Shannon Theory for Compressed Sensing

Yihong Wu

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Electrical Engineering
Adviser: Professor Sergio Verdú

September 2011

Abstract

Compressed sensing is a signal processing technique to encode analog sources by real numbers rather than bits, dealing with efficient recovery of a real vector from the information provided by linear measurements. By leveraging the prior knowledge of the signal structure (e.g., sparsity) and designing efficient non-linear reconstruction algorithms, effective compression is achieved by taking a much smaller number of measurements than the dimension of the original signal. However, partially due to the non-discrete nature of the problem, none of the existing models allow a complete understanding of the theoretical limits of the compressed sensing paradigm.

As opposed to the conventional worst-case (Hamming) approach, this thesis presents a statistical (Shannon) study of compressed sensing, where signals are modeled as random processes rather than individual sequences. This framework encompasses more general signal models than sparsity. Focusing on optimal decoders, we investigate the fundamental tradeoff between measurement rate and reconstruction fidelity gauged by error probability and noise sensitivity in the absence and presence of measurement noise, respectively. Optimal phase transition thresholds are determined as a functional of the input distribution and compared to suboptimal thresholds achieved by various popular reconstruction algorithms. In particular, we show that Gaussian sensing matrices incur no penalty on the phase transition threshold of noise sensitivity with respect to optimal nonlinear encoding. Our results also provide a rigorous justification of previous results based on replica heuristics in the weak-noise regime.

Acknowledgements

I am deeply indebted to my advisor Prof. Sergio Verdú for his constant guidance and support at every stage of my Ph.D. studies, without which this dissertation would be impossible. A superb researcher and mentor, he taught me how to formulate a problem in the cleanest possible way. His brilliant insights, mathematical elegance and aesthetic taste greatly shaped my research. In addition, I appreciate his endless patience and the freedom he gave, which allow me to pursue my research interests that sometimes are well outside the convex hull of the conventional information theory.

I thank Prof. Sanjeev Kulkarni and Prof. Paul Cuff for serving as my thesis reader and their valuable feedbacks. I thank Prof. Shlomo Shamai and Prof. Phillipe Rigollet for being on my defense committee. I thank Prof. Mung Chiang and Prof. Robert Calderbank for serving as my oral general examiners. I am also grateful to Prof. Mung Chiang for being my first-year academic advisor and my second-year research co-advisor.

I would like to thank the faculty of the Department of Electrical Engineering, Department of Mathematics and Department of Operational Research and Financial Engineering for offering great curricula and numerous inspirations, especially, Prof. Robert Calderbank, Prof. Erhan Çinlar, Prof. Andreas Hamel, Prof. Elliot Lieb and Dr. Tom Richardson, for their teaching.

I am grateful to Dr. Erik Ordentlich and Dr. Marcelo Weinberger for being my mentors during my internship at HP labs, where I spent a wonderful summer in 2010. I would also like to acknowledge Prof. Andrew Barron, Prof. Dongning Guo, Dr. Jouni Luukkainen, Prof. Shlomo Shamai and Dr. Pablo Shmerkin for stimulating discussions and fruitful collaborations.

I have been fortunate to have many friends and colleagues who made my five years at Princeton an unforgettable experience: Zhenxing Wang, Yiyue Wu, Zhen Xiang, Yuejie Chi, Chunxiao Li, Can Sun, Yu Yao, Lin Han, Chao Wang, Yongxin Xi, Guanchun Wang, Tiance Wang, Lorne Applebaum, Tiffany Tong, Jiming Yu, Qing Wang, Ying Li, Vaneet Aggarwal, Aman Jain, Victoria Kostina among them. In particular, I thank Yury Polyanskiy for numerous rewarding discussions, during every one of which I gained new knowledge. I thank Ankit Gupta for his philosophical teaching and endless wisdom. I thank Haipeng Zheng for a decade of true friendship. My dearest friends outside Princeton, Yu Zhang, Bo Jiang, Yue Zhao, Sheng Zhou, are always there for me during both joyful and stressful times, to whom I will always be thankful.

Finally, it is my greatest honor to thank my family: my mother, my father, my grandparents, my uncle and auntie. No words could possibly express my deepest gratitude for their love, self-sacrifice and unwavering support. To them I dedicate this dissertation.

To my family.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Compressed sensing

Current data acquisition methods are often extremely wasteful: massive amounts of data are acquired then discarded by a subsequent compression stage. In contrast, a recent approach known as compressed sensing [1, 2], combines the data collection and compression and reduces enormously the volume of data at much greater efficiency and lower costs. A paradigm of encoding of analog sources by real numbers rather than bits, compressed sensing deals with efficient recovery of sparse vectors from the information provided by linear measurements. By leveraging the prior knowledge of the signal structure (e.g., sparsity) and designing efficient non-linear reconstruction algorithms, effective compression is achieved by taking a much smaller number of measurements than the dimension of the original signal.

An abstract setup of compressed sensing is shown in Fig. 1.1: A real vector $x^n \in \mathbb{R}^n$ is mapped into $y^k \in \mathbb{R}^k$ by an encoder (or compressor) $f : \mathbb{R}^n \to \mathbb{R}^k$. The decoder (or decompressor) $g : \mathbb{R}^k \to \mathbb{R}^n$ receives $\hat{y}^k$, a possibly noisy version of the measurement, and outputs $\hat{x}^n$ as the reconstruction. The *measurement rate*, i.e., the number of measurements per sample or the dimensionality compression ratio, is given by

$$R = \frac{k}{n}. \tag{1.1}$$

The goal is to reconstruct the unknown vector x^n as faithfully as possible. In the classical setup of compressed sensing, the encoder f is constrained to be a linear mapping characterized by a $k \times n$ matrix $\mathbf{A}$, called the *sensing* or *measurement matrix*. Furthermore, $\mathbf{A}$ is usually assumed to be a random matrix, for example, consisting of i.i.d. entries or randomly sampled Fourier coefficients. In order to guarantee stable reconstruction in the presence of measurement inaccuracy, we usually also require the decoder to be robust, i.e., *Lipschitz continuous*. The fundamental question is: under practical constraints such as linearity of the encoder and/or robustness of the decoder, what is the minimal measurement rate so that the input vector can be accurately reconstructed?

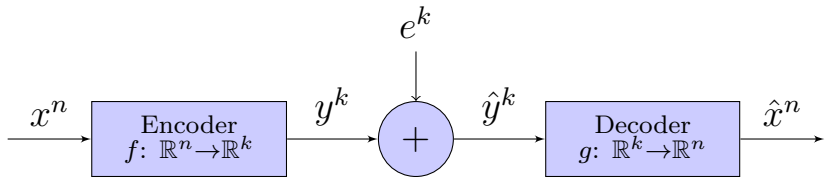

Figure 1.1: Compressed sensing: an abstract setup.

As opposed to the conventional worst-case (Hamming) approach, in this thesis we introduce a statistical (Shannon) framework for compressed sensing, where signals are modeled as random processes rather than individual sequences and performance is measured on an average basis. This setup can also be understood as a Bayesian formulation of compressed sensing where the input distribution serves as the *prior* of the signal. Our goal is to seek the fundamental limit of the measurement rate as a function of the input statistics. Similar approaches have been previously adopted in the literature, for example, [3, 4, 5, 6, 7, 8, 9, 10, 11].

In addition to linear encoders, we are also concerned with non-linear encoders mainly for the following reasons. First, studying the fundamental limits of nonlinear encoders provides converse results (lower bound on the measurement rate) for linear encoding. Second, it enables us to gauge the suboptimality incurred by linearity constraints on the encoding procedures.

1.2 Phase transition

When the measurements are noiseless, our goal is to reconstruct the original signal as perfectly as possible by driving the error probability to zero as the ambient dimension n grows. For many input processes (e.g., memoryless), it turns out that there exists a threshold for the measurement rate, above which it is possible to achieve a vanishing error probability and below which the error probability will eventually approaches one for any sequence of encoder-decoder pairs. Such a phenomenon is known as *phase transition* in statistical physics. In information-theoretic terms, we say that *strong converse* holds.

When the measurement is noisy, exact signal recovery is obviously impossible. Instead, our objective is to pursue robust reconstruction. To this end, we consider another benchmark called the *noise sensitivity*, defined as the ratio between mean-square reconstruction error and the noise variance. Similar to the behavior of error probability in the noiseless case, there exists a phase transition threshold of measurement rate which only depends on the input statistics, above which noise sensitivity is bounded for all noise variance, and below which the noise sensitivity blows up as noise variance tends to zero.

Clearly, to every sequence of encoders and decoders we might associate a phase transition threshold. In our paper we will be primarily concerned with *optimal* decoding procedures, as opposed to pragmatic low-complexity algorithms. For instance, the optimal *decoder* in the noisy case is the minimum mean-square error (MMSE)

estimator, i.e., the conditional expectation of the input vector given the noisy measurements. In view of the prohibitive computational complexity of optimal decoders, various practical suboptimal reconstruction algorithms have been proposed, most notably, decoders based on convex optimizations such as ℓ_1-minimization (i.e., Basis Pursuit) [12] and ℓ_1-penalized least-squares (i.e. LASSO) [13], greedy algorithms such as matching pursuit [14], graph-based iterative decoders such as approximate message passing (AMP) [4], etc. These algorithms in general have strictly higher phase-transition thresholds. In addition to optimal decoders, we also consider the case where we have the freedom to optimize over measurement matrices.

In the case of noiseless compressed sensing, the interesting regime of measurement rates is between zero and one. When the measurements are noisy, in principle it makes sense to consider measurement rate greater than one in order to combat the noise. Nevertheless, the optimal phase transition for noise sensitivity is always less than one, because with $k = n$ and an invertible measurement matrix, the linear MMSE estimator achieves bounded noise sensitivity for any noise variance.

1.3 Signal models: sparsity and beyond

Sparse vectors, supported on a subspace with dimension much smaller than the ambient dimension, are of great significance in signal processing and statistical modeling. One suitable probabilistic model for sparse signals is the following mixture distribution [5, 15, 4, 7]:

$$P = (1-\gamma)\delta_0 + \gamma\, P_{\mathrm{c}}, \tag{1.2}$$

where δ_0 denotes the Dirac measure at 0, P_{c} is a probability measure absolutely continuous with respect to the Lebesgue measure, and $0 \le \gamma \le 1$. Consider a random vector $X^n = [X_1, \ldots, X_n]^{\mathrm{T}}$ independently drawn from P. By the weak law of large numbers, if $\gamma < 1$, X^n is sparse with overwhelming probability. More precisely, $\frac{1}{n}\left\|X^n\right\|_0 \xrightarrow{\mathbb{P}} \gamma$, where the "$\ell_0$ norm"[1] $\left\|x^n\right\|_0 = \sum_{i=1}^n \mathbf{1}_{\{|x_i|>0\}}$ denotes the support size of a vector. This corresponds to the regime of *proportional (or linear) sparsity*. In (1.2), the weight on the continuous part γ parametrizes the signal sparsity: less γ corresponds to more sparsity, while P_{c} serves as the prior distribution of non-zero entries.

Apart from sparsity, there are other signal structures that have been previously explored in the compressed sensing literature [15, 16]. For example, the so-called *simple* signal in infrared absorption spectroscopy [16, Example 3, p. 914] is such that each entry of of the signal vector is constrained in the unit interval, with most of the entries saturated at the boundaries (0 or 1). Similar to the rationale that leads to (1.2), an appropriate statistical model for simple signals is a memoryless input process drawn from

$$P = (1-\gamma)\Big(p\delta_0 + (1-p)\delta_1\Big) + \gamma P_{\mathrm{c}}, \tag{1.3}$$

[1]Strictly speaking, $\|\cdot\|_0$ is not a norm because $\|\lambda x^n\|_0 \neq |\lambda|\,\|x^n\|_0$ for $\lambda \neq 0$ or ± 1. However, $\|x-y\|_0$ gives a valid metric on $\mathbb{R}^n$.

where P_c is an absolutely continuous probability measure supported on the unit interval and $0 \leq p \leq 1$.

Generalizing both (1.2) and (1.3), we consider *discrete-continuous mixed* distributions (i.e., *elementary distributions* in the parlance of [17]) of the following form:

$$P = (1-\gamma)P_d + \gamma P_c, \tag{1.4}$$

where P_d and P_c are discrete and absolutely continuous respectively. Although most of our results in this thesis hold for arbitrary input distributions, we will be focusing on discrete-continuous mixtures because of its relevance to the compressed sensing applications.

1.4 Three dimensions

There are three dimension concepts that are involved in various coding theorems throughout this thesis, namely the *information dimension*, the *Minkowski dimension*, and the *MMSE dimension*, summarized in Chapters 2, 3 and 4, respectively. They all measure how fractal subsets or probability measures are in Euclidean spaces in certain senses.

In fractal geometry, the Minkowski dimension (also known as the box-counting dimension) [18] gauges the fractality of a subset in metric spaces, defined as the exponent with which the covering number grows. The relevance of Minkowski dimension to compressed sensing are twofold: First, subsets of low Minkowski dimension can be linearly embedded into low-dimensional Euclidean spaces [19, 20, 21].[2] This result is further improved and extended to a probabilistic setup in Chapter 5. Second, Minkowski dimension never increases under Lipschitz mappings. Consequently, knowledge about Minkowski dimension provides *converse* results for analog compression when the decompressor are constrained to be Lipschitz.

Another key concept in fractal geometry is the *information dimension* (also known as the *entropy dimension* [23]) defined by Alfréd Rényi in 1959 [24]. It measures the rate of growth of the entropy of successively finer discretizations of a probability distribution. Although a fundamental information measure, the Rényi dimension is far less well-known than either the Shannon entropy or the Rényi entropy. Rényi showed that, under certain conditions, the information dimension of an absolutely continuous n-dimensional probability distribution is n. Hence he remarked in [24] that "*... the geometrical (or topological) and information-theoretical concepts of dimension coincide for absolutely continuous probability distributions*". However, as far as we know, the operational role of Rényi information dimension has not been addressed before except in the work of Kawabata and Dembo [25] and Guionnet and Shlyakhtenko [26], which relates it to the rate-distortion function and mutual information, respectively. In this thesis we provide several new operational characterizations of information dimension in the context of lossless compression of analog sources. In Chapter 3 we show that information dimension is connected to how product measures are concen-

[2] This connection has also been recognized in [22].

trated on subsets of low Minkowski dimension. Coupled with the linear embeddability result mentioned above, this connection suggests the almost lossless compressibility of memoryless sources via linear encoders, which makes noiseless compressed sensing possible. Moreover, Chapter 5 relates information dimension to the fundamental limit of almost lossless compression of analog sources under various regularity constraints of the encoder/decoder.

The MMSE dimension is an information measure for probability distributions that governs the asymptotics of the minimum mean-square error (MMSE) of estimating a random variable corrupted by additive Gaussian noise when the signal to noise ratio (SNR) or the sample size is large. This concept is closely related to information dimension, and will be useful when we discuss noisy compressed sensing in Chapter 6.

1.5 Main contributions

Compressed sensing, as an analog compression paradigm, imposes two basic requirements: the linearity of the encoder and the robustness of the decoder; the rationale is that low complexity of encoding operations and noise resilience of decoding operations are indispensable in dealing with analog sources. To better understand the fundamental limits imposed by the requirements of low complexity and noise resilience, it is pedagogically sound to study them both separately and jointly as well as in a more general framework than compressed sensing.

Motivated by this observation, in Chapter 5 we introduce the framework of almost lossless analog compression as a Shannon-theoretic formulation of noiseless compressed sensing. Abstractly, the approach boils down to probabilistic dimension reduction with smooth embedding. Under regularity conditions on the encoder or the decoder, various coding theorems for the minimal measurement rate are derived involving the information dimension of the input distribution. The most interesting regularity constraints are the linearity of the compressor and Lipschitz continuity (robustness) of the decompressor. The fundamental limits when these two constraints are imposed either separately or simultaneously are all considered. For memoryless discrete-continuous mixtures that are of the form (1.4), we show that the minimal measurement rate is given by the input *information dimension*, which coincides with the weight γ of the absolutely continuous part. Moreover, the Lipschitz constant of the decoder can be chosen independent of n, as a function of the gap between the measurement rate and γ. This resolves the optimal phase transition threshold of error probability in noiseless compressed sensing.

Chapter 6 deals with the noisy case where the measurements are corrupted by additive noise. We consider three formulations of noise sensitivity, corresponding to optimal nonlinear, optimal linear and random linear (with i.i.d. entries) encoders and the associated optimal decoders, respectively. In the case of memoryless input processes, we show that for any input distribution, the phase transition threshold for optimal encoding is given by the input information dimension. Moreover, this result also holds for discrete-continuous mixtures with optimal linear encoders and Gaussian random linear encoders. Invoking the results on MMSE dimension in Chapter 4,

we show that the calculation of the reconstruction error with random measurement matrices based on *heuristic* replica methods in [7] predicts the correct phase transition threshold. These results also serve as a rigorous verification of the replica calculations in [7] in the high-SNR regime (up to $o(\sigma^2)$ as the noise variance σ^2 vanishes).

In Chapter 7, we compare the optimal phase transition threshold to the suboptimal threshold of several practical reconstruction algorithms under various input distributions. In particular, we show that the thresholds achieved by decoders based on linear-programming [16] and the AMP decoder [15, 6] lie far from the optimal boundary, especially in the highly sparse regime that is most relevant to compressed sensing applications.

Chapter 2

Rényi information dimension

This chapter is devoted to an overview of Rényi information dimension and its properties. In particular, in Section 2.5 we give a novel interpretation in terms of the entropy rate of the dyadic expansion of the random variable. We also discuss the connection between information dimension and rate-distortion theory established by Kawabata and Dembo [25], as well as the high-SNR asymptotics of mutual information with additive noise due to Guionnet and Shlyakhtenko [26]. The material in this chapter has been presented in part in [27, 5, 28, 29]; Theorems 1 – 3 and the content in Section 2.9 are new.

2.1 Definitions

A key concept in fractal geometry, in [24] Rényi defined the *information dimension* (also known as the *entropy dimension* [23]) of a probability distribution. It measures the rate of growth of the entropy of successively finer discretizations.

Definition 1 (Information dimension)**.** Let X be an arbitrary real-valued random variable. For $m \in \mathbb{N}$, denote the uniformly quantized version of X by[1]

$$\langle X \rangle_m \triangleq \frac{\lfloor mX \rfloor}{m}. \tag{2.1}$$

Define

$$\underline{d}(X) = \liminf_{m\to\infty} \frac{H\left(\langle X \rangle_m\right)}{\log m} \tag{2.2}$$

and

$$\overline{d}(X) = \limsup_{m\to\infty} \frac{H\left(\langle X \rangle_m\right)}{\log m}, \tag{2.3}$$

where $\underline{d}(X)$ and $\overline{d}(X)$ are called lower and upper information dimensions of X respectively. If $\underline{d}(X) = \overline{d}(X)$, the common value is called the information dimension of

[1] Throughout this thesis, $\langle \cdot \rangle_m$ denotes the quantization operator, which can also be applied to real numbers, vectors or subsets of $\mathbb{R}^n$ componentwise.

X, denoted by $d(X)$, i.e.

$$d(X) = \lim_{m\to\infty} \frac{H\left(\langle X\rangle_m\right)}{\log m} \tag{2.4}$$

Rényi also defined the "entropy of dimension $d(X)$" as

$$\hat{H}(X) = \lim_{m\to\infty} \left(H(\langle X\rangle_m) - d(X)\log m\right), \tag{2.5}$$

provided the limit exists.

Definition 1 can be readily extended to random vectors, where the floor function $\lfloor\cdot\rfloor$ is taken componentwise [24, p. 208]. Since $d(X)$ only depends on the distribution of X, we also denote $d(P_X) = d(X)$. Similar convention also applies to entropy and other information measures.

