The Dispersion of the Mean Excess Distortion

Yuval Kochman
Hebrew University of Jerusalem
Jerusalem, Israel
Email: yuvalko@cs.huji.ac.il

Gregory W. Wornell
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: gww@mit.edu

***Abstract*—The problem of finite-blocklength lossy compression is considered. Motivated by troubling behavior of the rate expressions under an excess-distortion probability constraint, we define the mean excess distortion criterion. We evaluate the asymptotic performance of various settings under this criterion, and show that sharp and insightful rate bounds can be derived.**

I. Introduction

Excess-distortion analysis is a powerful tool for understanding the finite-blocklength behaviour of lossy compression. If a source block x^n is reconstructed as y^n, and under a single-letter distortion measure $d(x, y)$, an excess-distortion event $\mathcal{E}(n, D)$ happens when the empirical distortion

$$d(x^n, y^n) = \frac{1}{n}\sum_{i=1}^{n} d(x_i, y_i) \tag{1}$$

exceeds some prescribed threshold D; the excess-distortion probability $p_e(n, D)$ is the probability of this event. The reason for considering excess distortion is that from an operational point of view, the mean is not enough; for example, the fact that one source block is very accurately described may not compensate for another block that is distorted to the extent that it cannot be used. More generally, such considerations can be taken into account by applying any vector distortion measure to source blocks. However, the fact that the excess distortion is constructed from a scalar measure using the simple additive relation (1) maintains tractability.

The excess-distortion probability was introduced by Marton in [1], where its exponential behavior for fixed-rate compression was derived. Much more recently, the second-order (dispersion) behavior was found in [2], [3]: for fixed excess-distortion probability ϵ and blocklength n, the required rate of fixed-rate coding is given by:

$$R = R(D) + \sqrt{\frac{V(D)}{n}}Q^{-1}(\epsilon) + O\left(\frac{\log n}{n}\right), \tag{2}$$

where $R(D)$ and $V(D)$ are the source rate-distortion function (RDF) and dispersion, respectively, and $Q^{-1}(\cdot)$ is the inverse Gaussian tail probability function. This analysis extends to joint source-channel coding [4], [5] as well as to some network settings. In [6], it is extended the case where the quantizer dimension, and the rate over which the fixed-rate constraint is measured, may differ from n.

We are motivated by a recent work by Kostina et al. [7] where the basic fixed excess-distortion probability was studied under an *average* rate constraint. The required rate was found to differ fundamentally from (2):

$$R \cong (1-\epsilon)R(D) - \sqrt{\frac{V(D)}{n}}\Phi\left(Q^{-1}(\epsilon)\right), \tag{3}$$

where $\Phi(\cdot)$ is the Guassian probability density function. The main $(1-\epsilon)$ rate gain is due to allocating zero rate in case of an excess-distortion event. The second-order term is always negative, reflecting a further gain; this is since one can choose to have the excess-distortion in case of a "complicated" source block that requires high rate, thus the expected required rate given no excess distortion is lower than $R(D)$. A troubling phenomenon, observed in [7], is that the asymptotic rate is approached from below: the smaller the blocklength n, the bigger the second-order rate reduction. Could it be that it is *desirable* to work with small blocks?

It should be noted, that a similar phenomenon occurs in the context of lossless source coding under an average rate constraint [7]. In this *digital* case, it may not come as a big surprise, a longer blocklength n means a more stringent requirement of correctly reconstructing a larger amount of data.[1] However, distortion does satisfy the additive relation (1); thus with growing blocklength it retains the same average, while the variance should decrease thanks to statistical averaging. It is only the act of

[1]Indeed, a longer blocklength also means a richer family of encoders and decoders, which is favorable, although apparently less significant than the more strict requirement; in the case of channel coding under an average cost constraint [8], the situation is reversed and shorter blocklength means worse performance. Making the distinction between different blocklengths as in [6] allows to decouple these effects.

