t,n)}) \geq 1 - 2r \exp \left[-\frac{2n\epsilon^2}{r^2 \Delta_{\max}^2} \right] = \frac{3}{4} \quad (51)$$

when $\epsilon = \frac{r\Delta_{\max}}{\sqrt{n}} \sqrt{\frac{\log(8r)}{2}}$. Since we assumed $\epsilon < \delta$, the proof is valid when

$$\delta > \frac{r\Delta_{\max}}{\sqrt{n}} \sqrt{\frac{\log(8r)}{2}}.$$

So we require,

$$n > \frac{r^2 \Delta_{\max}^2 \log(8r)}{2\delta^2}. \quad (52)$$

Thus, combining (42) and (51) yields:

$$\begin{aligned}
P(E) &\geq \exp \left[-n \left(2\lambda_0(\delta)\epsilon + \Psi_{(n)}^*(D) \right) \right] \frac{3}{4} \\
&\geq \exp \left[-n \left(\Psi_{(n)}^*(D) + 2\lambda_0(\delta) \frac{r\Delta_{\max}}{\sqrt{n}} \sqrt{\frac{\log(8r)}{2}} - \frac{\log(\frac{3}{4})}{n} \right) \right] \\
&= \exp \left[-n \left(\Psi_{(n)}^*(D) + \frac{\lambda_0(\delta)\Delta_{\max}r\sqrt{2\log(8r)}}{\sqrt{n}} - \frac{\log(\frac{3}{4})}{n} \right) \right],
\end{aligned}$$

concluding the proof of the restricted form of the lemma, with $N = \frac{r^2 \Delta_{\max}^2 \log(8r)}{2\delta^2}$ and $o(n) = \frac{\lambda_0(\delta)\Delta_{\max}r\sqrt{2\log(8r)}}{\sqrt{n}} - \frac{\log(\frac{3}{4})}{n}$. We thus have the lower bound of Lemma 1 for $\bar{z}_n \leq D < \Delta_{(n)}$ since δ was arbitrary. ■

We now prove the lower bound of Lemma 1.

Proof: Observe that for $D \geq \Delta_{(n)}$, $P(S_n > D) = 0$. Since $\Psi_{(n)}^*(D) = \infty$ for such D , the lemma is true for all n and with $o(n) = 0$.

For $D < \bar{z}_n$, $P(S_n > D) \geq P(S_n > \bar{z}_n) = e^{-o(n)}$, for some N , $o(n)$ found in the manner described in the proof above. For such D , $\Psi_{(n)}(D) = 0$, so the lemma is true. This concludes the proof of the lower bound. ■

APPENDIX II: DETAILS OF METHOD TO COMPUTE REGION OF SYMBOL-BY-SYMBOL SUBOPTIMALITY

We have shown that the only candidates for the optimal denoiser and optimal symbol-by-symbol denoiser are the SWYS and ONES denoisers. So, in order to determine whether symbol-by-symbol performance is optimal, we must determine when minimizing the rate function over all types Q_0 and then maximizing over the choice of denoiser, SWYS or ONES, is equivalent to maximizing the rate function over the choice of denoiser, SYWS or ONES, and then minimizing over the type Q_0 .

We have also demonstrated that the rate function for a particular denoiser is a convex function of Q_0 .

So the question, illustrated in Figure 11, becomes: when is the max (over the two schemes SWYS and ONES) of the min (over Q_0) of the two convex functions of Q_0 strictly less than the min (over Q_0) of the max (over SWYS and ONES)? We need only compute the value of the exponents at two points each in order to determine when the

Fig. 11. Comparing two convex functions. The solid dots represent maxmin or minmax points. The shaded dots represent the minimum of a function. The vertical lines connect the minimum of a particular function to the corresponding point on the other function. We see that when the maxmin is strictly less than the minmax, the difference between each minimum point and the corresponding point on the other function is negative. When this does not hold, the maxmin and minmax are equal. Of course, this is just an intuitive argument. Rigorous reasoning is given in the text.

minmax equals the maxmin since the functions are convex in the minimizing variable. We will soon show that the only way the minmax is not equal to the maxmin is if the value of each function at its minimizing Q_0 is less than the value of the other function at that same Q_0 (see Figure 11).