Apart from discretization, information dimension can be defined from a more general viewpoint: the mesh cubes of size $\epsilon > 0$ in $\mathbb{R}^k$ are the sets $C_{z,\epsilon} = \prod_{j=1}^k [z_j\epsilon, (z_j+1)\epsilon)$ for $z^k \in \mathbb{Z}^k$. For any $\epsilon > 0$, the collection $\{C_{z,\epsilon} : z \in \mathbb{Z}^k\}$ partitions $\mathbb{R}^k$. Hence for any probability measure μ on $\mathbb{R}^k$, this partition generates a discrete probability measure μ_ϵ on $\mathbb{Z}^k$ by assigning $\mu_\epsilon(\{z\}) = \mu(\{C_{z,\epsilon}\})$. Then the information dimension of μ can be expressed as

$$d(X) = \lim_{\epsilon\downarrow 0} \frac{H\left(\mu_\epsilon\right)}{\log\frac{1}{\epsilon}}. \tag{2.6}$$

It should be noted that there exist alternative definitions of information dimension in the literature. For example, in [30] the lower and upper information dimensions are defined by replacing $H\left(\mu_\epsilon\right)$ with the ϵ-entropy $H_\epsilon\left(\mu\right)$ [31, 32] with respect to the ℓ_∞ distance. This definition essentially allows unequal partition of the whole space, since $H_\epsilon\left(\mu\right) \le H\left(\mu_\epsilon\right)$. However, the resulting definition is equivalent (see Theorem 8). As another example, the following definition is adopted in [33, Definition 4.2],

$$d(X) = \lim_{\epsilon\downarrow 0} \frac{\mathbb{E}\left[\log \mu(B_p(X,\epsilon))\right]}{\log\epsilon} \tag{2.7}$$

$$= \lim_{\epsilon\downarrow 0} \frac{\mathbb{E}\left[\log \mathbb{P}\left\{X \in B_p(\bar{X},\epsilon)|\bar{X}\right\}\right]}{\log\epsilon}, \tag{2.8}$$

where μ denotes the distribution of X, $\bar{X}$ is an independent copy of X and $B_p(x,\epsilon)$ is the ℓ_p-ball[2] of radius ϵ centered at x, with the ℓ_p norm on $\mathbb{R}^n$ defined as

$$\|x^n\|_p = \begin{cases} \left(\sum_{i=1}^n |x_i|^p\right)^{\frac{1}{p}}, & p \in [1,\infty). \\ \max\{|x_1|,\ldots,|x_n|\}, & p = \infty. \end{cases} \tag{2.9}$$

[2]By the equivalence of norms on finite-dimensional space, the the value of the limit on the right-hand side of (2.7) is independent of the underlying norm.

Note that using (2.7), information dimension can be defined for random variables valued on an arbitrary metric space. The equivalence of (2.7) to (2.4) is shown in Appendix A.1.

2.2 Characterizations and properties

The lower and upper information dimension of a random variable might not always be finite, because $H(\langle X \rangle_m)$ can be infinity for all m. However, as pointed out in [24], if the mild condition

$$H(\lfloor X^n \rfloor) < \infty \tag{2.10}$$

is satisfied, we have

$$0 \leq \underline{d}(X^n) \leq \overline{d}(X^n) \leq n. \tag{2.11}$$

The necessity of (2.10) and its equivalence to finite mutual information with additive noise is shown next in Theorem 1. In the scalar case, a sufficient condition for finite information dimension is the existence of the Shannon transform [34, Definition 2.12], i.e.,

$$\mathbb{E}\left[\log(1 + |X|)\right] < \infty. \tag{2.12}$$

See Appendix A.2 for a proof. Consequently, if $\mathbb{E}\left[|X|^\epsilon\right] < \infty$ for some $\epsilon > 0$, then $\overline{d}(X) < \infty$.

Theorem 1. *Let N_{G}^n be standard normal distributed independent of X^n. Define*

$$I(X^n, \mathsf{snr}) \ = I(X^n; \sqrt{\mathsf{snr}} X^n + N_{\mathsf{G}}^n). \tag{2.13}$$

Then the following are equivalent:

1) $0 \leq \underline{d}(X^n) \leq \overline{d}(X^n) \leq n$*;*

2) $H(\lfloor X^n \rfloor) < \infty$*;*

3) $\mathbb{E}\left[\log \frac{1}{P_{X^n}(B_p(X^n, \epsilon))}\right] < \infty$ *for some* $\epsilon > 0$*;*

4) $I(X^n, \mathsf{snr}) < \infty$ *for some* $\mathsf{snr} > 0$*;*

5) $I(X^n, \mathsf{snr}) < \infty$ *for all* $\mathsf{snr} > 0$*.*

If $H(\lfloor X^n \rfloor) = \infty$*, then*

$$d(X^n) = \infty. \tag{2.14}$$

Proof. See Appendix A.2. □

To calculate the information dimension in (2.2) and (2.3), it is sufficient to restrict to the exponential subsequence $m = 2^l$, as a result of the following proposition.

Lemma 1. *Denote*

$$[\cdot]_l \triangleq \langle \cdot \rangle_{2^l}. \tag{2.15}$$

Then

$$\underline{d}(X) = \liminf_{l\to\infty} \frac{H([X]_l)}{l}, \tag{2.16}$$

$$\overline{d}(X) = \limsup_{l\to\infty} \frac{H([X]_l)}{l}. \tag{2.17}$$

Proof. See Appendix A.3. □

By analogous arguments that lead to Theorem 1, we have

Lemma 2. *$\underline{d}(X)$ and $\overline{d}(X)$ are unchanged if rounding or ceiling functions are used in Definition 1.*

Proof. See Appendix A.3. □

The following lemma states additional properties of information dimension, the last three of which are inherited from Shannon entropy. For simplicity we assume that all information dimensions involved exist and are finite.

Lemma 3.

- *Translation-invariance: for all $x^n \in \mathbb{R}^n$,*

$$d(x^n + X^n) = d(X^n). \tag{2.18}$$

- *Scale-invariance: for all $\alpha \neq 0$,*

$$d(\alpha X^n) = d(X^n). \tag{2.19}$$

- *If X^n and Y^n are independent, then*[3]

$$\max\{d(X^n), d(Y^n)\} \leq d(X^n + Y^n) \tag{2.20}$$

$$\leq d(X^n) + d(Y^n). \tag{2.21}$$

- *If $\{X_i\}$ are independent and $d(X_i)$ exists for all i, then*

$$d(X^n) = \sum_{i=1}^{n} d(X_i). \tag{2.22}$$

- *If X^n, Y^n, Z^n are independent, then*

$$d(X^n + Y^n + Z^n) + d(Z^n) \leq d(X^n + Z^n) + d(Y^n + Z^n). \tag{2.23}$$

[3]The upper bound (2.21) can be improved by defining conditional information dimension.

Proof. Appendix A.4. □

The scale-invariance property of information dimension in (2.19) can in fact be generalized to invariance under (the pushforward of) all bi-Lipschitz functions, a property shared by Minkowski and packing dimension [35, Exercise 7.6]. The proof is straightforward in view of the alternative definition of information dimension in (2.7).

Theorem 2. *For any Lipschitz function $f : \mathbb{R} \to \mathbb{R}$, the (lower and upper) information dimension of X do not exceed those of $f(X)$. Consequently, for any bi-Lipschitz function $f : \mathbb{R} \to \mathbb{R}$, the (lower and upper) information dimension of X coincide with those of $f(X)$.*

To conclude this section, we discuss the regularity of information dimension as a functional of probability distributions. Similar to the Shannon entropy, information dimension is also discontinuous with respect to the weak-* topology. This can be verified using Theorem 4 presented in the next section, which shows that the information dimension of discrete and absolutely continuous distributions are zero and one respectively. However, the next result shows that information dimension is Lipschitz continuous with respect to the total variation distance.

Theorem 3. *Let X^n and Y^n be random vectors such that $H(\lfloor X^n \rfloor)$ and $H(\lfloor Y^n \rfloor)$ are both finite. Then*

$$\max\{|\underline{d}(X^n) - \underline{d}(Y^n)|, |\overline{d}(X^n) - \overline{d}(Y^n)|\} \leq n\mathsf{TV}(P_{X^n}, P_{Y^n}). \tag{2.24}$$

Proof. Let $\epsilon = \mathsf{TV}(P_{X^n}, P_{Y^n})$. Since dropping the integer part does not change the information dimension, we shall assume that X^n and Y^n are both valued in the unit hypercube $[0,1]^n$. Let $m \in \mathbb{N}$. Since the total variation distance, as an f-divergence, satisfies the data processing inequality [36], we have $\mathsf{TV}(P_{[X^n]_m}, P_{[Y^n]_m}) \leq \epsilon$. By the variational characterization of the total variation distance [37], we have

$$\mathsf{TV}(P_{[X^n]_m}, P_{[Y^n]_m}) = \min_{P_{[X^n]_m[Y^n]_m}} \mathbb{P}\{[X^n]_m \neq [Y^n]_m\}, \tag{2.25}$$

where the minimum is taken with respect to all joint distributions with the prescribed marginals $P_{[X^n]_m}$ and $P_{[Y^n]_m}$. Now let $[X^n]_m$ and $[Y^n]_m$ be jointly distributed according to the minimizer of (2.25). Then

$$\begin{aligned} H([X^n]_m) &\leq H([X^n]_m|[Y^n]_m) + H([Y^n]_m) && (2.26)\\ &\leq H([X^n]_m|[Y^n]_m, [X^n]_m \neq [Y^n]_m)\epsilon + H([Y^n]_m) && (2.27)\\ &\leq nm\epsilon + H([Y^n]_m) && (2.28) \end{aligned}$$

where (2.27) and (2.28) are due to (2.25) and that $X^n \in [0,1]^n$ respectively. Dividing both sides of (2.28) by m and taking lim inf and lim sup with respect to $m \to \infty$, we obtain (2.24). □

2.3 Evaluation of information dimension

By the Lebesgue Decomposition Theorem [38], a probability distribution can be uniquely represented as the mixture

$$\nu = p\nu_{\mathrm{d}} + q\nu_{\mathrm{c}} + r\nu_{\mathrm{s}}, \tag{2.29}$$

where

- $p + q + r = 1$, $p, q, r \geq 0$.
- ν_{d} is a purely atomic probability measure (discrete part).
- ν_{c} is a probability measure absolutely continuous with respect to Lebesgue measure, i.e., having a probability density function (continuous part[4]).
- ν_{s} is a probability measure singular with respect to Lebesgue measure but with no atoms (singular part).

As shown in [24], the information dimension for the mixture of discrete and absolutely continuous distribution can be determined as follows (see [24, Theorem 1 and 3] or [17, Theorem 1, pp. 588 – 592]):

Theorem 4. *Let X be a random variable such that $H(\lfloor X \rfloor) < \infty$. Assume the distribution of X can be represented as*

$$\nu = (1-\rho)\nu_{\mathrm{d}} + \rho\nu_{\mathrm{c}}, \tag{2.30}$$

where ν_{d} is a discrete measure, ν_{c} is an absolutely continuous measure and $0 \leq \rho \leq 1$. Then

$$d(X) = \rho. \tag{2.31}$$

Furthermore, given the finiteness of $H(\nu_{\mathrm{d}})$ and $h(\nu_{\mathrm{c}})$, $\hat{H}(X)$ admits a simple formula

$$\hat{H}(X) = (1-\rho)H(\nu_{\mathrm{d}}) + \rho h(\nu_{\mathrm{c}}) + h(\rho), \tag{2.32}$$

where $H(\nu_{\mathrm{d}})$ is the Shannon entropy of ν_{d}, $h(\nu_{\mathrm{c}})$ is the differential entropy of ν_{c}, and $h(\rho) \triangleq \rho \log \frac{1}{\rho} + (1-\rho) \log \frac{1}{1-\rho}$ is the binary entropy function.

Some consequences of Theorem 4 are as follows: as long as $H(\lfloor X \rfloor) < \infty$,

1. X is discrete: $d(X) = 0$, and $\hat{H}(X)$ coincides with the Shannon entropy of X.
2. X is continuous: $d(X) = 1$, and $\hat{H}(X)$ is equal to the differential entropy of X.
3. X is discrete-continuous-mixed: $d(X) = \rho$, and $\hat{H}(X)$ is the weighted sum of the entropy of discrete and continuous parts plus a term of $h_2(\rho)$.

For mixtures of countably many distributions, we have the following theorem.

[4] In measure theory, sometimes a measure is called continuous if it does not have any atoms, and a singular measure is called singularly continuous. Here we say a measure continuous if and only if it is absolutely continuous.

Theorem 5. *Let Y be a discrete random variable with $H(Y) < \infty$. If $d(P_{X|Y=i})$ exists for all i, then $d(X)$ exists and is given by $d(X) = \sum_{i=1}^{\infty} P_Y(i) d(P_{X|Y=i})$. More generally,*

$$\overline{d}(X) \leq \sum_{i=1}^{\infty} P_Y(i) \overline{d}(P_{X|Y=i}), \tag{2.33}$$

$$\underline{d}(X) \geq \sum_{i=1}^{\infty} P_Y(i) \underline{d}(P_{X|Y=i}). \tag{2.34}$$

Proof. For any m, the conditional distribution of $\langle X \rangle_m$ given $Y = i$ is the same as $\left\langle X^{(i)} \right\rangle_m$. Then

$$H(\langle X \rangle_m | Y) \leq H(\langle X \rangle_m) \leq H(\langle X \rangle_m | Y) + H(Y), \tag{2.35}$$

where

$$H\left(\langle X \rangle_m | Y\right) = \sum_{i=1}^{\infty} P_Y(i) H\left(\langle X \rangle_m | Y = i\right). \tag{2.36}$$

Since $H(Y) < \infty$, dividing both sides of (2.35) by $\log m$ and sending $m \to \infty$ yields (2.33) and (2.34). □

To summarize, when X has a discrete-continuous mixed distribution, the information dimension of X is given by the weight of the continuous part. When the distribution of X has a singular component, its information dimension does not admit a simple formula in general. For instance, it is possible that $\underline{d}(X) < \overline{d}(X)$ (see [24, p. 195] or the example at the end of Section 2.5). However, for the important class of *self-similar* singular distributions, the information dimension can be explicitly determined. See Section 2.8.

2.4 Information dimension of order α

With Shannon entropy replaced by Rényi entropy in (2.2) – (2.3), the generalized notion of information dimension of order α can be defined similarly as follows [39, 17]:[5]

Definition 2 (Information dimension of order α)**.** Let $\alpha \in [0, \infty]$. The lower and upper dimensions of X of order α are defined as

$$\underline{d}_\alpha(X) = \liminf_{m \to \infty} \frac{H_\alpha\left(\langle X \rangle_m\right)}{\log m} \tag{2.37}$$

and

$$\overline{d}_\alpha(X) = \limsup_{m \to \infty} \frac{H_\alpha\left(\langle X \rangle_m\right)}{\log m} \tag{2.38}$$

[5] d_α is also known as the L^α dimension or spectrum [23, 33]. In particular, d_2 is called the correlation dimension.

respectively, where $H_\alpha(Y)$ denotes the Rényi entropy of order α of a discrete random variable Y with probability mass function $\{p_y : y \in \mathcal{Y}\}$, defined as

$$H_\alpha(Y) = \begin{cases} \sum_{y\in\mathcal{X}} p_y \log \frac{1}{p_y}, & \alpha = 1, \\ \log \frac{1}{\max\limits_{y\in\mathcal{Y}} p_y}, & \alpha = \infty, \\ \frac{1}{1-\alpha} \log \left(\sum_{y\in\mathcal{Y}} p_y^\alpha \right), & \alpha \neq 1, \infty. \end{cases} \tag{2.39}$$

If $\underline{d}_\alpha(X) = \overline{d}_\alpha(X)$, the common value is called the information dimension of X of order α, denoted by $d_\alpha(X)$. The "the entropy of X of order α and dimension $d_\alpha(X)$" as

$$\hat{H}_\alpha(X) = \lim_{m\to\infty} \left(H_\alpha(\langle X \rangle_m) - d_\alpha(X) \log m \right), \tag{2.40}$$

provided the limit exists.

Analogous to (2.7), d_α has the following "continuous" version of definition that is equivalent to Definition 2 (see, for example, [21, Section 2.1] or [40, Definition 2.6]):

$$d_\alpha(X) = \lim_{\epsilon\downarrow 0} \frac{\log \mathbb{E}\left[\mu(B_p(X,\epsilon))^{\alpha-1}\right]}{(\alpha-1)\log\epsilon}. \tag{2.41}$$

As a consequence of the monotonicity of Rényi entropy, information dimensions of different orders satisfy the following result.

Lemma 4. *For $\alpha \in [0,\infty]$, $\underline{d}_\alpha$ and $\overline{d}_\alpha$ both decrease with α. Define*

$$\hat{d}(X) = \lim_{\alpha\uparrow 1} \overline{d}_\alpha(X). \tag{2.42}$$

Then

$$0 \le \underline{d}_\infty \le \lim_{\alpha\downarrow 1} \underline{d}_\alpha \le \underline{d} \le \overline{d} \le \hat{d} \le \overline{d}_0. \tag{2.43}$$

Proof. All inequalities follow from the fact that for a fixed random variable Y, $H_\alpha(Y)$ decreases with α in $[0,\infty]$. □

For dimension of order α we highlight the following result from [39].

Theorem 6 ([39, Theorem 3])**.** *Let X be a random variable whose distribution has Lebesgue decomposition as in (2.29). Assume that $H_\alpha(\lfloor X \rfloor) < \infty$. Then*

1. *For $\alpha > 1$: if $p > 0$, that is, X has a discrete component, then $d_\alpha(X) = 0$ and $\hat{H}_\alpha(X) = H_\alpha(\nu_\mathrm{d}) + \frac{\alpha}{1-\alpha}\log p$.*
2. *For $\alpha < 1$: if $q > 0$, that is, X has an absolutely continuous component, then $d_\alpha(X) = 1$ and $\hat{H}_\alpha(X) = h_\alpha(\nu_\mathrm{c}) + \frac{\alpha}{1-\alpha}\log q$, where $h_\alpha(\nu_\mathrm{c})$ is the differential Rényi entropy of ν_c, defined in terms of its density f_c as*

$$h_\alpha(\nu_\mathrm{c}) = \frac{1}{1-\alpha} \log \int f_\mathrm{c}^\alpha(x)\mathrm{d}x. \tag{2.44}$$

As a consequence of the monotonicity, $\alpha \mapsto d_\alpha(X)$ can have at most countably infinite discontinuities. Remarkably, the only possible discontinuity is at $\alpha = 1$, as a result of the following inequality due to Beck [41]:

$$\frac{\alpha' - 1}{\alpha'} d_{\alpha'}(X) \geq \frac{\alpha - 1}{\alpha} d_\alpha(X), \quad \alpha < \alpha', \alpha, \alpha' \neq 0, 1. \tag{2.45}$$

Recall that $d_1(X) = d(X)$. According to Theorem 6. for discrete-continuous mixed distributions, $d_\alpha(X) = \hat{d}(X) = 1$ for all $\alpha < 1$, while $d(X)$ equals to the weight of the continuous part. However, for Cantor distribution (see Section 2.8), $d_\alpha(X) = \hat{d}(X) = d(X) = \log_3 2$ for all α.

2.5 Interpretation of information dimension as entropy rate

Let $(x)_i$ denote the ith bit in the binary expansion of $0 \leq x < 1$, that is,

$$(x)_i = 2^i([x]_i - [x]_{i-1}) = \lfloor 2^i x \rfloor - 2\lfloor 2^{i-1} x \rfloor \in \mathbb{Z}_2. \tag{2.46}$$

Let $X \in [0, 1]$ a.s. Observe that $H(\langle X \rangle_m) \leq \log m$, since the range of $\langle X \rangle_m$ contains at most m values. Then $0 \leq \underline{d}(X) \leq \overline{d}(X) \leq 1$. The dyadic expansion of X can be written as

$$X = \sum_{j=1}^{\infty} (X)_j 2^{-j}, \tag{2.47}$$

where each $(X)_j$ is a binary random variable. Therefore there is a one-to-one correspondence between X and the binary random process $\{(X)_j : j \in \mathbb{N}\}$. Note that the partial sum in (2.47) is

$$[X]_i = \sum_{j=1}^{i} (X)_j 2^{-j}, \tag{2.48}$$

and $[X]_i$ and $((X)_1, \ldots, (X)_i)$ are in one-to-one correspondence, therefore

$$H([X]_i) = H\left((X)_1, \ldots, (X)_i\right). \tag{2.49}$$

By Lemma 1, we have

$$\underline{d}(X) = \liminf_{i \to \infty} \frac{H\left((X)_1, \ldots, (X)_i\right)}{i}, \tag{2.50}$$

$$\overline{d}(X) = \limsup_{i \to \infty} \frac{H\left((X)_1, \ldots, (X)_i\right)}{i}. \tag{2.51}$$

Thus the information dimension of X is precisely the *entropy rate* of its dyadic expansion, or the entropy rate of any M-ary expansion of X, divided by $\log M$.

This interpretation of information dimension enables us to gain more intuition about the result in Theorem 4. When X has a discrete distribution, its dyadic expansion has zero entropy rate. When X is uniform on $[0,1]$, its dyadic expansion is i.i.d. equiprobable, and therefore it has unit entropy rate in bits. If X is continuous, but nonuniform, its dyadic expansion still has unit entropy rate. Moreover, from (2.50), (2.51) and Theorem 4 we have

$$\lim_{i\to\infty} D\left((X)_1,\ldots,(X)_i \,||\, \text{equiprobable}\right) = -h(X), \tag{2.52}$$

where $D(\cdot\,||\,\cdot)$ denotes the relative entropy and the differential entropy is $h(X) \le 0$ since $X \in [0,1]$ a.s. The information dimension of a discrete-continuous mixture is also easily understood from this point of view, because the entropy rate of a mixed process is the weighted sum of entropy rates of each component. Moreover, random variables whose lower and upper information dimensions differ can be easily constructed from processes with different lower and upper entropy rates.