978-1-5090-1090-5/16/$31.00 ©2016 IEEE

translating the "analog" distortion into the binary excess-distortion event that may cause performance degradation with increasing blocklength.

We would like, then, to have a measure that reflects finite-blocklength behavior and is tractable (i.e., building on a scalar measure) as the excess-distortion probability, but has the property that it is easier to achieve as blocklength grows. To that end, we introduce the (soft) excess distortion for threshold D:

$$\delta(n, D) = [d(x^n, y^n) - D]_+, \tag{4}$$

where $[\cdot]_+$ denotes truncation of negative values. In the rest of this paper we study the rate required in order obtain a mean excess distortion $E[\delta(n, D)] \leq \delta$ for some fixed D and δ, as a function of the blocklength n. After giving some definitions and background, we turn in Section III to the average-rate problem of [7]. Then in Section IV we consider the "classical" fixed-rate setting, before extending the view in Section V to a multiple-blocklength scenario following [6].

II. Definitions and Background

We consider a memoryless source which takes values in some finite alphabet of size M, and some single-letter distortion measure $d(\cdot, \cdot) \leq d_{\max} < \infty$. For any distribution Q over this alphabet and any distortion level $\tilde{D} \geq 0$, the RDF is given by:

$$R(Q, \tilde{D}) = \min_{W: E_{Q,W} d(X,Y) \leq \tilde{D}} I(Q, W).$$

Further, we say that a distribution P and distortion level D form a regular point if the RDF is twice differentiable w.r.t. Q and differentiable w.r.t. $\tilde{D}$ in a neighborhood of (P, D);[2] we will only be interested in such points, and we denote the slope w.r.t. D by

$$\lambda = -\left. \frac{\partial R(Q, \tilde{D})}{\partial \tilde{D}} \right|_{\tilde{D}=D}. \tag{5}$$

For a regular point, the source dispersion $V(P, D)$ is well-defined and finite; for alternative expressions for this quantity, see [2], [3]. When the identity of the source is clear from the context, we write $R(D)$ and $V(D)$.

A fixed-rate code of blocklength n and rate R assigns to any source sequence X^n an index in $1, \ldots, 2^{nR}$. A variable-rate code assigns a codeword of length $R(X^n)$ bits where the set of all possible codewords is prefix-free.; the average rate of the code is the mean of $R(X^n)$ w.r.t. the source distribution.

[2]The neighborhood contains distributions outside the simplex, this is of no concern.

We use the method of types; specifically, the type class T_Q is the set of all n-length sequences that have type Q, where the blocklength is supressed.

We also use Gaussian distributions. Beyond the density $\Phi(x)$ and the tail probability $Q(x)$, we use:

$$K(x) = \int_x^\infty Q(x')dx' = \Phi(x) - xQ(x).$$

A. Optimal Type-Dependent Coding

In the following we consider a code where both rate and distortion are allowed to depend on the source realization.

Proposition 1: Consider a source which emits exchangeable sequences[3] over an alphabet size M. Consider a variable-rate code that allocates to each sequence a rate $R(X^n)$. Let $0 \leq D(Q)$ be a finite function of distributions. For any finite M there exist finite $a(M)$ and $b(M)$, s.t.:

1) Direct: There exists a code, such that for any $X^n \in T_Q$, $d(X^n, Y^n) \leq D(Q)$ and
$$R(X^n) \leq R(Q, D(Q)) + a(M)\frac{\log n}{n}$$
2) Converse: For any code, and for any type Q, if
$$E[R(X^n)|X^n \in T_Q] \leq R(Q, D(Q)) - b(M)\frac{\log n}{n}$$
then $E[d(X^n, Y^n)|X^n \in T_Q] \geq D(Q)$.

A few remarks are in place.