Although the region of such (δ, D) pairs is computed numerically, each point in the region can be verified analytically to show the suboptimality of symbol by symbol denoising.

Specifically, the program is used to compute the Q_0 that minimizes the exponent for each denoiser. This is easy to determine given the expression (29). We find the maximizing λ for (30) and then set Q_0 so the divergence term in (29) is zero. At this point, we need the following simple fact:

Claim 8: The minimizing Q_0 of the exponent for a particular denoiser is unique.

Proof: This follows from the fact that the divergence of two discrete distributions is zero if and only if the distributions are everywhere equal. ■

The minimum value of the exponent for a given denoiser is compared to the value of the exponent using the same Q_0 in the other denoiser. If the minimum exponent for each denoiser is strictly less than the exponent using that Q_0 in the other denoiser, then symbol-by-symbol is strictly suboptimal by the uniqueness of the minimizers of the exponents of the denoisers, the continuity of the exponent and the mean value theorem. These imply that the two exponent functions must cross for some value of Q_0 that lies strictly in between the minimizers of the two exponent functions. The value of the functions at this Q_0 is the minmax and is strictly greater than the maxmin since the minimizers of the exponents of the denoisers are unique.

The value of the supremum over λ is computed by taking a derivative and setting it equal to zero. We solve the

resulting equation numerically by using the roots function in Matlab. Since the functions are continuous in λ , the roots function will give an accurate value for λ . Each point in the region can be verified by computing the values of λ analytically and substituting into the corresponding rate function expressions.

3) *Sample Calculation of Symbol by Symbol Suboptimality:* We now provide an example of a particular channel and threshold where symbol by symbol denoising is suboptimal. In fact, we compute the maxmin and minmax values and show the strict inequality.

Use the above problem setup and set $\delta = 0.1$ and $D = 0.98$. Fixing the denoiser to be SWYS, (30) becomes:

$$\sup_{\lambda \geq 0} \lambda D - \log \left(\bar{\delta} + \frac{1}{2} \delta e^\lambda + \frac{1}{2} \delta e^{2\lambda} \right). \quad (53)$$

Differentiating with respect to λ and setting the result equal to zero, we get:

$$\delta \left(\frac{D}{2} - 1 \right) e^{2\lambda} + \delta \left(\frac{D-1}{2} \right) e^\lambda + D\bar{\delta} = 0. \quad (54)$$

Now, substituting the values for δ and D and scaling, we get:

$$-51e^{2\lambda} - e^\lambda + 882 = 0. \quad (55)$$

Thus, the expression is maximized by the nonnegative root of this equation, so $e^\lambda \cong 4.1488$ and $\lambda \cong 1.4228$. This yields a value of 0.7173 for the objective. Now, as described above, (29) is minimized by setting

$$\begin{aligned} Q_0 &= \frac{\sum_x p(x, 0) e^{\lambda \Lambda(X, \hat{x}(0))}}{E_{X,Z} e^{\lambda \Lambda(X, \hat{x}(Z))}} \\ &\cong 0.6659. \end{aligned}$$

We now compute the value of the objective for this value of Q_0 in the ONES denoiser. We use (23) to get:

$$\sup_{\lambda \geq 0} \lambda D - Q_0 \log Q_0 - Q_0 \log \left(\bar{p} \bar{\delta} e_p^\lambda \delta \right) + Q_1 \log Q_1 - Q_1 \log \left(\bar{p} \delta e^\lambda + p \bar{\delta} \right). \quad (56)$$

We take a derivative, set it to zero, set $p = 0.5$, and simplify to get:

$$(\delta \bar{\delta}) (D-1) e^{2\lambda} + (D \delta^2 + D \bar{\delta}^2 - Q_0 \bar{\delta}^2 - Q_1 \delta^2) e^\lambda + D \delta \bar{\delta} = 0. \quad (57)$$

We substitute the values for δ , D , and Q_0 to get:

$$-18e^{2\lambda} + 2609e^\lambda + 882 = 0. \quad (58)$$

So that $e^\lambda \cong 145.28$, $\lambda = 4.9787$ and the objective is 0.7754. This is greater than the value of the objective for the SWYS denoiser.