2.6 Connections with rate-distortion theory

The asymptotic behavior of the rate-distortion function, in particular, the asymptotic tightness of the Shannon lower bound in the high-rate regime, has been addressed in [42, 43] for continuous sources. In [25] Kawabata and Dembo generalized it to real-valued sources that do not necessarily possess a density, and showed that the information dimension plays a central role.

The asymptotic tightness of the Shannon lower bound in the high-rate regime is shown by the following result.

Theorem 7 ([43])**.** *Let X be a random variable on the normed space $(\mathbb{R}^k, \|\cdot\|)$ with a density such that $h(X) > -\infty$. Let the distortion function be $\rho(x,y) = \|x-y\|^r$ with $r > 0$, and the single-letter rate-distortion function is given by*

$$R_X(D) = \inf_{\mathbb{E}\|Y-X\|^r \le D} I(X;Y). \tag{2.53}$$

Suppose that there exists an $\alpha > 0$ such that $\mathbb{E}\,\|X\|^\alpha < \infty$. Then $R_X(D) \ge R_{\mathrm{L}}(D)$ and

$$\lim_{D\to 0}(R_X(D) - R_{\mathrm{L}}(D)) = 0, \tag{2.54}$$

where the Shannon lower bound $R_{\mathrm{L}}(D)$ takes the following form

$$R_{\mathrm{L}}(D) = \frac{k}{r}\log\frac{1}{D} + h(X) + \frac{k}{r}\log\frac{k}{r\mathrm{e}} - \frac{k}{r}\Gamma\left(\frac{k}{r}\right)\log V_k, \tag{2.55}$$

where V_k denotes the volume of the unit ball $B_k = \{x \in \mathbb{R}^k : \|x\| \le 1\}$.

Special cases of Theorem 7 include:

- MSE distortion and scalar source: $\rho(x,y) = (x-y)^2$, $\|x\| = |x|$, $r = 2$, $k = 1$ and $V_k = 2$. Then the Shannon lower bound takes the familiar form

$$R_{\mathrm{L}}(D) = h(X) - \frac{1}{2}\log(2\pi \mathrm{e} D). \tag{2.56}$$

- Absolute distortion and scalar source: $\rho(x,y) = |x-y|$, $\|x\| = |x|$, $r = 1$, $k = 1$ and $V_k = 2$. Then

$$R_{\mathrm{L}}(D) = h(X) - \log(2\mathrm{e} D). \tag{2.57}$$

- ℓ_∞ norm and vector source: $\rho(x,y) = \|x-y\|_\infty$, $\|x\| = \|x\|_\infty$, $r = 1$ and $V_k = 2^k$. Then

$$R_{\mathrm{L}}(D) = h(X) + \log\left(\frac{1}{k!}\left(\frac{k}{2\mathrm{e}D}\right)^k\right). \tag{2.58}$$

For general sources, Kawabata and Dembo introduced the concept of *rate-distortion dimension* in [25]. The rate-distortion dimension of a measure μ (or a random variable X with the distribution μ) on the metric space $(\mathbb{R}^k, d)$ is defined as follows:

$$\overline{\dim}_{\mathrm{R}}(X) = \limsup_{D\downarrow 0} \frac{R_X(D)}{\frac{1}{r}\log\frac{1}{D}}, \tag{2.59}$$

$$\underline{\dim}_{\mathrm{R}}(X) = \liminf_{D\downarrow 0} \frac{R_X(D)}{\frac{1}{r}\log\frac{1}{D}}, \tag{2.60}$$

where $R_X(D)$ is the rate-distortion function of X with distortion function $\rho(x,y) = d(x,y)^r$ and $r > 0$. Then under the equivalence of the metric and the ℓ_∞ norm as in (2.61), the rate-distortion dimension coincides with the information dimension of X:

Theorem 8 ([25, Proposition 3.3])**.** *Consider the metric space $(\mathbb{R}^k, d)$. If there exists $a_1, a_2 > 0$, such that for all $x, y \in \mathbb{R}^k$,*

$$a_2 \|x-y\|_\infty \le d(x,y) \le a_1 \|x-y\|_\infty, \tag{2.61}$$

then

$$\overline{\dim}_{\mathrm{R}}(X) = \overline{d}(X), \tag{2.62}$$

$$\underline{\dim}_{\mathrm{R}}(X) = \underline{d}(X), \tag{2.63}$$

Moreover, (2.62) and (2.63) hold even if the ϵ-entropy $H_\epsilon(X)$ instead of the rate-distortion function $R_X(\epsilon)$ is used in the definition of $\overline{\dim}_{\mathrm{R}}(X)$ and $\underline{\dim}_{\mathrm{R}}(X)$.

In particular, consider the special case of scalar source and MSE distortion $d(x,y) = |x-y|^2$. Then whenever $d(X)$ exists and is finite,[6]

$$R_X(D) = \frac{d(X)}{2} \log \frac{1}{D} + o(\log D), \quad D \to 0. \tag{2.64}$$

Therefore $\frac{d(X)}{2}$ is the scaling constant of $R_X(D)$ with respect to $\log \frac{1}{D}$ in the high-rate regime. This fact gives an operational characterization of information dimension in Shannon theory. Note that in the most familiar cases we can sharpen (2.64) to show the following: as $D \to 0$,

- X is discrete and $H(X) < \infty$: $R_X(D) = H(X) + o(1)$.
- X is continuous and $h(X) > -\infty$: $R_X(D) = \frac{1}{2} \log \frac{1}{2\pi e D} + h(X) + o(1)$.
- X is discrete-continuous mixed: if P_X is given by the mixture (2.30) with finite variance, then [42, Theorem 1]

$$R_X(D) = \frac{\rho}{2} \log \frac{1}{D} + \hat{H}(X) + o(1). \tag{2.65}$$

  where $\hat{H}(X)$ is the entropy of X of dimension ρ given in (2.32). In view of the definition of $\hat{H}$ in (2.5), we see that for mixtures the rate-distortion function coincides with the entropy of the quantized version up to $o(1)$ term, i.e., if $m = \frac{1}{\sqrt{D}}(1 + o(1))$, then $R_X(D) - H(\langle X \rangle_m) = o(1)$.

Remark 1. Information dimension is also related to the fine quantization of random variables. See the notion of *quantization dimension* introduced by Zador [44] and discussed in [45, Chapter 11].

2.7 High-SNR asymptotics of mutual information

The high-SNR asymptotics of mutual information with additive noise is governed by the *input information dimension*. Recall the notation $I(X, \mathsf{snr})$ defined in (2.13) which denotes the mutual information between X and $\sqrt{\mathsf{snr}}X + N_{\mathsf{G}}$, where N_{G} is standard additive Gaussian noise. Using [26, Theorems 2.7 and 3.1], Theorem 1 and the equivalent definition of information dimension in (2.7), we can relate the scaling law of mutual information under weak noise to the information dimension:

Theorem 9. *For any random variable X, Then*

$$\limsup_{\mathsf{snr}\to\infty} \frac{I(X, \mathsf{snr})}{\frac{1}{2}\log \mathsf{snr}} = \overline{d}(X), \tag{2.66}$$

$$\liminf_{\mathsf{snr}\to\infty} \frac{I(X, \mathsf{snr})}{\frac{1}{2}\log \mathsf{snr}} = \underline{d}(X). \tag{2.67}$$

[6] We use the following asymptotic notations: $f(x) = O(g(x))$ if $\limsup \frac{|f(x)|}{|g(x)|} < \infty$, $f(x) = \Omega(g(x))$ if $g(x) = O(f(x))$, $f(x) = \Theta(g(x))$ if $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$, $f(x) = o(g(x))$ if $\lim \frac{|f(x)|}{|g(x)|} = 0$, $f(x) = \omega(g(x))$ if $g(x) = o(f(x))$.

The original results in [26] are proved under the assumption of (2.12), which can in fact be dropped. Also, the authors of [26] did not point out that the "classical analogue" of the free entropy dimension they defined in fact coincides with Rényi's information dimension. For completeness, we give a proof in Appendix A.5, which includes a non-asymptotic version of Theorem 9.

Theorem 9 also holds for non-Gaussian noise N with finite non-Gaussianness:

$$D(N) \triangleq D(N||\Phi_N) < \infty, \tag{2.68}$$

where Φ_N is Gaussian distributed with the same mean and variance as N. This is because

$$I(X; \sqrt{\mathsf{snr}}X + \Phi_N) \leq I(X; \sqrt{\mathsf{snr}}X + N) \leq I(X; \sqrt{\mathsf{snr}}X + \Phi_N) + D(N). \tag{2.69}$$

Moreover, Theorem 9 can be readily extended to random vectors.

Remark 2. Consider the MSE distortion criterion. As a consequence of Theorems 9 and 8, when $d(X)$ exists, we have

$$I(X, \mathsf{snr}) = \frac{d(X)}{2} \log \mathsf{snr} + o(\log \mathsf{snr}), \tag{2.70}$$

$$R_X(D) = \frac{d(X)}{2} \log \frac{1}{D} + o\left(\log \frac{1}{D}\right), \tag{2.71}$$

when $\mathsf{snr} \to \infty$ and $D \to 0$ respectively. Therefore in the minimization problem (2.53), choosing the forward channel $P_{Y|X}$ according to $Y = X + \sqrt{D}N_{\mathsf{G}}$ is first-order optimal. However, it is unclear whether such a choice is optimal in the following stronger sense:

$$\lim_{D \downarrow 0} R_X(D) - I(X; X + \sqrt{D}N_{\mathsf{G}}) = 0. \tag{2.72}$$

2.8 Self-similar distributions

2.8.1 Definitions

An *iterated function system* (IFS) is a family of contractions $\{F_1, \ldots, F_m\}$ on $\mathbb{R}^n$, where $2 \leq m < \infty$, and $F_j : \mathbb{R}^n \to \mathbb{R}^n$ satisfies $\|F_j(x) - F_j(y)\|_2 \leq r_j \|x - y\|_2$ for all $x, y \in \mathbb{R}^n$ with $0 < r_j < 1$. By [46, Theorem 2.6], given an IFS there is a unique nonempty compact set E, called the *invariant set* of the IFS, such that $E = \bigcup_{j=1}^m F_j(E)$. The corresponding invariant set is called *self-similar*, if the IFS consists of *similarity transformations*, that is,

$$F_j(x) = r_j \boldsymbol{O}_j x + w_j \tag{2.73}$$

with $\boldsymbol{O}_j$ an orthogonal matrix, $r_j \in (0,1)$ and $w_j \in \mathbb{R}^n$, in which case

$$\frac{\|F_j(x) - F_j(y)\|_2}{\|x-y\|_2} = r_j \tag{2.74}$$

is called the *similarity ratio* of F_j. If r_j is constant for all j, we say that E is *homogeneous* [47] (or *equicontractive* [48]). Self-similar sets are usually fractal [18]. For example, consider the IFS on $\mathbb{R}$ with

$$F_1(x) = \frac{x}{3}, \quad F_2(x) = \frac{x+2}{3}. \tag{2.75}$$

The resulting invariant set is the middle-third Cantor set.

Now we define measures supported on a self-similar set E associated with the IFS $\{F_1, \ldots, F_m\}$.

Definition 3. Let $P = \{p_1, \ldots, p_m\}$ be a probability vector and $\{F_1, \ldots, F_m\}$ an IFS. A Borel probability measure μ is called *self-similar* if

$$\mu(A) = \sum_{j=1}^{m} p_j \mu(F_j^{-1}(A)), \tag{2.76}$$

for all Borel subsets A.

Classical results from fractal geometry ([49], [46, Theorem 2.8]) assert the existence and uniqueness of μ, which is supported on the self-similar invariance set E that satisfies $\mu(F_j(E)) = p_j$ for each j. The usual Cantor distribution [38] can be defined through the IFS in (2.75) and $P = \{1/2, 1/2\}$.

Remark 3. The fixed-point equation in (2.76) can be *equivalently* understood as follows. Consider a random transformation on $\mathbb{R}$ defined by mapping x to $F_j(x)$ with probability p_j. Then μ is the (unique) fixed point of this random transformation, i.e., the input distribution μ induces the same output distribution. For one-dimensional self-similar measure with IFS (2.73), the corresponding random transformation is given by

$$Y = RX + W \tag{2.77}$$

where X is independent of (R, W) and R is correlated with W according to $\mathbb{P}\{R = r_j, W = w_j\} = p_j, j = 1, \ldots, m$. It is easy to check that the distribution of

$$X = \sum_{k=1}^{\infty} W_k \prod_{i=1}^{k-1} R_i \tag{2.78}$$

is the fixed-point of (2.77), where $\{(R_k, W_k)\}$ in (2.77) is an i.i.d. sequence. In the homogeneous case, (2.77) becomes

$$Y = rX + W \tag{2.79}$$

and (2.78) simplifies to a randomly modulated geometric series[7]

$$X = \sum_{k \geq 1} W_k r^{k-1}, \tag{2.80}$$

where $\{W_k\}$ are i.i.d.

From (2.80) we see that distributions defined by M-ary expansion with independent digits can be equivalently defined via IFS, in which case $r = \frac{1}{M}$ and W_k is $\{0, \ldots, M-1\}$-valued. For example, the ternary expansion of the Cantor distribution consists of independent digits which take value 0 or 2 equiprobably.

2.8.2 Information dimension

It is shown in [23] that[8] the information dimension (or dimension of order α) of self-similar measures always exists.

Theorem 10. *Let X be distributed according to μ in (2.76). Then the sequence $\{H([X]_m) - 10\}$ is superadditive. Consequently, $d(X)$ exists and satisfies*

$$d(X) = \sup_{m \geq 1} \frac{H([X]_m)}{m}. \tag{2.81}$$

Proof. See [23, Theorem 1.1, Lemma 2.2 and Equation (3.10)]. □

In spite of the general existence result in [23], there is no known formula for the information dimension unless the IFS satisfies certain *separation* conditions, for example [18, p. 129]:

- *strong separation condition*: $\{F_j(E) : j = 1, \ldots, m\}$ are disjoint.
- *open set condition*: there exists a nonempty bounded open set $U \subset \mathbb{R}^n$, such that $\bigcup_j F_j(U) \subset U$ and $F_i(U) \cap F_j(U) = \emptyset$ for $i \neq j$.

Note that the open set condition is weaker than the strong separation condition. The following lemma (proved in Appendix A.6) gives a sufficient condition for a self-similar homogeneous measure (or random variables of the form (2.80)) to satisfy the open set condition.

Lemma 5. *Let the IFS $\{F_1, \ldots, F_m\}$ consist of similarity transformations $F_j(x) = rx + w_j$, where $0 < r < 1$ and $\mathcal{W} \triangleq \{w_1, \ldots, w_m\} \subset \mathbb{R}$. Then the open set condition is satisfied if*

$$r \leq \frac{\mathsf{m}(\mathcal{W})}{\mathsf{m}(\mathcal{W}) + \mathsf{M}(\mathcal{W})}, \tag{2.82}$$

where the minimum distance *and* spread *of $\mathcal{W}$ are defined by*

$$\mathsf{m}(\mathcal{W}) \triangleq \min_{w_i \neq w_j \in \mathcal{W}} |w_i - w_j|. \tag{2.83}$$

[7] When W_k is binary, (2.80) is known as the *Bernoulli convolution*, an important object in fractal geometry and ergodic theory [50].

[8] The results in [23] are more general, encompassing nonlinear IFS as well.

and

$$\mathsf{M}(\mathcal{W}) \triangleq \max_{w_i, w_j \in \mathcal{W}} |w_i - w_j| \tag{2.84}$$

respectively.

The next result gives a formula for the information dimension of a self-similar measure μ with IFS satisfying the open set condition:

Theorem 11 ([51], [18, Chapter 17].)**.** *Let μ be the self-similar measure (2.76) generated by the IFS (2.73) which satisfies the open set condition. Then*

$$d(\mu) = \frac{H(P)}{\sum_{j=1}^{m} p_j \log \frac{1}{r_j}} \tag{2.85}$$

and

$$d_\alpha(\mu) = \frac{\beta(\alpha)}{\alpha - 1}, \quad \alpha \neq 1, \tag{2.86}$$

where $\beta(\alpha)$ is the unique solution to the following equation of β:

$$\sum_{j=1}^{m} p_j^\alpha r_j^\beta = 1 \tag{2.87}$$

In particular, for self-similar homogeneous μ with $r_j = r$ for all j,

$$d_\alpha(\mu) = \frac{H_\alpha(P)}{\log \frac{1}{r}}, \quad \alpha \geq 0. \tag{2.88}$$

Theorem 11 can be generalized to

- non-linear IFS, with similarity ratio r_j in (2.85) replaced by the Lipschitz constant of F_j [52];
- random variables of the form (2.78) where the sequence $\{(W_k)\}$ is stationary and ergodic, with $H(P)$ in (2.85) replaced by the entropy rate [25, Theorem 4.1].
- IFS with overlaps [53].

The open set condition implies that $\sum_{j=1}^{m} \mathrm{Leb}(F_j(U)) = \mathrm{Leb}(U) \sum_{j=1}^{m} r_j^n \leq \mathrm{Leb}(U)$. Since $0 < \mathrm{Leb}(U) < \infty$, it follows that

$$\sum_{j=1}^{m} r_j^n \leq 1. \tag{2.89}$$

In view of (2.85), we have

$$d(X) = \frac{n\, H(P)}{H(P) + D(P \,||\, R)}, \tag{2.90}$$

where $R = (r_1^n, \ldots, r_m^n)$ is a *sub-probability* measure. Since $D(P \,||\, R) \geq 0$, we have $0 \leq d(X) \leq n$, which agrees with (2.11). Note that as long as the open set condition is satisfied, the information dimension does *not* depend on the translates $\{w_1, \ldots, w_m\}$.

To conclude this subsection, we note that in the homogeneous case, it can be shown that (2.88) always holds with inequality without imposing any separation condition.

Lemma 6. *For self-similar homogeneous measure μ with similarity ratio r,*

$$d_\alpha(\mu) \leq \frac{H_\alpha(P)}{\log \frac{1}{r}}, \quad \alpha \geq 0. \tag{2.91}$$

Proof. See Appendix A.7. □

2.9 Information dimension under projection

Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m \leq n$. Then for any X^n,

$$d(\mathbf{A}X^n) \leq \min\{d(X^n), \mathsf{rank}(\mathbf{A})\}. \tag{2.92}$$

Understanding how the dimension of a measure behaves under projections is a basic problem in fractal geometry. It is well-known that *almost every* projection preserves the dimension, be it Hausdorff dimension (Marstrand's projection theorem [35, Chapter 9]) or information dimension [33, Theorems 1.1 and 4.1]. However, computing the dimension for individual projections is in general difficult.

The preservation of information dimension of order $1 < \alpha \leq 2$ under typical projections is shown in [33, Theorem 1], which is relevant to our investigation of the almost sure behavior of degrees of freedom. The proof of the following result is potential-theoretic.

Lemma 7. *Let $\alpha \in (1, 2]$ and $m \leq n$. Then for almost every $\mathbf{A} \in \mathbb{R}^{m \times n}$,*

$$d_\alpha(\mathbf{A}X^n) = \min\left\{d_\alpha(X^n), m\right\}, \tag{2.93}$$

It is easy to see that (2.93) fails for information dimension ($\alpha = 1$) [33, p. 1041]: let $P_{X^n} = (1-\rho)\delta_0 + \rho Q$, where Q is an absolutely continuous distribution on $\mathbb{R}^n$ and $\rho < \frac{m}{n}$. Then $d(X^n) = \rho n$ and $\min\left\{d_\alpha(X^n), m\right\} = m$. Note that $\mathbf{A}X^n$ also has a mixed distribution with an atom at zero of mass at least $1 - \rho$. Then $d(\mathbf{A}X^n) = \rho m < m$. Examples of nonpreservation of d_α for $0 \leq \alpha < 1$ or $\alpha > 2$ are given in [33, Section 5].

A problem closely related to the degrees of freedom of interference channel is to determine the *dimension difference* of a two-dimensional product measure under two projections [54]. To ease the exposition, we assume that all information dimension in

the sequel exist. Let p, q, p', q' be non-zero real numbers. Then

$$\begin{aligned} & d(pX+qY) - d(p'X+q'Y) \\ \leq & \frac{1}{2}(d(pX+qY) + d(pX+qY) - d(X) - d(Y)) \qquad (2.94) \\ \leq & \frac{1}{2}. \qquad (2.95) \end{aligned}$$

where (2.94) and (2.95) follow from (2.20) and (2.21) respectively. Therefore the dimension of a two-dimensional product measure under two projections can differ by at most *one half*. However, the next result shows that if the coefficients are rational, then the dimension difference is *strictly* less than one half. The proof is based on new entropy inequalities for sums of dilated independent random variables derived in Appendix A.8.