1) Note that the direct is given in terms of maximum per-type rate and distortion, while the converse is given in terms of average ones.
2) The direct part can be seen as a "type covering" code, plus a prefix that describes the type; by the polynomial number of types, this prefix is included in the logarithmic redundancy term. The converse part holds even if the source is uniformly distributed over a single type class. See e.g. [2].
3) Taking $D(Q)$ constant results in a fixed-distortion code. Similarly, taking $R(X^n)$ constant results in a fixed-rate code.

B. A simple source

Throughout the paper we use the following example, which suffices to capture much of the essence of the results. It is a version of "erasure" source/distortion. The source and reproduction are both over the ternary alphabet $\{0, 1, E\}$. The source is 0 or 1 with probability

[3]That is, the probability of a sequence is invariant to permutations, e.g., an i.i.d. source.

$p/2$, E with probability $1-p$. The distortion measure is: $d(x,y)=1$ if $x \neq E$ and $y=E$, $d(x,y)=0$ if $x=E$ or $y=x$, infinite distortion otherwise. That is, source symbols can be "erased" with unit cost, except when the source has the "don't care" value E, then it never suffers from distortion.

The rate-distortion function is simply

$$R(p,D)=p-D.$$

That is, it takes 1 bit to accurately describe a source symbol that matters. It is the linearity of this function that makes this source convenient; in the general case, we will need linearizations.

For this source, it is convenient to consider "super type classes". That is, let the E-type class T_q contain all sequences that contain qn E-symbols, regardless of the number of zeros and ones. As the probability of sequences within T_q that are far from type $(q/2, q/2, 1-q)$ is exponentially small, Proposition 1 holds with respect to T_q. That is, the type-dependent rate $R(q)$ needed to achieve a type-dependent distortion $D(q)$ in the sense of Proposition 1 is:

$$R(q)=q-D(q)+O\left(\frac{\log n}{n}\right). \tag{6}$$

Also, it is not difficult to verify that the source dispersion is only "with respect to" the E-symbols, that is, it is the variance of a Bernoulli-p variable: $V(p,D)=p(1-p)$.

III. Average Rate Constraint

The following shows that for an average rate constraint, excess distortion is a "zero-sum game": there is no loss of generality in restricting attention to expected distortion. Indeed, with $\delta=0$ the following result reduces to the expected distortion redundancy [9].

Theorem 1: Let R_n be the required expected rate at blocklength n for any distortion threshold $D \geq 0$ and mean excess-distortion $\delta \geq 0$. Then

$$R_n=R(P,D+\delta)+O\left(\frac{\log n}{n}\right).$$

Proof outline: First consider the erasure source. For achievability, we cover all source sequences with fixed distortion $\tilde{D}$. For a source realization with E-type q, The rate needed is $R(q)$ of (6). Thus the expected rate is:

$$E[R(q)]=E[q]-\tilde{D}+O\left(\frac{\log n}{n}\right)=p-\tilde{D}+O\left(\frac{\log n}{n}\right).$$

Since the distortion is fixed, we may take $\tilde{D}=D+\delta$ for any $\delta \geq 0$. For the converse, by Jensen's inequality for the convex $[x]_+$,

$$D+\delta \geq E[d(X^n,Y^n)]$$

and then:

$$\begin{aligned} E[d(X^n,Y^n)] &= E[E[d(X^n,Y^n)|q]] \\ &\geq E\left[q-R(q)+O\left(\frac{\log n}{n}\right)\right] \\ &= p-R+O\left(\frac{\log n}{n}\right), \end{aligned}$$

where for the inequality we have used again (6).

For generalizing to any source, one can define a Q-neighborhood of P that has vanishing radius, and at the same time very high probability. Within that neighborhood $R(Q,D)$ is approximately linear in Q, with mean $R(P,D)$. Finally, Jensen's inequality is applied to $R(Q,D)$ which is convex in D, in order to assert that it is indeed optimal to have $E[d(X^n,Y^n)|Q]$ independent of Q.