Now, we find the optimal value of (30) for the ONES denoiser. We have:

$$\sup_{\lambda \geq 0} \lambda D - \log \left(p + \bar{p} e^\lambda \right). \quad (59)$$

Differentiating with respect to λ and substituting $p = 0.5$ leads to

$$e^\lambda = \frac{D}{1-D}. \quad (60)$$

That is, $e^\lambda = 49$ and $\lambda \cong 3.8918$. Then, the objective is 0.5951. Furthermore, $Q_0 = 0.884$. We now compute the value of (29) for the SWYS denoiser and this value of Q_0 . (29) becomes:

$$\sup_{\lambda \geq 0} \lambda D + Q_0 \log Q_0 - Q_0 \log (\bar{p}\bar{\delta} + p\delta e^{2\lambda}) - Q_1 \log (p\bar{\delta} + p\delta e^\lambda) + Q_1 \log Q_1. \quad (61)$$

Differentiating with respect to λ and substituting $p = 0.5$ leads to

$$(D\delta^2 - Q_0\delta^2 - \delta^2) e^{3\lambda} + (D\delta\bar{\delta} - 2Q_0\delta\bar{\delta}) e^{2\lambda} + (D\delta\bar{\delta} - Q_1\delta\bar{\delta}) e^\lambda + D\bar{\delta}^2 = 0. \quad (62)$$

Substituting the δ , D and Q_0 values and simplifying yields:

$$-90.4e^{3\lambda} - 709.2e^{2\lambda} + 777.6e^\lambda + 7938 = 0. \quad (63)$$

Thus, the objective is maximized by $e^\lambda = 3.2306$, i.e. $\lambda = 1.1727$. The value of the objective is 0.8727, which is greater than that for the ONES denoiser.

By the results of this section, since the minimum value of the rate function of each denoiser is less than that of the other denoiser for the minimizing Q_0 , the rate functions must cross and thus $\max\min < \min\max$. So, we have a concrete example of the suboptimality of symbol-by-symbol denoising schemes.

ACKNOWLEDGMENT

The authors would like to thank Prof. Robert Gray, James Mammen and Young-Han Kim for helpful discussions.

REFERENCES

- [1] T. Weissman, E. Ordentlich, G. Seroussi, S. Verdú, and M. Weinberger, "Universal discrete denoising: Known channel," *IEEE Trans. Inform. Theory*, vol. 51, no. 1, pp. 5–28, January 2005.
- [2] E. Ordentlich and T. Weissman, "On the optimality of symbol by symbol filtering and denoising," *IEEE Trans. Inform. Theory*, vol. 52, no. 1, pp. 19–40, January 2006.
- [3] K. Marton, "Error exponent for source coding with a fidelity criterion," *IEEE Trans. Inform. Theory*, vol. 20, pp. 197–199, March 1974.
- [4] T. Weissman and N. Merhav, "Tradeoffs between the excess-code-length exponent and the excess-distortion exponent in lossy source coding," *IEEE Trans. Inform. Theory*, vol. 48, no. 2, pp. 396–415, February 2002.
- [5] N. Merhav and I. Kontoyiannis, "Source coding exponents for zero-delay coding with finite memory," *IEEE Trans. Inform. Theory*, vol. 49, pp. 609–624, March 2003.
- [6] B. R. M. Gastpar and M. Vetterli, "To code or not to code: Lossy source-channel communication revisited," *IEEE Trans. Inform. Theory*, vol. 49, no. 5, pp. 1147–1158, May 2003.
- [7] T. Weissman, "Universally attainable error-exponents for rate-distortion coding of noisy sources," *IEEE Trans. Inform. Theory*, vol. 50, no. 6, pp. 1229–1246, June 2004.
- [8] A. Puhalskii and V. Spokoiny, "On large deviation efficiency in statistical inference," *Bernoulli*, vol. 4, no. 2, pp. 203–272, 1998.
- [9] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [10] I. Csiszar and J. Korner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. New York: Academic Press, 1981.
- [11] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. New York: Springer-Verlag, 1998.