Theorem 12. *Let p, p', q, q' be non-zero integers. Then*

$$d(pX+qY) - d(p'X+q'Y) \leq \frac{1}{2} - \epsilon(p'q, pq'), \qquad (2.96)$$

where

$$\epsilon(a,b) \triangleq \frac{1}{36(\lfloor \log|a| \rfloor + \lfloor \log|b| \rfloor) + 74}. \qquad (2.97)$$

Proof. Appendix A.8. □

If $p' = q' = 1$, for the special case of $(p, q) = (1, 2)$ and $(1, -1)$, the right-hand side of (2.96) can be improved to $\frac{3}{7}$ and $\frac{2}{5}$ respectively. See Remark 32 at the end of Appendix A.8.

For irrational coefficients, (2.95) is tight in the following sense:

Theorem 13. *Let $\frac{p'q}{pq'}$ be irrational. Then*

$$\sup_{X \perp\!\!\!\perp Y} d(pX+qY) - d(p'X+q'Y) = \frac{1}{2}. \qquad (2.98)$$

Proof. Appendix A.9. □

Chapter 3

Minkowski dimension

In this chapter we introduce the notion of Minkowski dimension, a central concept in fractal geometry which gauges the degree of fractality of sets and measures. The relevance of Minkowski dimension to compressed sensing is mainly due to the linear embeddability of sets of low Minkowski dimension into low-dimensional Euclidean spaces. Section 3.1 gives the definition of Minkowski dimension and overviews its properties. Section 3.2 contains definitions and coding theorems of lossless Minkowski dimension compression, which are important intermediate results for Chapter 5. New type of concentration-of-measure type of results are proved for memoryless sources, which shows that overwhelmingly large probability is concentrated on subsets of low (Minkowski) dimension. The material in this chapter has been presented in part in [5].

3.1 Minkowski dimension of sets and measures

Definition 4 (Covering number)**.** Let A be a nonempty bounded subset of the metric space (X, d). For $\epsilon > 0$, define $N_A(\epsilon)$, the ϵ-covering number of A, to be the smallest number of ϵ-balls needed to cover A, that is,

$$N_A(\epsilon) = \min \left\{ k : A \subset \bigcup_{i=1}^{k} B(x_i, \epsilon), x_i \in X \right\}. \tag{3.1}$$

Definition 5 (Minkowski dimensions)**.** Let A be a nonempty bounded subset of metric space (X, d). Define the lower and upper Minkowski dimensions of A as

$$\underline{\dim}_{\mathrm{B}} A = \liminf_{\epsilon \to 0} \frac{\log N_A(\epsilon)}{\log \frac{1}{\epsilon}} \tag{3.2}$$

and

$$\overline{\dim}_{\mathrm{B}} A = \limsup_{\epsilon \to 0} \frac{\log N_A(\epsilon)}{\log \frac{1}{\epsilon}} \tag{3.3}$$

respectively. If $\underline{\dim}_{\mathrm{B}} A = \overline{\dim}_{\mathrm{B}} A$, the common value is called the Minkowski dimension of A, denoted by $\dim_{\mathrm{B}} A$.

The (ϵ-)Minkowski dimension of a probability measure is defined as the lowest Minkowski dimension among all sets with measure at least $1-\epsilon$ [30].

Definition 6. Let μ be a probability measure on $(\mathbb{R}^n, \mathcal{B}_{\mathbb{R}^n})$. Define the ($\epsilon$-)Minkowski dimension of μ as

$$\overline{\dim}_{\mathrm{B}}^{\epsilon}(\mu) = \inf\{\overline{\dim}_{\mathrm{B}}(A) : \mu(A) \geq 1-\epsilon\}. \tag{3.4}$$

It should be pointed out that the Minkowski dimension depends on the underlying metric. Nevertheless, equivalent metrics result in the same dimension. A few examples are:

- $\dim_{\mathrm{B}} A = 0$ for any finite set A.
- $\dim_{\mathrm{B}} A = n$ for any bounded set A of *nonempty interior* in Euclidean space $\mathbb{R}^n$.
- Let C be the middle-third Cantor set in the unit interval. Then $\dim_{\mathrm{B}} C = \log_3 2$ [18, Example 3.3].
- $\dim_{\mathrm{B}} \left\{\frac{1}{i} : i \in \mathbb{N}\right\} = \frac{1}{2} > 0$ [18, Example 3.5]. From this example wee see that Minkowski dimension lacks certain stability properties one would expect of a dimension, since it is often desirable that adding a countable set would have no effect on the dimension. This property fails for Minkowski dimension. On the contrary, we observe that Rényi information dimension exhibits stability with respect to adding a discrete component as long as the entropy is finite. However, mixing any distribution with a discrete measure with unbounded support and infinite entropy will necessarily result in infinite information dimension.

The upper Minkowski dimension satisfies the following properties (see [18, p. 48 (iii) and p. 102 (7.9)]), which will be used in the proof of Theorem 14.

Lemma 8. *For bounded sets $A_1, \ldots, A_k$,*

$$\overline{\dim}_{\mathrm{B}} \bigcup_{i=1}^{k} A_i = \max_i \overline{\dim}_{\mathrm{B}} A_i, \tag{3.5}$$

$$\overline{\dim}_{\mathrm{B}} A_1 \times \cdots \times A_k \leq \sum_i \overline{\dim}_{\mathrm{B}} A_i. \tag{3.6}$$

The following lemma shows that in Euclidean spaces, without loss of generality we can restrict attention to covering A with *mesh cubes*: For $z^n \in \mathbb{Z}^n$,

$$C_m(z^n) \triangleq \left\{x^n \in \mathbb{R}^n : [x^n]_m = \frac{z^n}{2^m}\right\} \tag{3.7}$$

$$= \prod_{i=1}^{n} \left[\frac{z_i}{2^m}, \frac{z_i+1}{2^m}\right). \tag{3.8}$$

Since all the mesh cubes partition the whole space, to calculate the lower or upper Minkowski dimension of a set, it suffices to count the number of mesh cubes it intersects, hence justifying the name of box-counting dimension. Denote by $\tilde{N}_A(2^{-m})$ the smallest number of mesh cubes of size 2^{-m} that covers A, that is,

$$\tilde{N}_A(2^{-m}) = \min\left\{k : A \subset \bigcup_{i=1}^{k} C_m(z_i), z_i \in \mathbb{Z}^n\right\} \tag{3.9}$$

$$= |\{z \in \mathbb{Z}^n : A \cap C_m(z) \neq \emptyset\}| \tag{3.10}$$

$$= |[A]_m|, \tag{3.11}$$

where $[A]_m \triangleq \{[x]_m : x \in A\}$.

Lemma 9. *Let A be a bounded subset in $(\mathbb{R}^n, \|\cdot\|_p)$, $1 \leq p \leq \infty$. The Minkowski dimensions satisfy*

$$\underline{\dim}_{\mathrm{B}} A = \liminf_{m\to\infty} \frac{\log |[A]_m|}{m} \tag{3.12}$$

and

$$\overline{\dim}_{\mathrm{B}} A = \limsup_{m\to\infty} \frac{\log |[A]_m|}{m}. \tag{3.13}$$

Proof. See Appendix C.2. □

3.2 Lossless Minkowski-dimension compression

As a counterpart to lossless data compression, in this section we investigate the problem of lossless Minkowski dimension compression for general sources, where the minimum ϵ-achievable rate is defined as $\mathsf{R}_{\mathsf{B}}(\epsilon)$. This is an important intermediate tool for studying the fundamental limits of lossless linear encoding and Lipschitz decoding in Chapter 5. We present bounds for $\mathsf{R}_{\mathsf{B}}(\epsilon)$ for general sources, as well as tight results for discrete-continuous mixtures and self-similar sources.

Consider a source X^n in $\mathbb{R}^n$ equipped with an ℓ_p-norm. We define the minimum ϵ-achievable rate for Minkowski-dimension compression as follows:

Definition 7 (Minkowski-dimension compression rate)**.** Let $\{X_i : i \in \mathbb{N}\}$ be a stochastic process on $(\mathbb{R}^{\mathbb{N}}, \mathcal{B}^{\otimes\mathbb{N}})$. Define

$$\mathsf{R}_{\mathsf{B}}(\epsilon) = \limsup_{n\to\infty} \frac{\overline{\dim}^{\epsilon}_{\mathrm{B}}(P_{X^n})}{n} \tag{3.14}$$

$$= \limsup_{n\to\infty} \frac{1}{n} \inf\left\{\overline{\dim}_{\mathrm{B}} S : S \subset \mathbb{R}^n, \mathbb{P}\{X^n \in S\} \geq 1-\epsilon\right\}. \tag{3.15}$$

Remark 4 (Connections to data compression)**.** Definition 7 deals with the minimum upper Minkowski dimension of high-probability events of source realizations. This can

be seen as a counterpart of lossless source coding, which seeks the smallest cardinality of high-probability events. Indeed, when the source alphabet is some countable set $\mathcal{X}$, the conventional minimum source coding rate is defined like in (3.15) replacing $\overline{\dim}_{\mathrm{B}} S$ by $\log |S|$ (see Definition 12 in Section 5.2.1). For memoryless sources, the strong source coding theorem states that

$$H(X) = \lim_{n\to\infty} \frac{1}{n} \inf \left\{ \log |S| : S \subset \mathcal{X}^n, \mathbb{P}\left\{X^n \in S\right\} \geq 1-\epsilon \right\}, \quad 0 < \epsilon < 1. \tag{3.16}$$

In general, $0 \leq \mathsf{R}_{\mathsf{B}}(\epsilon) \leq 1$ for any $0 < \epsilon \leq 1$. This is because for any n, there exists a compact subset $S \subset \mathbb{R}^n$, such that $\mathbb{P}\left\{X^n \in S\right\} \geq 1-\epsilon$, and $\overline{\dim}_{\mathrm{B}} S \leq n$ by definition. Several coding theorems for $\mathsf{R}_{\mathsf{B}}(\epsilon)$ are given as follows.

Theorem 14. *Suppose that the source is memoryless with distribution such that $\overline{d}(X) < \infty$. Then*

$$\mathsf{R}_{\mathsf{B}}(\epsilon) \leq \lim_{\alpha\uparrow 1} \overline{d}_\alpha(X) \tag{3.17}$$

and

$$\mathsf{R}_{\mathsf{B}}(\epsilon) \geq \overline{d}(X) \tag{3.18}$$

for $0 < \epsilon < 1$.

Proof. See Section 3.3. □

For the special cases of discrete-continuous mixed and self-similar sources, we have the following tight results.

Theorem 15. *Suppose that the source is memoryless with a discrete-continuous mixed distribution. Then*

$$\mathsf{R}_{\mathsf{B}}(\epsilon) = \rho \tag{3.19}$$

for all $0 < \epsilon < 1$, where ρ is the weight of the continuous part of the distribution. If $d(X)$ is finite, then

$$\mathsf{R}_{\mathsf{B}}(\epsilon) = d(X). \tag{3.20}$$

Theorem 16. *Suppose that the source is memoryless with a self-similar distribution that satisfies the strong separation condition. Then*

$$\mathsf{R}_{\mathsf{B}}(\epsilon) = d(X) \tag{3.21}$$

for all $0 < \epsilon < 1$.

Theorem 17. *Suppose that the source is memoryless and bounded, whose M-ary expansion defined in (2.47) consists of independent (not necessarily identically distributed) digits. Then*

$$\mathsf{R}_{\mathsf{B}}(\epsilon) = \overline{d}(X) \tag{3.22}$$

for all $0 < \epsilon < 1$.

When $\overline{d}(X) = \infty$, $\mathsf{R}_{\mathsf{B}}(\epsilon)$ can take any value in $[0,1]$ for all $0 < \epsilon < 1$. Such a source can be constructed using Theorem 15 as follows: Let the distribution of X be a mixture of a continuous and a discrete distribution with weights d and $1-d$ respectively, where the discrete part is supported on $\mathbb{N}$ and has infinite entropy. Then $\overline{d}(X) = \infty$ by Theorem 1 and Theorem 4, but $\mathsf{R}_{\mathsf{B}}(\epsilon) = d$ by Theorem 15.

3.3 Proofs

Before showing the converse part of Theorem 14, we state two lemmas which are of independent interest in conventional lossless source coding theory.

Lemma 10. *Assume that $\{X_i : i \in \mathbb{N}\}$ is a discrete memoryless source with common distribution P on the alphabet $\mathcal{X}$. $H(P) < \infty$. Let $\delta \geq -H(P)$. Denote by $p_n^*(P,\delta)$ the block error probability of the optimal $(n, \lfloor (H(X)+\delta)n \rfloor)$-code, i.e.,*

$$p_n^*(P,\delta) \triangleq 1 - \sup\left\{P^n(S) : |S| \leq \exp(\lfloor (H(X)+\delta)n \rfloor)\right\}. \tag{3.23}$$

Then for any $n \in \mathbb{N}$,

$$p_n^*(P,\delta) \leq \exp[-nE_0(P,\delta)], \tag{3.24}$$
$$p_n^*(P,\delta) \geq 1 - \exp[-nE_1(P,\delta)],. \tag{3.25}$$

where the exponents are given by

$$E_0(P,\delta) = \inf_{Q:H(Q)>H(P)+\delta} D(Q||P) \tag{3.26}$$
$$= \max_{\lambda \geq 0} \lambda \left[H(P) + \delta - H_{\frac{1}{1+\lambda}}(P)\right], \tag{3.27}$$
$$E_1(P,\delta) = \inf_{Q:H(Q)<H(P)+\delta} D(Q||P) \tag{3.28}$$
$$= \max_{\lambda \geq 0} \lambda \left[H_{\frac{1}{1-\lambda}}(P) - H(P) - \delta\right]. \tag{3.29}$$

Lemma 11. *For $\delta \geq 0$,*

$$\min\left\{E_0(P,\delta), E_1(P,-\delta)\right\} \geq \min\left\{\frac{1}{8}, \frac{1}{2}\left[\frac{(\delta - \mathrm{e}^{-1}\log \mathrm{e})^+}{\log|\mathcal{X}|}\right]^2\right\} \log \mathrm{e}. \tag{3.30}$$

Proof. See Appendix C.3. □

Lemma 10 shows that the error exponents for lossless source coding are not only asymptotically tight, but also apply to *every* block length. This has been shown for rates above entropy in [55, Exercise 1.2.7, pp. 41 – 42] via a combinatorial method.

A unified proof can be given through the method of Rényi entropy, which we omit for conciseness. The idea of using Rényi entropy to study lossless source coding error exponents was previously introduced by [56, 57, 58].

Lemma 11 deals with universal lower bounds on the source coding error exponents, in the sense that these bounds are independent of the source distribution. A better bound on $E_0(P,\delta)$ has been shown in [59]: for $\delta > 0$

$$E_0(P,\delta) \geq \frac{1}{2}\left(\frac{\delta}{\log|\mathcal{X}|}\right)^2 \log \mathrm{e}. \tag{3.31}$$

However, the proof of (3.31) was based on the dual expression (3.27) and a similar lower bound of random channel coding error exponent due to Gallager [60, Exercise 5.23], which cannot be applied to $E_1(P,\delta)$. Here we give a common lower bound on both exponents, which is a consequence of Pinsker's inequality [61] combined with the lower bound on entropy difference by variational distance [55].

Proof of Theorem 14. (Converse) Let $0 < \epsilon < 1$ and abbreviate $\overline{d}(X)$ as $\overline{d}$. Suppose $\mathsf{R}_{\mathsf{B}}(\epsilon) < \overline{d} - 4\delta$ for some $\delta > 0$. Then for sufficiently large n there exists $S^n \subset \mathbb{R}^n$, such that $\mathbb{P}\{X^n \in S^n\} \geq 1-\epsilon$ and $\overline{\dim}_{\mathrm{B}} S^n \leq (\overline{d} - 3\delta)n$.

First we assume that the source has *bounded* support, that is, $|X| \leq K$ a.s. for some $K > 0$. By Lemma 1, choose M such that for all $m > M$,

$$H([X]_m) \geq (\overline{d} - \delta)m. \tag{3.32}$$

By Lemma 9, for all n, there exists $M_n > M$, such that for all $m > M_n$,

$$|[S^n]_m| \leq 2^{m(\overline{\dim}_{\mathrm{B}}(S^n)+\delta n)} \leq 2^{mn(\overline{d}-2\delta)}, \tag{3.33}$$

$$\mathbb{P}\{[X^n]_m \in [S^n]_m\} \geq \mathbb{P}\{X^n \in S^n\} \geq 1-\epsilon. \tag{3.34}$$

Then by (3.32), we have

$$|[S^n]_m| \leq 2^{(H([X]_m)-m\delta)n}. \tag{3.35}$$

Note that $[X^n]_m$ is a memoryless source with alphabet size at most $\log K + m$. By Lemmas 10 and Lemma 11, for all m, n, for all A_m^n such that $|A_m^n| \leq 2^{(H([X]_m)-m\delta)n}$, we have

$$\mathbb{P}\{[X^n]_m \in A_m^n\} \leq 2^{-nE_1(m\delta)} \tag{3.36}$$

$$\leq 2^{-\frac{n}{2}\left(\frac{(m\delta - \mathrm{e}^{-1}\log \mathrm{e})^+}{m+\log K}\right)^2}. \tag{3.37}$$

Choose m and n so large that the right hand side of (3.37) is less than $1-\epsilon$. In the special case of $A_m^n = [S^n]_m$, (3.37) contradicts (3.34) in view of (3.35).

Next we drop the restriction that X is almost surely bounded. Denote by μ the distribution of X. Let K be so large that

$$p = \mu(B(0,K)) > 1-\delta. \tag{3.38}$$

Denote the normalized restriction of μ on $B(0,K)$ by μ_1, that is,

$$\mu_1(A) = \frac{\mu(A \cap B(0,K))}{\mu(B(0,K))}, \tag{3.39}$$

and the normalized restriction of μ on $B(0,K)^{\mathrm{c}}$ by μ_2. Then

$$\mu = p\mu_1 + (1-p)\mu_2. \tag{3.40}$$

By Theorem 5,

$$\begin{aligned} \overline{d} &\leq p\overline{d}(\mu_1) + (1-p)\overline{d}(\mu_2) && (3.41)\\ &\leq p\overline{d}(\mu_1) + (1-p), && (3.42) \end{aligned}$$

where $\overline{d}(\mu_2) \leq 1$ because of the following: since $\overline{d}$ is finite, we have $\mathbb{E}\left[\log(1+|X|)\right] < \infty$ in view of Theorem 1. Consequently, $\mathbb{E}\left[\log(1+|X|)\mathbf{1}_{\{|X|>K\}}\right] < \infty$, hence $\overline{d}(\mu_2) \leq 1$.

The distribution of $\tilde{X} \triangleq X\mathbf{1}_{\{X\in B(0,K)\}}$ is given by

$$\tilde{\mu} = p\mu_1 + (1-p)\delta_0, \tag{3.43}$$

where δ_0 denotes the Dirac measure with atom at 0. Then

$$\begin{aligned} \overline{d}(\tilde{X}) &= p\overline{d}(\mu_1) && (3.44)\\ &> \overline{d} - \delta. && (3.45) \end{aligned}$$

where

- (3.44): by (3.43) and (2.35), since Theorem 4 implies $d(\delta_0) = 0$.
- (3.45): by (3.42) and (3.38).

For $T \subset \{1,\ldots,n\}$, let $S_T^n = \{x_T^n : x^n \in S^n\}$. Define

$$\tilde{S}^n = \bigcup_{T\subset\{1,\ldots,n\}} \{x^n \in \mathbb{R}^n : x_T^n \in S_T^n, x_{T^{\mathrm{c}}}^n = 0\}. \tag{3.46}$$

Then for all n,

$$\mathbb{P}\left\{\tilde{X}^n \in \tilde{S}^n\right\} \geq \mathbb{P}\left\{X^n \in S^n\right\} \geq 1-\epsilon, \tag{3.47}$$

and

$$
\begin{aligned}
&\overline{\dim}_{\mathrm{B}} \tilde{S}^n \\
&= \max_{T \subset \{1,\dots,n\}} \overline{\dim}_{\mathrm{B}} \{x^n \in \mathbb{R}^n : x_T^n \in S_T^n, x_{T^c}^n = 0\} && (3.48)\\
&\le \overline{\dim}_{\mathrm{B}} S^n && (3.49)\\
&\le (\overline{d} - 3\delta)n && (3.50)\\
&< (\overline{d}(\tilde{X}) - 2\delta)n, && (3.51)
\end{aligned}
$$

where (3.51), (3.48) and (3.49) follow from (3.45), (3.5) and (3.6) respectively. But now (3.51) and (3.47) contradict the converse part of Theorem 14 for the bounded source $\tilde{X}$, which we have already proved.

(Achievability) Recall that

$$\hat{d} = \lim_{\alpha \uparrow 1} \overline{d}_\alpha. \tag{3.52}$$

We show that for all $0 < \epsilon < 1$, for all $\delta > 0$, there exists N such that for all $n > N$, there exists $S_n \subset \mathbb{R}^n$ with $\mathbb{P}\{X^n \in S^n\} \ge 1 - \epsilon$ and $\overline{\dim}_{\mathrm{B}} S^n \le n(\hat{d} + \delta)$. Therefore $\mathsf{R}_{\mathsf{B}}(\epsilon) \le \hat{d}$ readily follows.