IV. Fixed Rate Constraint

We now turn to the "classical" setting of fixed-rate quantization. Let us first interpret the excess-distortion probability dispersion (2). Consider the erasure source. A fixed-rate code of rate R_n can achieve for a given E-type q a distortion:[4]

$$D(q)=q-R_n+O\left(\frac{\log n}{n}\right). \tag{7}$$

The excess-distortion probability is thus approximately the probability of $D(q)$ to exceed the threshold D:

$$\epsilon \cong \Pr\left[q>D+R_n+O\left(\frac{\log n}{n}\right)\right].$$

By the CLT, q is approximately normal with mean p and variance $p(1-p)/n$. The correction to the CLT, given by the Berry-Essen theorem, only gives a correction of order $1/n$, thus we have:

$$R_n=p-D+\sqrt{\frac{p(1-p)}{n}}Q^{-1}(\epsilon)+O\left(\frac{\log n}{n}\right),$$

which is exactly (2) applied to the erasure source.

The same nearly optimal scheme can be applied to mean excess-distortion analysis as well. Since the distribution of q scales as $1/\sqrt{n}$, the mean excess distortion also scales as $1/\sqrt{n}$ when ϵ is fixed. Indeed, we use this normalization in the following.

Theorem 2: Let R_n be the required rate at blocklength n for distortion threshold D and mean excess-distortion $\delta=\delta_0/\sqrt{n}$. For any fixed $D, \delta_0 \geq 0$,

$$R_n=R(P,D)+\sqrt{\frac{V(P,D)}{n}}K^{-1}(\lambda\delta_0)+O\left(\frac{\log n}{n}\right).$$

[4] For excess-distortion probability, one uses a version of Proposition 1 where the second part states that with very high probability, the distortion will be above D.

Proof outline: First consider the erasure source. By Proposition 1 there exists a fixed-rate code of rate R_n, such that the maximum distortion given a type class satisfies (7), and there does not exist a code such that the average distortion is better than (7). Thus we have:

$$\frac{\delta_0}{\sqrt{n}} = E\left[\left[q - D - R_n + O\left(\frac{\log n}{n}\right)\right]_+\right].$$

Now, since for X that is Gaussian $(\mu, \sigma^2/n)$,

$$E\left[[X - \mu - x_0]_+\right] = \frac{1}{\sqrt{n}} \cdot K\left(\sqrt{\frac{n}{\sigma^2}} x_0\right),$$

it follows that under a Gaussian approximation for q,

$$\delta_0 = K\left(\sqrt{\frac{n}{p(1-p)}}\left(R_n - p + D + O\left(\frac{\log n}{n}\right)\right)\right),$$

which is equivalent to the required result. For applying the Gaussian approximation, the Berry-Essen theorem does not suffice (as the uniform gap to the Gaussian cumulative distribution function explodes in the calculation of the expected value), thus a stronger approximation is applied, as in the proof of [7, Lemma 1].

Generalization to any source follows by linearization, closely following [2]. Locally, we have

$$R(Q, D) = R(P, D) + \sum_{i=1}^{M} R_i'(Q_i - P_i),$$

where R_i' is the derivative of $R(Q, D)$ w.r.t. the i-th coordinate, evaluated at $P = Q$. The dispersion $V(P, D)$ is the variance of R_i when the source is distributed according to P. The RDF slope λ (5) serves to normalize from distortion to rate.

V. General Result

We now consider the setting defined and motivated in [6], where a distinction is made between three blocklengths:

1) Processing blocklength k: the dimension of the code.
2) Fidelity blocklength n: the dimension of the distortion measurement.
3) Resource blocklength m: the dimension over which the rate constraint is enforced.