Suppose for now that we can construct a sequence of subsets $V_m^n \subset 2^{-m}\mathbb{Z}^n$, such that for any n the following holds:

$$|V_m^n| \le 2^{mn(\hat{d}+\delta)}, \tag{3.53}$$

$$\sum_{m \ge 1} \mathbb{P}\{[X^n]_m \notin V_m^n\} < \infty. \tag{3.54}$$

Recalling the mesh cube defined in (3.8), we denote

$$U_m^n = \sum_{z^n \in V_m^n} C_m(2^m z^n), \tag{3.55}$$

$$T_m^n = \bigcap_{j \ge m} U_j^n, \tag{3.56}$$

$$T^n = \bigcup_{m \ge 1} T_m^n = \liminf_{m \to \infty} U_m^n. \tag{3.57}$$

Then $[U_m^n]_m = V_m^n$. Now we claim that for each m, $\overline{\dim}_{\mathrm{B}} T_m^n \le (\hat{d} + \delta)n$ and $\mathbb{P}\{X^n \in T^n\} = 1$. First observe that for each $j \ge m$, $T_m^n \subset U_j^n$, therefore $\{C_j(2^j z^n) : z^n \in V_j^n\}$ covers T_m^n. Hence $|[T_m^n]_j| \le |V_j^n| \le 2^{jn(\hat{d}+\delta)}$, by (3.53). Therefore by Lemma 9,

$$\overline{\dim}_{\mathrm{B}} T_m^n = \limsup_{j \to \infty} \frac{\log |[T_m^n]_j|}{j} \le n(\hat{d} + \delta). \tag{3.58}$$

By the Borel-Cantelli lemma, (3.54) implies that $\mathbb{P}\{X^n \in T^n\} = 1$. Let $S^n = \bigcup_{m=1}^{M} T_m^n$ where M is so large that $\mathbb{P}\{X^n \in S^n\} \geq 1-\epsilon$. By the finite subadditivity[1] of upper Minkowski dimension in Lemma 8, $\overline{\dim}_{\mathrm{B}} S^n = \max_{m=1,\ldots,M} \overline{\dim}_{\mathrm{B}} T_m^n \leq n(\hat{d}+\delta)$. By the arbitrariness of δ, the ϵ-achievability of rate $\hat{d}$ is proved.

Now let us proceed to the construction of the required $\{V_m^n\}$. To that end, denote by μ_m the probability mass function of $[X]_m$. Let

$$V_m^n = \left\{ z^n \in 2^{-m}\mathbb{Z}^n : \frac{1}{mn} \log \frac{1}{\mu_m^n(\{z^n\})} \leq \hat{d} + \delta \right\} \tag{3.59}$$

$$= \left\{ z^n \in 2^{-m}\mathbb{Z}^n : \mu_m^n(\{z^n\}) \geq 2^{-mn(\hat{d}+\delta)} \right\} \tag{3.60}$$

Then, immediately (3.53) follows from $\mathbb{P}\{[X^n]_m \in V_m^n\} \leq 1$. Also, for $\alpha < 1$,

$$\mathbb{P}\{[X^n]_m \notin V_m^n\}$$

$$= \sum_{z^n \in 2^{-m}\mathbb{Z}^n} \mu_m^n(\{z^n\}) \mathbf{1}_{\{z^n \notin V_m^n\}} \tag{3.61}$$

$$\leq \sum_{z^n \in 2^{-m}\mathbb{Z}^n} \mu_m^n(\{z^n\}) \left(\frac{2^{-mn(\hat{d}+\delta)}}{\mu_m^n(\{z^n\})} \right)^{1-\alpha} \tag{3.62}$$

$$= 2^{-nm(1-\alpha)(\hat{d}+\delta)} \sum_{z^n \in 2^{-m}\mathbb{Z}^n} [\mu_m^n(\{z^n\})]^{\alpha} \tag{3.63}$$

$$= 2^{(1-\alpha)\left(H_\alpha([X^n]_m) - mn(\hat{d}+\delta)\right)} \tag{3.64}$$

$$= 2^{n(1-\alpha)\left(H_\alpha([X]_m) - m\hat{d} - m\delta\right)}. \tag{3.65}$$

where (3.65) follows from the fact that $[X^n]_m = [[X_1]_m, \ldots, [X_n]_m]^{\mathrm{T}}$ are i.i.d. and the joint Rényi entropy is the sum of individual Rényi entropies. Hence

$$\frac{\log \mathbb{P}\{[X^n]_m \notin V_m^n\}}{m} \leq (1-\alpha) \left(\frac{H_\alpha([X]_m)}{m} - \hat{d} - \delta \right) n. \tag{3.66}$$

Choose $0 < \alpha < 1$ such that $\overline{d}_\alpha < \hat{d} + \delta/2$, which is guaranteed to exist in view of (3.52) and the fact that $\overline{d}_\alpha$ is nonincreasing in α according to Lemma 4. Then

$$\limsup_{m\to\infty} \frac{\log \mathbb{P}\{[X^n]_m \notin V_m^n\}}{m}$$

$$\leq (1-\alpha) \left(\limsup_{m\to\infty} \frac{H_\alpha([X]_m)}{m} - \hat{d} - \delta \right) n \tag{3.67}$$

$$= (1-\alpha) \left(\overline{d}_\alpha - \hat{d} - \delta \right) n \tag{3.68}$$

$$\leq -(1-\alpha) n\delta/2 < 0, \tag{3.69}$$

[1] Countable subadditivity fails for upper Minkowski dimension. Had it been satisfied, we could have picked $S^n = T^n$ to achieve $\epsilon = 0$.

Hence $\mathbb{P}\{[X^n]_m \notin V_m^n\}$ decays at least exponentially with m. Accordingly, (3.54) holds, and the proof of $\mathsf{R}_{\mathsf{B}}(\epsilon) \le \hat{d}$ is complete. □

Next we prove $\mathsf{R}_{\mathsf{B}}(\epsilon) = d(X)$ for the special case of discrete-continuous mixtures. The following lemma is needed in the converse proof.

Lemma 12 ([62, Theorem 4.16])**.** *Any Borel set $A \subset \mathbb{R}^k$ whose upper Minkowski dimension is strictly less than k has zero Lebesgue measure.*

Proof of Theorem 15. (Achievability) Let the distribution of X be

$$\mu = (1-\rho)\mu_{\mathrm{d}} + \rho\mu_{\mathrm{c}}, \tag{3.70}$$

where $0 \le \rho \le 1$, μ_{c} is a probability measure on $(\mathbb{R}, \mathcal{B})$ absolutely continuous with respect to Lebesgue measure and μ_{d} is a discrete probability measure. Let $\mathcal{A}$ be the collection of all the atoms of μ_{d}, which is, by definition, a countable subset of $\mathbb{R}$.

Let $W_i = \mathbf{1}_{\{X_i \notin \mathcal{A}\}}$. Then $\{W_i\}$ is a sequence of i.i.d. binary random variables with expectation

$$\mathbb{E}W_i = \mathbb{P}\{X_i \notin \mathcal{A}\} = (1-\rho)\mu_{\mathrm{d}}(\mathcal{A}^{\mathrm{c}}) + \rho\mu_{\mathrm{c}}(\mathcal{A}^{\mathrm{c}}) = \rho. \tag{3.71}$$

By the weak law of large numbers (WLLN),

$$\frac{1}{n}|\mathrm{spt}(X^n)| = \frac{1}{n}\sum_{i=1}^{n} W_i \xrightarrow{\mathbb{P}} \rho. \tag{3.72}$$

where the *generalized support* of vector x^n is defined as

$$\mathrm{spt}(x^n) \triangleq \{i = 1, \ldots, n : x_i \notin \mathcal{A}\}. \tag{3.73}$$

Fix an arbitrary $\delta > 0$ and let

$$C_n = \{x^n \in \mathbb{R}^n : |\mathrm{spt}(x^n)| < (\rho+\delta)\, n\}. \tag{3.74}$$

By (3.72), for sufficiently large n, $\mathbb{P}\{X^n \in C_n\} \ge 1 - \epsilon/2$. Decompose C_n as:

$$C_n = \bigcup_{\substack{T \subset \{1,\ldots,n\} \\ |T| < (\rho+\delta)n}} \bigcup_{z \in \mathcal{A}^{n-|T|}} U_{T,z}, \tag{3.75}$$

where

$$U_{T,z} = \{x^n \in \mathbb{R}^n : \mathrm{spt}(x^n) = T, x^n_{T^{\mathrm{c}}} = z\}. \tag{3.76}$$

Note that the collection of all $U_{T,z}$ is *countable*, and thus we may relabel them as $\{U_j : j \in \mathbb{N}\}$. Then $C_n = \bigcup_j U_j$. Hence there exists $J \in \mathbb{N}$ and $R > 0$, such that $\mathbb{P}\{X^n \in S^n\} \ge 1 - \epsilon$, where

$$S_n = B(0, R) \cap \bigcup_{j=1}^{J} U_j. \tag{3.77}$$

Then

$$\overline{\dim}_{\mathrm{B}} S_n = \max_{j=1,\ldots,J} \overline{\dim}_{\mathrm{B}} U_j \cap B(0,R) \tag{3.78}$$

$$\leq \ (\rho+\delta)n, \tag{3.79}$$

where (3.78) is by Lemma 8, and (3.79) follows from $\overline{\dim}_{\mathrm{B}} U_j \cap B(0,R) \leq |T| < (\rho+\delta)n$ for each j. This proves the ϵ-achievability of $\rho+\delta$. By the arbitrariness of δ, $\mathsf{R}_{\mathsf{B}}(\epsilon) \leq \rho$.

(Converse) Since $\mathsf{R}_{\mathsf{B}}(\epsilon) \geq 0$, we can assume $\rho > 0$. Let $S_n \in \mathbb{R}^n$ be such that $\mathbb{P}\{X^n \in S_n\} \geq 1-\epsilon$. Define

$$T_n = \{x^n \in \mathbb{R}^n : |\mathrm{spt}(x^n)| > (\rho-\delta)\, n\}. \tag{3.80}$$

By (3.72), $\mathbb{P}\{X^n \in T_n\} \geq 1-\epsilon$ for sufficiently large n. Let $G_n = S_n \cap T_n$, then $\mathbb{P}\{X^n \in G_n\} \geq 1-2\epsilon$. Write G_n as the disjoint union:

$$G_n = \bigcup_{\substack{T \subset \{1,\ldots,n\} \\ |T| > (\rho-\delta)n}} \bigcup_{z \in \mathcal{A}^{n-|T|}} V_{T,z}, \tag{3.81}$$

where

$$V_{T,z} = U_{T,z} \cap G_n. \tag{3.82}$$

Also let $E_{T,z} = \{x_T^n : x^n \in V_{T,z}\}$. Since $\mathbb{P}\{X^n \in G_n\} \geq 1-2\epsilon > 0$, there exists $T \subset \{1,\ldots,n\}$ and $z \in \mathcal{A}^{n-|T|}$, such that $\mathbb{P}\{X^n \in V_{T,z}\} > 0$. Note that

$$\begin{aligned} &\mathbb{P}\{X^n \in V_{T,z}\} \\ &= \rho^{|T|}(1-\rho)^{n-|T|}\mu_{\mathrm{c}}^{|T|}(E_{T,z}) \prod_{i=1}^{n-|T|} \mu_{\mathrm{d}}(\{z_i\}) > 0. \end{aligned} \tag{3.83}$$

If $\rho < 1$, $\mu_{\mathrm{c}}^{|T|}(E_{T,z}) > 0$, which implies $E_{T,z} \subset \mathbb{R}^{|T|}$ has positive Lebesgue measure. By Lemma 12, $\overline{\dim}_{\mathrm{B}} E_{T,z} = |T| > (\rho-\delta)n$, hence

$$\overline{\dim}_{\mathrm{B}} S_n \geq \overline{\dim}_{\mathrm{B}} V_{T,z} = \overline{\dim}_{\mathrm{B}} E_{T,z} > (\rho-\delta)n. \tag{3.84}$$

If $\rho = 1$, $\mu = \mu_{\mathrm{c}}$ implies that $S_n \subset \mathbb{R}^n$ has positive Lebesgue measure. Thus $\overline{\dim}_{\mathrm{B}} S_n = n$. By the arbitrariness of $\delta > 0$, the proof of $\mathsf{R}_{\mathsf{B}}(\epsilon) \geq \rho$ is complete. □

Next we prove $\mathsf{R}_{\mathsf{B}}(\epsilon) = d(X)$ for self-similar sources under the same assumption of Theorem 11. This result is due to the stationarity and ergodicity of the underlying discrete process that generates the analog source distribution.

Proof of Theorem 16. By Theorem 11, $d(X)$ is finite. Therefore the converse follows from Theorem 14. To show achievability, we invoke the following definition. Define the *local dimension* of a Borel measure ν on $\mathbb{R}^n$ as the function (if the limit exists) [46, p. 169]

$$\dim_{\mathrm{loc}} \nu(x) = \lim_{r \downarrow 0} \frac{\log \nu(B(x,r))}{\log r} \tag{3.85}$$

Denote the distribution of X^n by the product measure μ^n, which is also self-similar and satisfies the strong separation theorem. By [46, Lemma 6.4(b) and Proposition 10.6],

$$\dim_{\text{loc}}\mu^n(x) = d(\mu^n) = nd(X) \tag{3.86}$$

holds for μ^n-almost every x. Define the sequence of random variables

$$D_m(X^n) = \frac{1}{m}\log\frac{1}{\mu^n(B(X^n, 2^{-m}))}. \tag{3.87}$$

Then (3.86) implies that $D_m \xrightarrow{\text{a.s.}} nd(X)$ as $m \to \infty$. Therefore for all $0 < \epsilon < 1$ and $\delta > 0$, there exists M, such that

$$\mathbb{P}\left\{\bigcap_{m=M}^{\infty}\{D_m(X^n) \le nd(X) + \delta\}\right\} \ge 1 - \epsilon. \tag{3.88}$$

Let

$$S_n = \bigcap_{m=M}^{\infty}\{x^n : D_m(x^n) \le nd(X) + \delta\} \tag{3.89}$$

$$= \bigcap_{m=M}^{\infty}\{x^n : \mu^n(B(x^n, 2^{-m})) \ge 2^{-m(nd(X)+\delta)}\}. \tag{3.90}$$

Then $\mathbb{P}\{X^n \in S_n\} \ge 1 - \epsilon$ in view of (3.88), and

$$\overline{\dim}_{\text{B}} S_n = \limsup_{m\to\infty}\frac{\log N_{S_n}(2^{-m})}{m} \tag{3.91}$$

$$\le nd(X) + \delta, \tag{3.92}$$

where (3.92) follows from (3.90) and $\mu^n(S_n) \le 1$. Hence $\mathsf{R}_{\mathsf{B}}(\epsilon) \le d(X)$ is proved. □

Finally we prove $\mathsf{R}_{\mathsf{B}}(\epsilon) = \overline{d}(X)$ for memoryless sources whose M-ary expansion consisting of independent digits.

Proof of Theorem 17. Without loss of generality, assume $M = 2$, that is, the binary expansion of X consists of independent bits. We follow the same steps as in the achievability proof of Theorem 14. Suppose for now that we can construct a sequence of subsets $\{V_m^n\}$, such that (3.54) holds and their cardinality does not exceed:

$$|V_m^n| \le 2^{n(H([X]_m)+m\delta)} \tag{3.93}$$

By the same arguments that lead to (3.58), the sets defined in (3.56) satisfy

$$\overline{\dim}_{\mathrm{B}} T_m^n = \limsup_{j\to\infty} \frac{\log |[T_m^n]_j|}{j} \tag{3.94}$$

$$= n \limsup_{j\to\infty} \frac{H([X]_j) + j\delta}{j} \tag{3.95}$$

$$= n(\overline{d} + \delta). \tag{3.96}$$

Since $\lim_{m\to\infty} \mathbb{P}\{X^n \in T_m^n\} = 1$, this shows the ϵ-achievability of $\overline{d}$.

Next we proceed to construct the required $\{V_m^n\}$. Applying Lemma 10 to the discrete memoryless source $[X]_m$ and blocklength n yields

$$p_n^*([X]_m, m\delta) \le 2^{-nE_0(m\delta)}. \tag{3.97}$$

By the assumption that X is bounded, without loss of generality, we shall assume $|X| \le 1$ a.s. Therefore the alphabet size of $[X]_m$ is at most 2^m. Simply applying (3.31) to $[X]_m$ yields the lower bound

$$E_0([X]_m, m\delta) \ge \frac{\delta^2}{2}, \tag{3.98}$$

which does not grow with m and cannot suffice for our purpose of constructing V_m^n. Exploiting the structure that $[X]_m$ consists of independent bits, we can show a much better bound:

$$E_0([X]_m, m\delta) \ge \frac{m\delta^2}{2}. \tag{3.99}$$

Then by (3.97) and (3.99), there exists V_m^n, such that (3.93) holds and

$$\mathbb{P}\{[X^n]_m \notin V_m^n\} \le 2^{-\frac{nm\delta^2}{2}}, \tag{3.100}$$

which implies (3.54). This yields the desired $\mathsf{R}_{\mathsf{B}}(\epsilon) \le \overline{d}$.

To finish the proof, it remains to establish (3.99). By (3.31), for all $\delta \in \mathbb{R}$,

$$E_0(P, \delta) \ge \frac{c(\delta)}{\log^2 |\mathcal{X}|}, \tag{3.101}$$

where

$$c(x) = \frac{\log \mathrm{e}}{2}(x^+)^2, \tag{3.102}$$

which is a nonnegative nondecreasing convex function.

Since $|X| \le 1$ a.s., $[X]_m$ is in one-to-one correspondence with $[(X)_1, \ldots, (X)_m]$. Denote the distribution of $[(X)_1, \ldots, (X)_m]$ and $(X)_i$ by P^m and P_i respectively. By assumption, $(X)_1, \ldots, (X)_m$ are independent, hence

$$P^m = P_1 \times \cdots \times P_m. \tag{3.103}$$

By (3.26),

$$E_0(P^m, m\delta) = \min_{Q^m : H(Q^m) \geq H([X]_m) + m\delta} D(Q^m || P^m), \tag{3.104}$$

where Q^m is a distribution on $\mathbb{Z}_2^m$. Denote the marginals of Q^m by $\{Q_1, \ldots, Q_m\}$. Combining (3.103) with properties of entropy and relative entropy, we have

$$D(Q^m || P^m) \geq \sum_{i=1}^{m} D(Q_i || P_i), \tag{3.105}$$

and

$$H(P^m) = \sum_{i=1}^{m} H(P_i) \tag{3.106}$$

$$H(Q^m) \leq \sum_{i=1}^{m} H(Q_i) \tag{3.107}$$

Therefore,

$$\begin{aligned}
E_0(P^m, m\delta) = & \min_{H(Q^m) \geq H([X]_m) + m\delta} D(Q^m || P^m) && (3.108)\\
\geq & \min_{H(Q^m) \geq H([X]_m) + m\delta} \sum_{i=1}^{m} D(Q_i || P_i) && (3.109)\\
\geq & \min_{\sum \alpha_i = H([X]_m) + m\delta} \sum_{i=1}^{m} \min_{H(Q_i) \geq \alpha_i} D(Q_i || P_i) && (3.110)\\
\geq & \min_{\sum \alpha_i = H([X]_m) + m\delta} \sum_{i=1}^{m} E_0(P_i, \alpha_i - H(P_i)) && (3.111)\\
\geq & \min_{\sum \alpha_i = H([X]_m) + m\delta} \sum_{i=1}^{m} c(\alpha_i - H(P_i)) && (3.112)\\
= & \min_{\sum \beta_i = m\delta} \sum_{i=1}^{m} c(\beta_i) && (3.113)\\
\geq & \; mc(\delta) && (3.114)\\
= & \; \frac{1}{2} m\delta^2 \log \mathrm{e}, && (3.115)
\end{aligned}$$

where

- (3.109): by (3.105).
- (3.110): by (3.107), we have

$$\begin{aligned}
& \{Q^m : H(Q^m) \geq H([X]_m) + m\delta\} \\
\subset & \left\{ Q^m : H(Q_i) \geq \alpha_i, \; \sum \alpha_i = H([X]_m) + m\delta \right\}. \qquad (3.116)
\end{aligned}$$

- (3.112): by (3.101).
- (3.113): let $\beta_i = \alpha_i - H(X_i)$. Then $\sum \beta_i = m\delta$, by (3.106).
- (3.114): due to the convexity of $x \mapsto (x^+)^2$.

$\square$

Chapter 4

MMSE dimension

In this chapter we introduce the MMSE dimension, an information measure that governs the high-SNR asymptotics of MMSE of estimating a random variable corrupted by additive noise. It is also the asymptotic ratio of nonlinear MMSE to linear MMSE. Section 4.4 explores the relationship between the MMSE dimension and the information dimension. Section 4.5 gives results on (conditional) MMSE dimension for various input distributions. When the noise is Gaussian, we show that the MMSE dimension and the information dimension coincide whenever the input distribution has no singular component. Technical proofs are relegated to Appendix B. The material in this chapter has been presented in part in [63, 28, 64].

4.1 Basic setup

The minimum mean square error (MMSE) plays a pivotal role in estimation theory and Bayesian statistics. Due to the lack of closed-form expressions for posterior distributions and conditional expectations, exact MMSE formulae are scarce. Asymptotic analysis is more tractable and sheds important insights about how the fundamental estimation-theoretic limits depend on the input and noise statistics. The theme of this chapter is the high-SNR scaling law of MMSE of estimating a random variable based on observations corrupted by additive noise.

The MMSE of estimating X based on Y is denoted by[1]

$$\mathsf{mmse}(X|Y) = \inf_f \mathbb{E}\left[(X - f(Y))^2\right] \tag{4.1}$$

$$= \mathbb{E}\left[(X - \mathbb{E}[X|Y])^2\right] = \mathbb{E}\left[\mathsf{var}(X|Y)\right], \tag{4.2}$$

[1]Throughout this chapter, the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is fixed. For $p \geq 1$, $L^p(\Omega)$ denotes the collection of all random variables defined on $(\Omega, \mathcal{F}, \mathbb{P})$ with finite p^{th} moments. $L^\infty(\Omega)$ denotes the collection of all almost surely (a.s.) bounded random variables. f_X denotes the density of a random variable X whose distribution is absolutely continuous with respect to Lebesgue measure. $\mu \perp \nu$ denotes that μ and ν are mutually singular, i.e., there exists a measurable set E such that $\mu(E) = 1$ and $\nu(E) = 0$. N_{G} denotes a standard normal random variable. For brevity, in this chapter natural logarithms are adopted and information units are nats.

where the infimum in (4.1) is over all Borel measurable f. When Y is related to X through an additive-noise channel with gain $\sqrt{\mathsf{snr}}$, i.e.,

$$Y = \sqrt{\mathsf{snr}}X + N \tag{4.3}$$

where N is independent of X, we denote

$$\mathsf{mmse}(X, N, \mathsf{snr}) \triangleq \mathsf{mmse}(X|\sqrt{\mathsf{snr}}X + N), \tag{4.4}$$

and, in particular, when the noise is Gaussian, we simplify

$$\mathsf{mmse}(X, \mathsf{snr}) \triangleq \mathsf{mmse}(X, N_{\mathsf{G}}, \mathsf{snr}). \tag{4.5}$$

When the estimator has access to some side information U such that N is independent of $\{X, U\}$, we define

$$\mathsf{mmse}(X, N, \mathsf{snr}|U) \triangleq \mathsf{mmse}(X|\sqrt{\mathsf{snr}}X + N, U) \tag{4.6}$$

$$= \mathbb{E}[(X - \mathbb{E}[X|\sqrt{\mathsf{snr}}X + N, U])^2]. \tag{4.7}$$

Assuming $0 < \mathsf{var}N < \infty$, next we proceed to define the (conditional) MMSE dimension formally. We focus particular attention on the case of Gaussian noise. First note the following general inequality:

$$0 \leq \mathsf{mmse}(X, N, \mathsf{snr}|U) \leq \frac{\mathsf{var}N}{\mathsf{snr}}. \tag{4.8}$$

where the rightmost side can be achieved using the affine estimator $f(y) = \frac{y - \mathbb{E}N}{\sqrt{\mathsf{snr}}}$. Therefore as $\mathsf{snr} \to \infty$, it holds that:

$$\mathsf{mmse}(X, N, \mathsf{snr}|U) = O\left(\frac{1}{\mathsf{snr}}\right). \tag{4.9}$$

We are interested in the exact scaling constant, which depends on the distribution of X, U and N. To this end, we introduce the following notion:

Definition 8. Define the *upper* and *lower MMSE dimension* of the pair (X, N) as follows:

$$\overline{\mathscr{D}}(X, N) = \limsup_{\mathsf{snr}\to\infty} \frac{\mathsf{snr} \cdot \mathsf{mmse}(X, N, \mathsf{snr})}{\mathsf{var}N}, \tag{4.10}$$

$$\underline{\mathscr{D}}(X, N) = \liminf_{\mathsf{snr}\to\infty} \frac{\mathsf{snr} \cdot \mathsf{mmse}(X, N, \mathsf{snr})}{\mathsf{var}N}. \tag{4.11}$$

If $\overline{\mathscr{D}}(X, N) = \underline{\mathscr{D}}(X, N)$, the common value is denoted by $\mathscr{D}(X, N)$, called the *MMSE dimension* of (X, N). In particular, when N is Gaussian, we denote these limits by $\overline{\mathscr{D}}(X), \underline{\mathscr{D}}(X)$ and $\mathscr{D}(X)$, called the *upper*, *lower* and *MMSE dimension* of X respectively.

Replacing $\mathsf{mmse}(X, N, \mathsf{snr})$ by $\mathsf{mmse}(X, N, \mathsf{snr}|U)$, the *conditional MMSE dimension* of (X, N) given U can be defined similarly, denoted by $\overline{\mathscr{D}}(X, N|U), \underline{\mathscr{D}}(X, N|U)$ and $\mathscr{D}(X, N|U)$ respectively. When N is Gaussian, we denote them by $\overline{\mathscr{D}}(X|U), \underline{\mathscr{D}}(X|U)$ and $\mathscr{D}(X|U)$, called the *upper*, *lower* and *conditional MMSE dimension* of X given U respectively.

The MMSE dimension of X governs the high-SNR scaling law of $\mathsf{mmse}(X, \mathsf{snr})$. If $\mathscr{D}(X)$ exists, we have

$$\mathsf{mmse}(X, \mathsf{snr}) = \frac{\mathscr{D}(X)}{\mathsf{snr}} + o\left(\frac{1}{\mathsf{snr}}\right). \tag{4.12}$$

As we show in Section 4.3.2, MMSE dimension also characterizes the high-SNR sub-optimality of linear estimation.

The following proposition is a simple consequence of (4.8):

Theorem 18.

$$0 \le \underline{\mathscr{D}}(X, N|U) \le \overline{\mathscr{D}}(X, N|U) \le 1. \tag{4.13}$$

The next result shows that MMSE dimension is invariant under translations and positive scaling of input and noise.

Theorem 19. *For any $\alpha, \beta, \gamma, \eta \in \mathbb{R}$, if either*

- $\alpha\beta > 0$

or

- $\alpha\beta \neq 0$ *and either X or N has a symmetric distribution,*

then

$$\underline{\mathscr{D}}(X, N) = \underline{\mathscr{D}}(\alpha X + \gamma, \beta N + \eta), \tag{4.14}$$

$$\overline{\mathscr{D}}(X, N) = \overline{\mathscr{D}}(\alpha X + \gamma, \beta N + \eta). \tag{4.15}$$

Proof. Appendix B.1. □

Remark 5. Particularizing Theorem 19 to Gaussian noise, we obtain the scale-invariance of MMSE dimension: $\mathscr{D}(\alpha X) = \mathscr{D}(X)$ for all $\alpha \neq 0$, which can be generalized to random vectors and all *orthogonal* transformations. However, it is unclear whether the invariance of information dimension under bi-Lipschitz mappings established in Theorem 2 holds for MMSE dimension.

Although most of our focus is on square integrable random variables, the functional $\mathsf{mmse}(X, N, \mathsf{snr})$ can be defined for infinite-variance noise. Consider $X \in L^2(\Omega)$ but $N \notin L^2(\Omega)$. Then $\mathsf{mmse}(X, N, \mathsf{snr})$ is still finite but (4.10) and (4.11) cease to make sense. Hence the scaling law in (4.9) could fail. It is instructive to consider the following example:

Example 1. Let X be uniformly distributed in $[0,1]$ and N have the following density:

$$f_{N_\alpha}(z) = \frac{\alpha-1}{z^\alpha}\mathbf{1}_{\{z>1\}}, \quad 1<\alpha\le 3. \tag{4.16}$$

Then $\mathbb{E}N_\alpha^2 = \infty$. As α decreases, the tail of N_α becomes heavier and accordingly $\mathsf{mmse}(X, N_\alpha, \mathsf{snr})$ decays slower. For instance, for $\alpha = 2$ and $\alpha = 3$ we obtain (see Appendix B.2)

$$\mathsf{mmse}(X, N_2, \mathsf{snr}) = \frac{3+\pi^2}{18}\frac{1}{\sqrt{\mathsf{snr}}} - \frac{\log^2 \mathsf{snr}}{4\mathsf{snr}} + \Theta\left(\frac{\log \mathsf{snr}}{\mathsf{snr}}\right), \tag{4.17}$$

$$\mathsf{mmse}(X, N_3, \mathsf{snr}) = \frac{\log \mathsf{snr}}{\mathsf{snr}} - \frac{2(2+\log 2)}{\mathsf{snr}} + \Theta\left(\frac{1}{\mathsf{snr}^{3/2}}\right) \tag{4.18}$$

respectively. Therefore in both cases the MMSE decays strictly slower than $\frac{1}{\mathsf{snr}}$, i.e., for $N = N_2$ or N_3,

$$\mathsf{mmse}(X, N, \mathsf{snr}) = \omega\left(\frac{1}{\mathsf{snr}}\right). \tag{4.19}$$

However, $\mathsf{var}N = \infty$ does not alway imply (4.19). For example, consider an arbitrary integer-valued N (none of whose moments may exist) and $0 < X < 1$ a.s. Then $\mathsf{mmse}(X, N, \mathsf{snr}) = 0$ for all $\mathsf{snr} > 1$.

4.2 Related work

The low-SNR asymptotics of $\mathsf{mmse}(X, \mathsf{snr})$ has been studied extensively in [65, 66]. In particular, it is shown in [66, Proposition 7] that if all moments of X are finite, then $\mathsf{mmse}(X, \cdot)$ is smooth on $\mathbb{R}_+$ and admits a Taylor expansion at $\mathsf{snr} = 0$ up to arbitrarily high order. For example, if $\mathbb{E}[X] = 0$ and $\mathbb{E}[X^2] = 1$, then as $\mathsf{snr} \to 0$,

$$\begin{aligned}\mathsf{mmse}(X, \mathsf{snr}) = 1 - \mathsf{snr} + (2 - (\mathbb{E}X^3)^2)\frac{\mathsf{snr}^2}{2} \\ - \left(15 - 12(\mathbb{E}X^3)^2 - 6\,\mathbb{E}X^4 + (\mathbb{E}X^4)^2\right)\frac{\mathsf{snr}^3}{6} + O(\mathsf{snr}^4).\end{aligned} \tag{4.20}$$

However, the asymptotics in the high-SNR regime remain underexplored in the literature. In [65, p. 1268] it is pointed out that the high-SNR behavior depends on the input distribution: For example, for binary X, $\mathsf{mmse}(X, \mathsf{snr})$ decays exponentially, while for standard Gaussian X, $\mathsf{mmse}(X, \mathsf{snr}) = \frac{1}{\mathsf{snr}+1}$. Unlike the low-SNR regime, the high-SNR behavior is considerably more complicated, as it depends on the measure-theoretical structure of the input distribution rather than moments. On the other hand, it can be shown that the high-SNR asymptotics of MMSE is equivalent to the low-SNR asymptotics when the input and noise distributions are switched. As we show in Section 4.3.4, this simple observation yields new low-SNR results for Gaussian input contaminated by non-Gaussian noise.

The asymptotic behavior of Fisher's information (closely related to MMSE when N is Gaussian) is conjectured in [26, p. 755] to satisfy

$$\lim_{\mathsf{snr}\to\infty} J(\sqrt{\mathsf{snr}}X + N_{\mathsf{G}}) = 1 - \overline{d}(X). \tag{4.21}$$

As a corollary to our results, we prove this conjecture when X has no singular components. The Cantor distribution gives a counterexample to the general conjecture (Section 4.6.1).

Other than our scalar Bayesian setup, the weak-noise asymptotics of optimal estimation/filtering error has been studied in various regimes in statistics. One example is filtering a deterministic signal observed in weak additive white Gaussian noise (AWGN): Pinsker's theorem ([67, 68]) establishes the exact asymptotics of the optimal minimax square error when the signal belongs to a Sobolev class with finite duration. For AWGN channels and stationary Markov input processes that satisfy a stochastic differential equation, it is shown in [69, p. 372] that, under certain regularity conditions, the filtering MMSE decays as $\Theta\left(\frac{1}{\sqrt{\mathsf{snr}}}\right)$.

4.3 Various connections

4.3.1 Asymptotic statistics

The high-SNR behavior of $\mathsf{mmse}(X, \mathsf{snr})$ is equivalent to the behavior of quadratic Bayesian risk for the Gaussian location model in the large sample limit, where P_X is the prior distribution and the sample size n plays the role of snr. To see this, let $\{N_i : i \in \mathbb{N}\}$ be a sequence of i.i.d. standard Gaussian random variables independent of X and denote $Y_i = X + N_i$ and $Y^n = (Y_1, \ldots, Y_n)$. By the sufficiency of sample mean $\bar{Y} = \frac{1}{n}\sum_{i=1}^n Y_i$ in Gaussian location models, we have

$$\mathsf{mmse}(X|Y^n) = \mathsf{mmse}(X|\bar{Y}) = \mathsf{mmse}(X, n), \tag{4.22}$$

where the right-hand side is the function $\mathsf{mmse}(X, \cdot)$ defined in (4.5) evaluated at n. Therefore as the sample size grows, the Bayesian risk of estimating X vanishes as $O\left(\frac{1}{n}\right)$ with the scaling constant given by the *MMSE dimension of the prior*[2]

$$\mathscr{D}(X) = \lim_{n\to\infty} n\,\mathsf{mmse}(X|Y^n). \tag{4.23}$$

The asymptotic expansion of $\mathsf{mmse}(X|Y^n)$ has been studied in [70, 71, 72] for absolutely continuous priors and general models where X and Y are not necessarily related by additive Gaussian noise. Further comparison to our results is given in Section 4.5.3.

[2]In view of Lemma 13 in Section 4.6, the limit of (4.23) is unchanged even if n takes non-integer values.

4.3.2 Optimal linear estimation

MMSE dimension characterizes the gain achievable by nonlinear estimation over linear estimation in the high-SNR regime. To see this, define the linear MMSE (LMMSE) as

$$\mathsf{lmmse}(X|Y) = \min_{a,b\in\mathbb{R}} \mathbb{E}[(aY+b-X)^2]. \tag{4.24}$$

When $Y = \sqrt{\mathsf{snr}}X + N$, direct optimization over a and b reveals that the best parameters are given by

$$\hat{a} = \frac{\sqrt{\mathsf{snr}}\,\mathsf{var}X}{\mathsf{snr}\,\mathsf{var}X + \mathsf{var}N}, \tag{4.25}$$

$$\hat{b} = \frac{\mathbb{E}[X]\,\mathsf{var}N - \sqrt{\mathsf{snr}}\,\mathbb{E}[N]\,\mathsf{var}X}{\mathsf{snr}\,\mathsf{var}X + \mathsf{var}N}. \tag{4.26}$$

Hence

$$\mathsf{lmmse}(X,N,\mathsf{snr}) \triangleq \mathsf{lmmse}(X|\sqrt{\mathsf{snr}}X+N) \tag{4.27}$$

$$= \frac{\mathsf{var}N\ \mathsf{var}X}{\mathsf{snr}\ \mathsf{var}X + \mathsf{var}N} \tag{4.28}$$

$$= \frac{\mathsf{var}N}{\mathsf{snr}} + o\left(\frac{1}{\mathsf{snr}}\right), \tag{4.29}$$

as $\mathsf{snr}\to\infty$. As long as $\mathsf{var}N<\infty$, the above analysis holds for any input X even if $\mathsf{var}X=\infty$, in which case (4.28) simplifies to $\mathsf{lmmse}(X|\sqrt{\mathsf{snr}}X+N) = \frac{\mathsf{var}N}{\mathsf{snr}}$. In view of Definition 8, (4.29) gives a more general definition of MMSE dimension:

$$\mathscr{D}(X,N) = \lim_{\mathsf{snr}\to\infty} \frac{\mathsf{mmse}(X,N,\mathsf{snr})}{\mathsf{lmmse}(X,N,\mathsf{snr})}, \tag{4.30}$$

which can be easily generalized to random vectors or processes.

It is also possible to show that (4.30) carries over to non-square loss functions that behaves quadratically near zero. To this end, consider the loss function $\ell(x,y) = \rho(x-y)$ where $\rho\in C^2(\mathbb{R})$ is convex with $\rho'(0)=0$ and $\rho''(0)>0$. The Bayesian risk of estimating X based on $\sqrt{\mathsf{snr}}X+N$ with respect to ρ is

$$R(X,N,\mathsf{snr}) = \inf_f \mathbb{E}\left[\rho(X - f(\sqrt{\mathsf{snr}}X+N))\right], \tag{4.31}$$

while the linear Bayesian risk is

$$R_{\mathrm{L}}(X,N,\mathsf{snr}) = \inf_{a,b\in\mathbb{R}} \mathbb{E}\left[\rho(X - a(\sqrt{\mathsf{snr}}X+N) - b)\right]. \tag{4.32}$$

We conjecture that (4.30) holds with mmse and lmmse replaced by R and R_{L} respectively.

4.3.3 Asymptotics of Fisher information

In the special case of Gaussian noise it is interesting to draw conclusions on the asymptotic behavior of Fisher's information based on our results. Recall that the Fisher information (with respect to the location parameter) of a random variable Z is defined as: [73, Definition 4.1]

$$J(Z) = \sup\{|\mathbb{E}\left[\psi'(Z)\right]|^2 : \psi \in C^1,\ \mathbb{E}\left[\psi^2(Z)\right] = 1\} \tag{4.33}$$

where C^1 denotes the collection of all continuously differentiable functions. When Z has an absolutely continuous density f_Z, we have $J(Z) = \int \frac{{f_Z'}^2}{f_Z}$. Otherwise, $J(Z) = \infty$ [73].

In view of the representation of MMSE by the Fisher information of the channel output with additive Gaussian noise [74, (1.3.4)], [65, (58)]:

$$\mathsf{snr} \cdot \mathsf{mmse}(X, \mathsf{snr}) = 1 - J(\sqrt{\mathsf{snr}}X + N_{\mathsf{G}}) \tag{4.34}$$

and $J(aZ) = a^{-2}J(Z)$, letting $\epsilon = \frac{1}{\sqrt{\mathsf{snr}}}$ yields

$$\mathsf{mmse}(X, \epsilon^{-2}) = \epsilon^2 - \epsilon^4 J(X + \epsilon N_{\mathsf{G}}). \tag{4.35}$$

By the lower semicontinuity of Fisher information [73, p. 79], when the distribution of X is *not* absolutely continuous, $J(X + \epsilon N_{\mathsf{G}})$ diverges as ϵ vanishes, but no faster than

$$J(X + \epsilon N_{\mathsf{G}}) \leq \epsilon^{-2}, \tag{4.36}$$

because of (4.35). Similarly to the MMSE dimension, we can define the *Fisher dimension* of a random variable X as follows:

$$\overline{\mathscr{J}}(X) = \limsup_{\epsilon \downarrow 0} \epsilon^2 \cdot J(X + \epsilon N_{\mathsf{G}}), \tag{4.37}$$

$$\underline{\mathscr{J}}(X) = \liminf_{\epsilon \downarrow 0} \epsilon^2 \cdot J(X + \epsilon N_{\mathsf{G}}). \tag{4.38}$$

Equation (4.35) shows Fisher dimension and MMSE dimension are complementary of each other:

$$\overline{\mathscr{J}}(X) + \underline{\mathscr{D}}(X) = \underline{\mathscr{J}}(X) + \overline{\mathscr{D}}(X) = 1. \tag{4.39}$$

In [26, p.755] it is conjectured that

$$\underline{\mathscr{J}}(X) = 1 - \overline{d}(X) \tag{4.40}$$

or equivalently

$$\overline{\mathscr{D}}(X) = \overline{d}(X). \tag{4.41}$$

According to Theorem 4, this holds for distributions without singular components but not in general. Counterexamples can be found for singular X. See Section 4.5.5 for more details.