For tractability we concentrate on a synchronized version of the problem, defined as follows. k, n and m are integers such that any two of them have an integer ratio. The order between them is arbitrary, except that $m \geq k$, that is, $m = Ak$ for an integer A. Using these blocklengths, the problem is defined as follows. Let

$$\overline{n} = \max(k, m, n) = \max(m, n) \tag{8a}$$

$$\underline{n} = \min(k, m, n) = \min(k, n). \tag{8b}$$

A source block $X_1^{\overline{n}}$ is processed by a coding scheme operating on blocks X_1^k, X_{k+1}^{2k}... and the reconstructions are pasted back together in the same order to create the reconstruction $Y_1^{\overline{n}}$. Let the rate of the r-th processing block be R_r. A sequence of A processing blocks is observed, and then a rate allocation is made, such that the following constraint is satisfied:[5]

$$\frac{1}{A} \sum_{r=(C-1)A+1}^{CA} R_r \leq R.$$

The mean excess distortion is given by:

$$\delta = \frac{1}{B} \sum_{f=1}^{B} E\left[\left[d\left(X_{(f-1)n+1}^{fn}, Y_{(f-1)n+1}^{fn}\right) - D\right]_+\right] \tag{9}$$

for $B = \overline{n}/n$. For these, we have the following.

Theorem 3: Let $R_{\overline{n},\underline{n}}$ be the required rate at blocklengths $\overline{n}$ and $\underline{n}$ for the synchronized setting defined above, and for distortion threshold D and mean excess-distortion $\delta = \delta_0/\sqrt{n}$. For any fixed $D, \delta_0 \geq 0$,

$$R_{\overline{n},\underline{n}} = R(P, D) + \sqrt{\frac{V(P, D)}{\overline{n}}} K^{-1}(\lambda \delta_0) + O\left(\frac{\log \underline{n}}{\underline{n}}\right).$$

Before giving a proof outline, let us discuss some consequences.

1) A similar expression was derived in [6] for the excess-distortion probability. However, in that work, for the case $m > k$, this was only an achievability result without a converse; the reason was that one could gain by allocating zero rate to some source blocks, as done in [7] for the case of average rate. Shifting to mean excess distortion allows to prove the converse result as well.
2) When all the blocklengths are equal ($k = n = m$), the result reduces to Theorem 2.
3) In general, we see that the second-order (dispersion) term is with respect to the maximum of the blocklengths. In this sense, he result is optimistic: if one starts from equal blocklengths, some of them can be reduced without hurting the second-order performance, until they are so small that $O(\log \underline{n}/\underline{n})$ is more significant than $O(1/\overline{n})$.
4) Taking the limit of infinite m or infinite n amounts to average distortion or average rate. Indeed it

[5] A scheme that allocates rates to the processing blocks sequentially, only keeping track of the cumulative rate consumption, can also be devised.

leads to a zero second-order term, reducing to the logarithmic redundancy of [9] and of Section III.

Proof outline: we will only treat the erasure source; the extension to other sources follows in a similar manner to Theorem 2. We make a distinction between two cases.

1) $\overline{n} = n \geq m \geq k = \underline{n}$. Since performance is monotonic in both m and k, we assume fixed-rate coding with rate $R_{\overline{n},\underline{n}}$ at blocklength ℓ, where for the direct $\ell = k = \underline{n}$ while for the converse $\ell = m$. Now, the whole source block is composed of $G = \overline{n}/\ell$ processing blocks. At each one, we have as in (7) a type-dependent distortion of:

$$D(q_a) = q_a - R_{\overline{n},\underline{n}} + O\left(\frac{\log \ell}{\ell}\right), \quad a = 1, \ldots, G$$

where q_a is the source E-type at the a-th block. For evaluating δ (9) we consider

$$\frac{1}{G}\sum_{a=1}^{G} D(q_a) = \frac{1}{G}\sum_{a=1}^{G} q_a - R_{\overline{n},\underline{n}} + O\left(\frac{\log \ell}{\ell}\right)$$
$$= q - R_{\overline{n},\underline{n}} + O\left(\frac{\log \ell}{\ell}\right),$$

where q is the E-type of the full source sequence of length $\overline{n}$. Now, following the Gaussian approximation as in the proof of Theorem 2 we have the desired result with the logarithmic term being a function of ℓ rather than $\underline{n}$; but since $\ell \geq \underline{n}$, the result follows.