4.3.4 Duality to low-SNR asymptotics

Note that

$$\mathsf{snr} \cdot \mathsf{mmse}(X, N, \mathsf{snr}) = \mathsf{mmse}(\sqrt{\mathsf{snr}}X | \sqrt{\mathsf{snr}}X + N) \tag{4.42}$$

$$= \mathsf{mmse}(N | \sqrt{\mathsf{snr}}X + N) \tag{4.43}$$

$$= \mathsf{mmse}\left(N, X, \mathsf{snr}^{-1}\right), \tag{4.44}$$

which gives an equivalent definition of the MMSE dimension:

$$\mathscr{D}(X, N) = \lim_{\epsilon \downarrow 0} \mathsf{mmse}(N, X, \epsilon). \tag{4.45}$$

This reveals an interesting duality: the high-SNR MMSE scaling constant is equal to the low-SNR limit of MMSE when the roles of input and noise are switched. Restricted to the Gaussian channel, it amounts to studying the asymptotic MMSE of estimating a Gaussian random variable contaminated with *strong noise with an arbitrary distribution*. On the other end of the spectrum, the asymptotic expansion of the MMSE of an arbitrary random variable contaminated with *strong Gaussian noise* is studied in [66, Section V.A]. The asymptotics of other information measures have also been studied: For example, the asymptotic Fisher information of Gaussian (or other continuous) random variables under *weak arbitrary noise* was investigated in [75]. The asymptotics of non-Gaussianness (see (2.68)) in this regime is studied in [76, Theorem 1]. The second-order asymptotics of mutual information under *strong Gaussian noise* is studied in [77, Section IV].

Unlike $\mathsf{mmse}(X, \mathsf{snr})$ which is monotonically decreasing with snr, $\mathsf{mmse}(N_\mathsf{G} | \sqrt{\mathsf{snr}}X + N_\mathsf{G})$ may be increasing in snr (Gaussian X), decreasing (binary-valued X) or oscillatory (Cantor X) in the high-SNR regime (see Fig. 4.3). In those cases in which $\mathsf{snr} \mapsto \mathsf{mmse}(N_\mathsf{G} | \sqrt{\mathsf{snr}}X + N_\mathsf{G})$ is monotone, MMSE dimension and information dimension exist and coincide, in view of (4.45) and Theorem 20.

4.4 Relationship between MMSE dimension and Information Dimension

The following theorem reveals a connection between the MMSE dimension and the information dimension of the input:

Theorem 20. *If $H(\lfloor X \rfloor) < \infty$, then*

$$\underline{\mathscr{D}}(X) \le \underline{d}(X) \le \overline{d}(X) \le \overline{\mathscr{D}}(X). \tag{4.46}$$

Therefore, if $\mathscr{D}(X)$ exists, then $d(X)$ exists and

$$\mathscr{D}(X) = d(X), \tag{4.47}$$

and equivalently, as $\mathsf{snr} \to \infty$,

$$\mathsf{mmse}(X, \mathsf{snr}) = \frac{d(X)}{\mathsf{snr}} + o\left(\frac{1}{\mathsf{snr}}\right). \tag{4.48}$$

In view of (4.13) and (2.14), (4.46) fails when $H(\lfloor X \rfloor) = \infty$. A Cantor-distributed X provides an example in which the inequalities in (4.46) are strict (see Section 4.5.5). However, whenever X has a discrete-continuous mixed distribution, (4.47) holds and the information dimension governs the high-SNR asymptotics of MMSE. It is also worth pointing out that (4.46) need not hold when the noise is non-Gaussian. See Fig. 4.2 for a counterexample when the noise is uniformly distributed on the unit interval.

The proof of Theorem 20 hinges on two crucial results:

- The I-MMSE relationship [65]:[3]

$$\frac{\mathrm{d}}{\mathrm{d}\mathsf{snr}} I(\mathsf{snr}) = \frac{1}{2}\mathsf{mmse}(\mathsf{snr}). \tag{4.49}$$

  where we have used the following short-hand notations:

$$I(\mathsf{snr}) = I(X; \sqrt{\mathsf{snr}}X + N_{\mathsf{G}}), \tag{4.50}$$
$$\mathsf{mmse}(\mathsf{snr}) = \mathsf{mmse}(X, \mathsf{snr}). \tag{4.51}$$

- The high-SNR scaling law of $I(\mathsf{snr})$ in Theorem 9.

Before proceeding to the proof, we first outline a naïve attempt at proving that MMSE dimension and information dimension coincide. Assuming

$$\lim_{\mathsf{snr}\to\infty} \frac{I(\mathsf{snr})}{\log \mathsf{snr}} = \frac{d(X)}{2}, \tag{4.52}$$

it is tempting to apply the l'Hôpital's rule to (4.52) to conclude

$$\lim_{\mathsf{snr}\to\infty} \frac{\frac{\mathrm{d}}{\mathrm{d}\mathsf{snr}} I(\mathsf{snr})}{\frac{1}{\mathsf{snr}}} = \frac{d(X)}{2}, \tag{4.53}$$

which, combined with (4.49), would produce the desired result in (4.47). However, this approach fails because applying l'Hôpital's rule requires establishing the existence of the limit in (4.53) in the first place. In fact, we show in Section 4.5.5 when X has certain singular (e.g., Cantor) distribution, the limit in (4.53), i.e. the MMSE

[3] The previous result in [65, Theorem 1] requires $\mathbb{E}\left[X^2\right] < \infty$, which, as shown in [64], can be weakened to $H(\lfloor X \rfloor) < \infty$.

dimension, does not exist because of oscillation. Nonetheless, because mutual information is related to MMSE through an integral relation, the information dimension does exist since oscillation in MMSE is smoothed out by the integration.

In fact it is possible to construct a function $I(\mathsf{snr})$ which satisfies all the monotonicity and concavity properties of mutual information [65, Corollary 1]:

$$I > 0, \quad I' > 0, \quad I'' < 0. \tag{4.54}$$

$$I(0) = I'(\infty) = 0, \quad I(\infty) = \infty; \tag{4.55}$$

yet the limit in (4.53) does not exist because of oscillation. For instance,

$$I(\mathsf{snr}) = \frac{d}{2}\left[\log(1+\mathsf{snr}) + \frac{1-\cos(\log(1+\mathsf{snr}))}{2}\right] \tag{4.56}$$

satisfies (4.52), (4.54) and (4.55), but (4.53) fails, since

$$I'(\mathsf{snr}) = \frac{d\left[2+\sin(\log(1+\mathsf{snr}))\right]}{4(1+\mathsf{snr})}. \tag{4.57}$$

In fact (4.57) gives a correct quantitative depiction of the oscillatory behavior of $\mathsf{mmse}(X, \mathsf{snr})$ for Cantor-distributed X (see Fig. 4.1).

Proof of Theorem 20. First we prove (4.46). By (4.49), we obtain

$$I(\mathsf{snr}) = \frac{1}{2}\int_0^{\mathsf{snr}} \mathsf{mmse}(\gamma)\mathrm{d}\gamma. \tag{4.58}$$

By definition of $\underline{\mathscr{D}}(X)$, for all $\delta > 0$, there exists γ_0, such that $\gamma\,\mathsf{mmse}(\gamma) \geq \underline{\mathscr{D}}(X) - \delta$ holds for all $\gamma > \gamma_0$. Then by (4.58), for any $\mathsf{snr} > \gamma_0$,

$$I(\mathsf{snr}) \geq \frac{1}{2}\int_0^{\gamma_0} \mathsf{mmse}(\gamma)\mathrm{d}\gamma + \frac{1}{2}\int_{\gamma_0}^{\mathsf{snr}} \frac{\underline{\mathscr{D}}(X)-\delta}{\gamma}\mathrm{d}\gamma \tag{4.59}$$

$$\geq \frac{1}{2}(\underline{\mathscr{D}}(X)-\delta)\log\mathsf{snr} + O(1). \tag{4.60}$$

In view of (2.67), dividing both sides by $\frac{1}{2}\log\mathsf{snr}$ and by the arbitrariness of $\delta > 0$, we conclude that $\underline{\mathscr{D}}(X) \leq \underline{d}(X)$. Similarly, $\overline{\mathscr{D}}(X) \geq \overline{d}(X)$ holds.

Next, assuming the existence of $\mathscr{D}(X)$, (4.47) simply follows from (4.46), which can also be obtained by applying l'Hôpital's rule to (4.52) and (4.53). □

4.5 Evaluation of MMSE dimension

In this section we drop the assumption of $H(\lfloor X \rfloor) < \infty$ and proceed to give results for various input and noise distributions.

4.5.1 Data processing lemma for MMSE dimension

Theorem 21. *For any X, U and any $N \in L^2(\Omega)$,*

$$\overline{\mathscr{D}}(X, N) \geq \overline{\mathscr{D}}(X, N|U), \tag{4.61}$$

$$\underline{\mathscr{D}}(X, N) \geq \underline{\mathscr{D}}(X, N|U). \tag{4.62}$$

When N is Gaussian and U is discrete, (4.61) and (4.62) hold with equality.

Proof. Appendix B.3. □

In particular, Theorem 21 states that no *discrete* side information can reduce the MMSE dimension. Consequently, we have

$$\mathscr{D}(X|[X]_m) = \mathscr{D}(X) \tag{4.63}$$

for any $m \in \mathbb{N}$, that is, knowing arbitrarily finitely many digits X does not reduce its MMSE dimension. This observation agrees with our intuition: when the noise is weak, $[X]_m$ can be estimated with exponentially small error (see (4.68)), while the fractional part $X - [X]_m$ is the main contributor to the estimation error.

It is possible to extend Theorem 21 to non-Gaussian noise (e.g., uniform, exponential or Cauchy distributions) and more general models. See Remark 33 at the end of Appendix B.3.

4.5.2 Discrete input

Theorem 22. *If X is discrete (finitely or countably infinitely valued), and $N \in L^2(\Omega)$ whose distribution is absolutely continuous with respect to Lebesgue measure, then*

$$\mathscr{D}(X, N) = 0. \tag{4.64}$$

In particular,

$$\mathscr{D}(X) = 0. \tag{4.65}$$

Proof. Since constants have zero MMSE dimension, $\mathscr{D}(X) = \mathscr{D}(X|X) = 0$, in view of Theorem 21. The more general result in (4.64) is proved in Appendix B.4. □

Remark 6. Theorem 22 implies that $\mathsf{mmse}(X, N, \mathsf{snr}) = o\left(\frac{1}{\mathsf{snr}}\right)$ as $\mathsf{snr} \to \infty$. As observed in Example 4, MMSE can decay much faster than polynomially. Suppose the alphabet of X, denoted by $\mathcal{A} = \{x_i : i \in \mathbb{N}\}$, has no accumulation point. Then

$$d_{\min} \triangleq \inf_{i \neq j} |x_i - x_j| > 0. \tag{4.66}$$

If N is almost surely bounded, say $|N| \leq A$, then $\mathsf{mmse}(X, N, \mathsf{snr}) = 0$ for all $\mathsf{snr} > \left(\frac{2A}{d_{\min}}\right)^2$. On the other hand, if X is almost surely bounded, say $|X| \leq \bar{A}$, then $\mathsf{mmse}(X, \mathsf{snr})$ decays exponentially: since the error probability, denoted by $p_{\mathrm{e}}(\mathsf{snr})$, of

a MAP estimator for X based on $\sqrt{\mathsf{snr}}X + N_{\mathsf{G}}$ is $O\left(Q\left(\frac{\sqrt{\mathsf{snr}}}{2}d_{\min}\right)\right)$, where $Q(t) = \int_t^\infty \varphi(x)\mathrm{d}x$ and φ is the standard normal density

$$\varphi(x) = \frac{1}{\sqrt{2\pi}}\mathrm{e}^{-\frac{x^2}{2}}. \tag{4.67}$$

Hence

$$\mathsf{mmse}(\mathsf{snr}) \leq \bar{A}^2 p_{\mathrm{e}}(\mathsf{snr}) = O\left(\frac{1}{\sqrt{\mathsf{snr}}}\mathrm{e}^{-\frac{\mathsf{snr}}{8}d_{\min}^2}\right). \tag{4.68}$$

If the input alphabet has accumulation points, it is possible that the MMSE decays polynomially. For example, when X takes values on $\left\{0, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots\right\}$ and N is uniform distributed on $[0, 1]$, it can be shown that $\mathsf{mmse}(X, N, \mathsf{snr}) = \Theta(\mathsf{snr}^{-2})$.

4.5.3 Absolutely continuous input

Theorem 23. *Suppose the density of $N \in L^2(\Omega)$ is bounded and is such that for some $\alpha > 1$,*

$$f_N(u) = O\left(|u|^{-\alpha}\right), \tag{4.69}$$

as $|u| \to \infty$. Then

$$\mathscr{D}(X, N) = 1 \tag{4.70}$$

holds for all X with an absolutely continuous distribution with respect to Lebesgue measure. In particular,

$$\mathscr{D}(X) = 1, \tag{4.71}$$

i.e.,

$$\mathsf{mmse}(X, \mathsf{snr}) = \frac{1}{\mathsf{snr}} + o\left(\frac{1}{\mathsf{snr}}\right). \tag{4.72}$$

Proof. Appendix B.5. □

If the density of X is sufficiently regular, then (4.70) holds for all noise distributions:

Theorem 24. *Suppose the density of X is continuous and bounded. Then*

$$\mathscr{D}(X, N) = 1 \tag{4.73}$$

holds for all (not necessarily absolutely continuous) $N \in L^2(\Omega)$.

Proof. Appendix B.5. □

In view of Theorem 24 and (4.30), we conclude that the linear estimator is *dimensionally optimal* for estimating absolutely continuous random variables contaminated by additive Gaussian noise, in the sense that it achieves the input MMSE dimension.

A refinement of Theorem 24 entails the second-order expansion for $\mathsf{mmse}(X, \mathsf{snr})$ for absolutely continuous input X. This involves the Fisher information of X. Suppose $J(X) < \infty$. Then

$$J(\sqrt{\mathsf{snr}}X + N_{\mathsf{G}}) \leq \int \varphi(z) J(\sqrt{\mathsf{snr}}X + z) \mathrm{d}z \tag{4.74}$$

$$= \int \varphi(z) J(\sqrt{\mathsf{snr}}X) \mathrm{d}z \tag{4.75}$$

$$= \frac{J(X)}{\mathsf{snr}} < \infty, \tag{4.76}$$

where (4.74) and (4.75) follow from the convexity and translation invariance of Fisher information respectively. In view of (4.35), we have

$$\mathsf{mmse}(X, \mathsf{snr}) = \frac{1}{\mathsf{snr}} + O\left(\frac{1}{\mathsf{snr}^2}\right), \tag{4.77}$$

which improves (4.72) slightly.

Under stronger regularity conditions the second-order term can be determined exactly. A result of Prelov and van der Meulen [75] states the following: if $J(X) < \infty$ and the density of X satisfies certain regularity conditions [75, (3) – (7)], then

$$J(X + \epsilon N_{\mathsf{G}}) = J(X) + O(\epsilon). \tag{4.78}$$

Therefore in view of (4.34), we have

$$\mathsf{mmse}(X, \mathsf{snr}) = \frac{1}{\mathsf{snr}} - \frac{J(X)}{\mathsf{snr}^2} + o\left(\frac{1}{\mathsf{snr}^2}\right). \tag{4.79}$$

This result can be understood as follows: Stam's inequality [78] implies that

$$\frac{1}{J(\sqrt{\mathsf{snr}}X + N_{\mathsf{G}})} \geq \frac{1}{J(\sqrt{\mathsf{snr}}X)} + \frac{1}{J(N_{\mathsf{G}})} \tag{4.80}$$

$$= \frac{\mathsf{snr}}{J(X)} + 1. \tag{4.81}$$

Using (4.34), we have

$$\mathsf{mmse}(X, \mathsf{snr}) \geq \frac{1}{J(X) + \mathsf{snr}} \tag{4.82}$$

$$= \frac{1}{\mathsf{snr}} - \frac{J(X)}{\mathsf{snr}^2} + o\left(\frac{1}{\mathsf{snr}^2}\right). \tag{4.83}$$

Inequality (4.82) is also known as the Bayesian Crámer-Rao bound (or the Van Trees inequality) [79, pp. 72–73], [80, Corollary 2.3]. In view of (4.79), we see that (4.82) is asymptotically tight for sufficiently regular densities of X.

Instead of using the asymptotic expansion of Fisher information and Stam's inequality, we can show that (4.79) holds for a much broader class of densities of X and non-Gaussian noise:

Theorem 25. *Suppose $X \in L^3(\Omega)$ with bounded density f_X whose first two derivatives are also bounded and $J(X) < \infty$. Furthermore, assume f_X satisfies the following regularity condition [81, Equation (2.5.16)]:*

$$\int_{\mathbb{R}} f_X''(x)\mathrm{d}x = 0. \tag{4.84}$$

Then for any $N \in L^2(\Omega)$,

$$\mathsf{mmse}(X, N, \mathsf{snr}) = \frac{\mathsf{var}N}{\mathsf{snr}} - J(X)\left(\frac{\mathsf{var}N}{\mathsf{snr}}\right)^2 + o\left(\frac{1}{\mathsf{snr}^2}\right). \tag{4.85}$$

Proof. Appendix B.5. □

Combined with (4.83), Theorem 25 implies that the Bayesian Crámer-Rao bound

$$\mathsf{mmse}(X, N, \mathsf{snr}) \geq \frac{1}{J(X) + \mathsf{snr}J(N)} \tag{4.86}$$

$$= \frac{1}{\mathsf{snr}\, J(N)} - \frac{J(X)}{\mathsf{snr}^2\, J^2(N)} + o\left(\frac{1}{\mathsf{snr}^2}\right) \tag{4.87}$$

is asymptotically tight *if and only if* the noise is Gaussian.

Equation (4.85) reveals a new *operational role* of $J(X)$. The regularity conditions imposed in Theorem 25 are much weaker and easier to check than those in [75]; however, (4.79) is slightly stronger than (4.85) because the $o(\mathsf{snr}^{-2})$ term in (4.79) is in fact $O(\mathsf{snr}^{-\frac{5}{2}})$, as a result of (4.78).

It is interesting to compare (4.85) to results on asymptotic Bayesian risk in the large sample limit [70, Theorem 5.1a]:

- When N is Gaussian, recalling (4.22), we have

$$\mathsf{mmse}(X|Y^n) = \mathsf{mmse}(X, n) \tag{4.88}$$

$$= \frac{1}{n} - \frac{J(X)}{n^2} + o\left(\frac{1}{n^2}\right). \tag{4.89}$$

  This agrees with the results in [70] but our proof requires less stringent conditions on the density of X.

- When N is non-Gaussian, under the regularity conditions in [70, Theorem 5.1a], we have

$$\mathsf{mmse}(X|Y^n) = \frac{1}{n\, J(N)} - \frac{J(X)}{n^2\, J^2(N)} + o\left(\frac{1}{n^2}\right), \tag{4.90}$$

which agrees with the Bayesian Crámer-Rao bound in the product case. On the other hand, by (4.85) we have

$$\mathsf{mmse}(X|\bar{Y}) = \mathsf{mmse}(X, n) \tag{4.91}$$

$$= \frac{\mathsf{var}N}{n} - J(X)\left(\frac{\mathsf{var}N}{n}\right)^2 + o\left(\frac{1}{n^2}\right). \tag{4.92}$$

whose first-order term is inferior to (4.90), due to the Crámer-Rao bound $\mathsf{var}N \geq \frac{1}{J(N)}$. This agrees with the fact that the sample mean is asymptotically suboptimal for non-Gaussian noise, and the suboptimality is characterized by the gap in the Crámer-Rao inequality.

To conclude the discussion of absolutely continuous inputs, we give an example where (4.85) fails:

Example 2. Consider X and N uniformly distributed on $[0, 1]$. It is shown in Appendix B.2 that

$$\mathsf{mmse}(X, N, \mathsf{snr}) = \mathsf{var}N\left(\frac{1}{\mathsf{snr}} - \frac{1}{2\mathsf{snr}^{\frac{3}{2}}}\right), \quad \mathsf{snr} \geq 4. \tag{4.93}$$

Note that (4.85) does not hold because X does not have a differentiable density, hence $J(X) = \infty$. This example illustrates that the $o\left(\frac{1}{\mathsf{snr}}\right)$ term in (4.72) is not necessarily $O\left(\frac{1}{\mathsf{snr}^2}\right)$.

4.5.4 Mixed distribution

Next we present results for general mixed distributions, which are direct consequences of Theorem 21. The following result asserts that MMSE dimensions are affine functionals, a fundamental property shared by Rényi information dimension Theorem 5.

Theorem 26.

$$\overline{\mathscr{D}}\left(\sum \alpha_i \mu_i\right) = \sum \alpha_i \overline{\mathscr{D}}(\mu_i), \tag{4.94}$$

$$\underline{\mathscr{D}}\left(\sum \alpha_i \mu_i\right) = \sum \alpha_i \underline{\mathscr{D}}(\mu_i). \tag{4.95}$$

where $\{\alpha_i : i \in \mathbb{N}\}$ is a probability mass function and each μ_i is a probability measure.

Another application of Theorem 21 is to determine the MMSE dimension of inputs with mixed distributions, which are frequently used in statistical models of *sparse signals* [5, 4, 7, 82]. According to Theorem 21, knowing the support does not decrease the MMSE dimension.