2) $\overline{n} = m \geq n$. Since performance is monotonic in both n and k,[6] we assume fixed-distortion coding with distortion $\tilde{D}(q)$ to be determined (where q is the E-type of the whole source sequence), at block-length ℓ, where for the direct $\ell = \min(k, n) = \underline{n}$ while for the converse $\ell = \max(k, n)$. Now, the whole source block is composed of $G = \overline{n}/\ell$ processing blocks. At each one, we need a type-dependent rate of:

$$R(q_a) = q_a - \tilde{D}(q) + O\left(\frac{\log \ell}{\ell}\right), \quad a = 1, \ldots, G$$

where q_a is the source E-type at the a-th block. In order to satisfy the rate constraint we have:

$$R_{\overline{n},\underline{n}} = q - \tilde{D}(q) + O\left(\frac{\log \ell}{\ell}\right).$$

Since we used fixed distortion, we have:

$$\frac{\delta_0}{\sqrt{n}} = E\left[\left[\tilde{D}(q) - D\right]_+\right]$$
$$= E\left[\left[q - D - R_{\overline{n},\underline{n}} + O\left(\frac{\log \ell}{\ell}\right)\right]_+\right].$$

The proof is completed as in the first case.

[6] Monotonicity in n holds true since we consider mean excess distortion; if one tries to prove a similar converse for excess-distortion probability, this is the point where it fails.

VI. Discussion

We have presented the mean excess distortion as a measure for finite-blocklength performance that provides tractability while avoiding the anomalies encountered when using the excess-distortion probability. Of course, one may wonder if other measures would be suitable as well. Indeed, from a technical point of view the problem with the excess-distortion probability is that it is the mean of an indicator function, which is not convex; any convex function of the empirical distortion would be adequate. However, the linearity of the mean distortion gives the simplest results, and it seems that the freedom to choose the distortion measure gives enough flexibility in tailoring the analysis to the application.

References

[1] K. Marton. Error exponent for source coding with a fidelity criterion. *IEEE Trans. Info. Theory*, IT-20:197–199, Mar. 1974.

[2] A. Ingber and Y. Kochman. The dispersion of lossy source coding. In *Proc. of the Data Compression Conference*, Snowbird, Utah, March 2011.

[3] V. Kostina and S. Verdú. Fixed-length lossy compression in the finite blocklength regime. *IEEE Trans. Info. Theory*, 58(6):3309–3338, June 2012.

[4] I. Csiszár. On the error exponent of source-channel transmission with a distortion threshold. *IEEE Trans. Info. Theory*, IT-28:823–838, Nov. 1982.

[5] D. Wang, A. Ingber, and Y. Kochman. The dispersion of joint source-channel coding. In *49th Annual Allerton Conference on Communication, Control and Computing*, Monticelo, IL, Sep. 2011.

[6] Y. Kochman and G. W. Wornell. Excess distortion in lossy compression: Beyond one-shot analysis. In *53rd Annual Allerton Conference on Communication, Control and Computing*, Monticelo, IL, Sep. 2015.

[7] V. Kostina, Y. Polyanskiy, and S. Verdú. Variable-length compression allowing errors. *IEEE Trans. Info. Theory*, 61(8):4316–4330, 2015.

[8] W. Yang, G. Caire, G. Durisi, and Y. Polyanskiy. Optimum power control at finite blocklength. *IEEE Trans. Info. Theory*, 61(9):4598–4615, Sept 2015.

[9] Z. Zhang, E.H. Yang, and V. Wei. The redundancy of source coding with a fidelity criterion - Part one: Known statistics. *IEEE Trans. Info. Theory*, IT-43:71–91, Jan. 1997.