Corollary 1. *Let $X = UZ$ where U is independent of Z, taking values in $\{0,1\}$ with $\mathbb{P}\{U=1\} = \rho$. Then*

$$\overline{\mathscr{D}}(X) = \rho\overline{\mathscr{D}}(Z), \tag{4.96}$$

$$\underline{\mathscr{D}}(X) = \rho\underline{\mathscr{D}}(Z). \tag{4.97}$$

Obtained by combining Theorems 22, 23 and 26, the next result solves the MMSE dimension of discrete-continuous mixtures completely. Together with Theorem 20, it also provides an alternative proof of Rényi's theorem on information dimension for mixtures (Theorem 4) via MMSE.

Theorem 27. *Let X have a discrete-continuous mixed distribution as defined in (1.4). Then its MMSE dimension equals the weight of the absolutely continuous part, i.e.,*

$$\mathscr{D}(X) = \rho. \tag{4.98}$$

The above results also extend to non-Gaussian noise or general channels which satisfy the condition in Remark 33 at the end of Appendix B.3.

To conclude this subsection, we illustrate Theorem 27 by the following examples:

Example 3 (Continuous input)**.** If $X \sim \mathcal{N}(0,1)$, then

$$\mathsf{mmse}(X, \mathsf{snr}) = \frac{1}{\mathsf{snr}+1} \tag{4.99}$$

and (4.71) holds.

Example 4 (Discrete input)**.** If X is equiprobable on $\{-1,1\}$, then by (4.68) or [83, Theorem 3],

$$\mathsf{mmse}(X, \mathsf{snr}) = O\left(\frac{1}{\sqrt{\mathsf{snr}}}\mathrm{e}^{-\frac{\mathsf{snr}}{2}}\right) \tag{4.100}$$

and (4.65) holds.

Example 5 (Mixed input)**.** Let N be uniformly distributed in $[0,1]$, and let X be distributed according to an equal mixture of a mass at 0 and a uniform distribution on $[0,1]$. Then

$$\frac{\mathsf{mmse}(X,N,\mathsf{snr})}{\mathsf{var}N}$$

$$= \frac{3}{2} - \frac{3}{4\sqrt{\mathsf{snr}}} + \frac{1}{\mathsf{snr}} - \frac{1}{4\mathsf{snr}^{\frac{3}{2}}} - \frac{3\sqrt{\mathsf{snr}}}{2}\log\frac{1+\sqrt{\mathsf{snr}}}{\sqrt{\mathsf{snr}}} \tag{4.101}$$

$$= \frac{1}{2\mathsf{snr}} + \frac{1}{4\mathsf{snr}^{\frac{3}{2}}} + o\left(\frac{1}{\mathsf{snr}^{\frac{3}{2}}}\right), \tag{4.102}$$

which implies $\mathscr{D}(X,N) = \frac{1}{2}$ and verifies Theorem 27 for non-Gaussian noise.

4.5.5 Singularly continuous distribution

In this subsection we consider atomless input distributions mutually singular with respect to Lebesgue measure. There are two new phenomena regarding MMSE dimensions of singular inputs. First, the lower and upper MMSE dimensions $\underline{\mathscr{D}}(X,N)$ and $\overline{\mathscr{D}}(X,N)$ depend on the noise distribution, even if the noise is restricted to be absolutely continuous. Second, the MMSE dimension of a singular input need not exist. For an important class of *self-similar* singular distributions (e.g., the Cantor distribution), the function $\mathsf{snr} \mapsto \mathsf{snr}\,\mathsf{mmse}(X,\mathsf{snr})$ oscillates between the lower and upper dimension periodically in $\log \mathsf{snr}$ (i.e., in dB). This periodicity arises from the *self-similarity* of the input, and the period can be determined exactly. Unlike the lower and upper dimension, the period does not depend on the noise distribution.

We focus on a special class of inputs with self-similar distributions [46, p.36]: inputs with i.i.d. digits. Consider $X \in [0,1]$ a.s. For an integer $M \geq 2$, the M-ary expansion of X is defined in (2.47). Assume that the sequence $\{(X)_j : j \in \mathbb{N}\}$ is i.i.d. with common distribution P supported on $\{0,\ldots,M-1\}$. According to (2.50), the information dimension of X exists and is given by the normalized entropy rate of the digits:

$$d(X) = \frac{H(P)}{\log M}. \tag{4.103}$$

For example, if X is Cantor-distributed, then the ternary expansion of X consists of i.i.d. digits, and for each j,

$$\mathbb{P}\{(X)_j = 0\} = \mathbb{P}\{(X)_j = 2\} = 1/2. \tag{4.104}$$

By (4.103), the information dimension of the Cantor distribution is $\log_3 2$. The next theorem shows that for such X the scaling constant of MMSE oscillates periodically.

Theorem 28. *Let $X \in [0,1]$ a.s., whose M-ary expansion defined in (2.47) consists of i.i.d. digits with common distribution P. Then for any $N \in L^2(\Omega)$, there exists a $2\log M$-periodic function*[4] *$\Phi_{X,N} : \mathbb{R} \to [0,1]$, such that as $\mathsf{snr} \to \infty$,*

$$\mathsf{mmse}(X,N,\mathsf{snr}) = \frac{\mathsf{var}N}{\mathsf{snr}}\Phi_{X,N}(\log \mathsf{snr}) + o\left(\frac{1}{\mathsf{snr}}\right). \tag{4.105}$$

The lower and upper MMSE dimension of (X,N) are given by

$$\overline{\mathscr{D}}(X,N) = \limsup_{b\to\infty} \Phi_{X,N}(b) = \sup_{0\leq b\leq 2\log M} \Phi_{X,N}(b), \tag{4.106}$$

$$\underline{\mathscr{D}}(X,N) = \liminf_{b\to\infty} \Phi_{X,N}(b) = \inf_{0\leq b\leq 2\log M} \Phi_{X,N}(b). \tag{4.107}$$

[4] Let $T > 0$. We say a function $f : \mathbb{R} \to \mathbb{R}$ is T-periodic if $f(x) = f(x+T)$ for all $x \in \mathbb{R}$, and T is called a period of f ([84, p. 183] or [85, Section 3.7]). This differs from the definition of the *least period* which is the infimum of all periods of f.

Moreover, when $N = N_{\mathsf{G}}$ is Gaussian, the average of $\Phi_{X,N_{\mathsf{G}}}$ over one period coincides with the information dimension of X:

$$\frac{1}{2\log M}\int_0^{2\log M}\Phi_{X,N_{\mathsf{G}}}(b)\,\mathrm{d}b = d(X) = \frac{H(P)}{\log M}. \tag{4.108}$$

Proof. Appendix B.6. □

Remark 7. Trivial examples of Theorem 28 include $\Phi_{X,N} \equiv 0$ ($X = 0$ or 1 a.s.) and $\Phi_{X,N} \equiv 1$ (X is uniformly distributed on $[0,1]$).

Theorem 28 shows that in the high-SNR regime, the function $\mathsf{snr}\cdot\mathsf{mmse}(X,N,\mathsf{snr})$ is periodic in snr (dB) with period $20\log M$ dB. Plots are given in Fig. 4.1 – 4.2. Although this reveals the oscillatory nature of $\mathsf{mmse}(X,N,\mathsf{snr})$, we do not have a general formula to compute the lower (or upper) MMSE dimension of (X,N). However, when the noise is Gaussian, Theorem 20 provides a sandwich bound in terms of the information dimension of X, which is reconfirmed by combining (4.106) – (4.108).

Remark 8 (Binary-valued noise)**.** One interesting case for which we are able to compute the lower MMSE dimension corresponds to binary-valued noise, with which all singular inputs (including discrete distributions) have *zero* lower MMSE dimension (see Appendix B.7 for a proof). This phenomenon can be explained by the following fact about negligible sets: a set with zero Lebesgue measure can be translated by an arbitrarily small amount to be disjoint from itself. Therefore, if an input is supported on a set with zero Lebesgue measure, we can perform a binary hypothesis test based on its noisy version, which admits a decision rule with zero error probability when SNR is large enough.

4.6 Numerical Results

4.6.1 Approximation by discrete inputs

Due to the difficulty of computing conditional expectation and estimation error in closed form, we capitalize on the regularity of the MMSE functional by computing the MMSE of successively finer discretizations of a given X. For an integer m we uniformly quantize X to $[X]_m$. Then we numerically compute $\mathsf{mmse}([X]_m,\mathsf{snr})$ for fixed snr. By the weak continuity of $P_X \mapsto \mathsf{mmse}(X,\mathsf{snr})$ [86, Corollary 3], as the quantization level m grows, $\mathsf{mmse}([X]_m,\mathsf{snr})$ converges to $\mathsf{mmse}(X,\mathsf{snr})$; however, one caveat is that to obtain the value of MMSE within a given accuracy, the quantization level needs to grow as snr grows (roughly as $\log\mathsf{snr}$) so that the quantization error is much smaller than the noise. Lastly, due to the following result, to obtain the MMSE dimension it is sufficient to only consider integer values of snr.

Lemma 13. *It is sufficient to restrict to integer values of* snr *when calculating* $\overline{\mathscr{D}}(X|U)$ *and* $\underline{\mathscr{D}}(X|U)$ *in (4.10) and (4.11) respectively.*

Proof. For any $\mathsf{snr} > 0$, there exists $n \in \mathbb{Z}_+$, such that $n \leq \mathsf{snr} < n+1$. Since the function $\mathsf{mmse}(\mathsf{snr}) = \mathsf{mmse}(X, \mathsf{snr}|U)$ is monotonically decreasing, we have

$$\mathsf{mmse}(n+1) \leq \mathsf{mmse}(\mathsf{snr}) \leq \mathsf{mmse}(n), \tag{4.109}$$

hence

$$n\ \mathsf{mmse}(n+1) \leq \mathsf{snr}\ \mathsf{mmse}(\mathsf{snr}) \tag{4.110}$$

$$\leq n\ \mathsf{mmse}(n) + \mathsf{mmse}(n). \tag{4.111}$$

Using that $\lim_{n\to\infty} \mathsf{mmse}(n) = 0$, the claim of Lemma 13 follows. □

4.6.2 Self-similar input distribution

We numerically calculate the MMSE for Cantor distributed X and Gaussian noise by choosing $[X]_m = \lfloor 3^m X \rfloor 3^{-m}$. By (4.104), $[X]_m$ is equiprobable on the set $\mathcal{A}_m$, which has cardinality 2^m and consists of all 3-adic fractionals whose ternary digits are either 0 or 2. According to Theorem 28, in the high-SNR regime, $\mathsf{snr}\,\mathsf{mmse}(X, \mathsf{snr})$ oscillates periodically in $\log \mathsf{snr}$ with period $2\log 3$, as plotted in Fig. 4.1. The lower and upper MMSE dimensions of the Cantor distribution turn out to be (to six decimals):

$$\underline{\mathscr{D}}(X) = 0.621102, \tag{4.112}$$

$$\overline{\mathscr{D}}(X) = 0.640861. \tag{4.113}$$

Note that the information dimension $d(X) = \log_3 2 = 0.630930$ is sandwiched between $\underline{\mathscr{D}}(X)$ and $\overline{\mathscr{D}}(X)$, according to Theorem 20. From this and other numerical evidence it is tempting to conjecture that

$$d(X) = \frac{\underline{\mathscr{D}}(X) + \overline{\mathscr{D}}(X)}{2} \tag{4.114}$$

when the noise is Gaussian.

It should be pointed out that the sandwich bounds in (4.46) need not hold when N is not Gaussian. For example, in Fig. 4.2 where $\mathsf{snr} \cdot \mathsf{mmse}(X, N, \mathsf{snr})$ is plotted against $\log_3 \mathsf{snr}$ for X Cantor distributed and N uniformly distributed in $[0, 1]$, it is evident that $d(X) = \log_3 2 > \overline{\mathscr{D}}(X, N) = 0.5741$.

4.6.3 Non-Gaussian noise

Via the input-noise duality in (4.44), studying high-SNR asymptotics provides insights into the behavior of $\mathsf{mmse}(X, N, \mathsf{snr})$ for non-Gaussian noise N. Combining various results from Sections 4.5.2, 4.5.3 and 4.5.5, we observe that $\mathsf{mmse}(X, N, \mathsf{snr})$ can behave very irregularly in general, unlike in Gaussian channels where $\mathsf{mmse}(X, \mathsf{snr})$ is strictly decreasing in snr. To illustrate this, we consider the case where standard Gaussian input is contaminated by various additive noises. For all N, it is evident

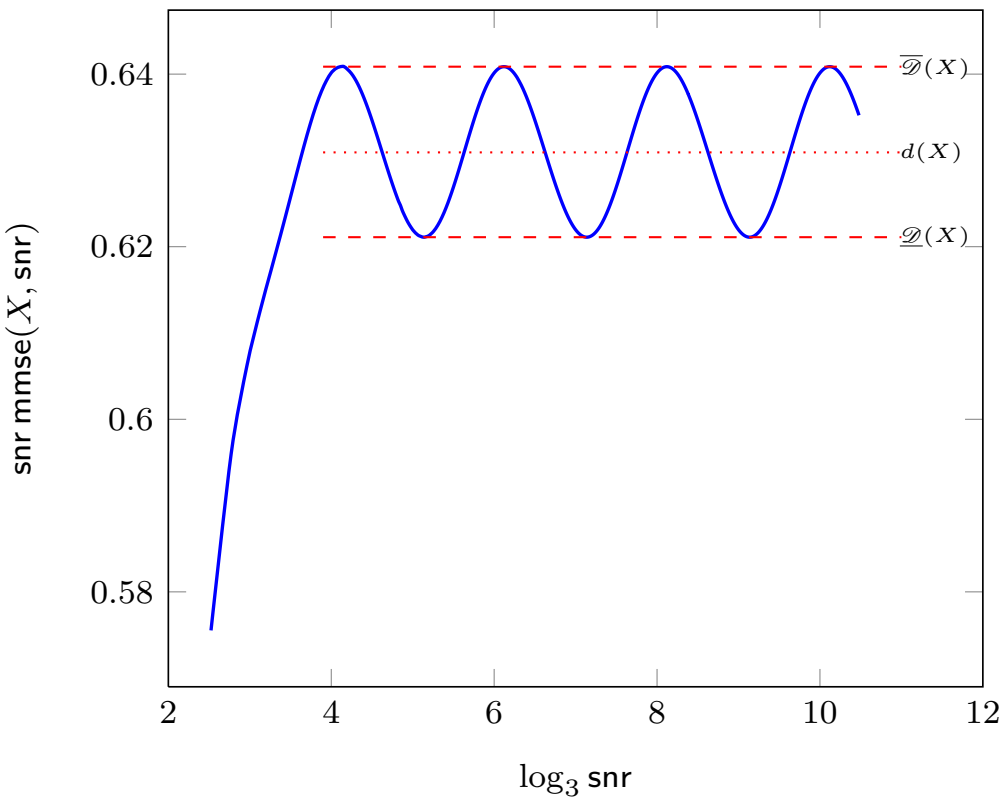

Figure 4.1: Plot of $\mathsf{snr}\,\mathsf{mmse}(X,\mathsf{snr})$ against $\log_3 \mathsf{snr}$, where X has a ternary Cantor distribution.

that $\mathsf{mmse}(X,N,0) = \mathsf{var}X = 1$. Due to Theorem 24, the MMSE vanishes as $\frac{\mathsf{var}N}{\mathsf{snr}}$ regardless of the noise. The behavior of MMSE associated with Gaussian, Bernoulli and Cantor distributed noises is (Fig. 4.3a):

- For standard Gaussian N, $\mathsf{mmse}(X,\mathsf{snr}) = \frac{1}{1+\mathsf{snr}}$ is *continuous* at $\mathsf{snr} = 0$ and decreases monotonically according to $\frac{1}{\mathsf{snr}}$. Recall that [86, Section III] this monotonicity is due to the MMSE data-processing inequality [87] and the stability of Gaussian distribution.
- For equiprobable Bernoulli N, $\mathsf{mmse}(X,N,\mathsf{snr})$ has a *discontinuity of first kind* at $\mathsf{snr} = 0$.

$$\mathsf{var}X = \mathsf{mmse}(X,N,0) > \lim_{\mathsf{snr}\downarrow 0} \mathsf{mmse}(X,N,\mathsf{snr}) = 0. \tag{4.115}$$

  As $\mathsf{snr} \to 0$, the MMSE vanishes according to $O\left(\frac{1}{\sqrt{\mathsf{snr}}}\mathrm{e}^{-\frac{1}{2\mathsf{snr}}}\right)$, in view of (4.44) and (4.68), and since it also vanishes as $\mathsf{snr} \to \infty$, it is not monotonic with $\mathsf{snr} > 0$.
- For Cantor distributed N, $\mathsf{mmse}(X,N,\mathsf{snr})$ has a *discontinuity of second kind* at $\mathsf{snr} = 0$. According to Theorem 28, as $\mathsf{snr} \to 0$, MMSE oscillates relentlessly and does not have a limit (See the zoom-in plot in Fig. 4.3b).

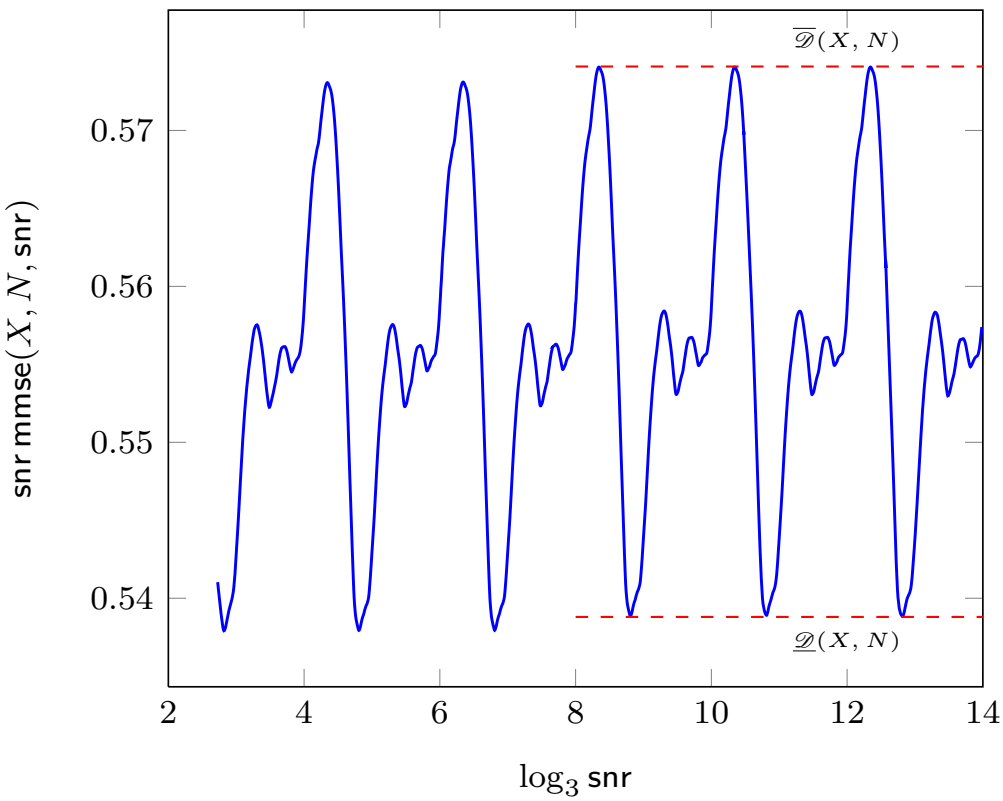

Figure 4.2: Plot of $\mathsf{snr} \cdot \mathsf{mmse}(X, N, \mathsf{snr})/\mathsf{var}N$ against $\log_3 \mathsf{snr}$, where X has a ternary Cantor distribution and N is uniformly distributed in $[0, 1]$.

4.7 Discussions

Through the high-SNR asymptotics of MMSE in Gaussian channels, we defined a new information measure called the MMSE dimension. Although stemming from estimation-theoretic principles, MMSE dimension shares several important features with Rényi's information dimension. By Theorem 21 and Theorem 5, they are both affine functionals. According to Theorem 20, information dimension is sandwiched between the lower and upper MMSE dimensions. For distributions with no singular components, they coincide to be the weight of the continuous part of the distribution. The high-SNR scaling law of mutual information and MMSE in Gaussian channels are governed by the information dimension and the MMSE dimension respectively. In Chapter 5, we show that the information dimension plays a pivotal role in almost lossless analog compression, an information-theoretic model for noiseless compressed sensing. In fact we show in Chapter 6 that the MMSE dimension is closely related to the fundamental limit of noisy compressed sensing and the optimal phase transition of noise sensitivity (see Theorem 48).

Characterizing the high-SNR suboptimality of linear estimation, (4.30) provides an alternative definition of MMSE dimension, which enables us to extend our results to random vectors or processes. In these more general setups, it is interesting to investigate how the *causality* constraint of the estimator affects the high-SNR behavior

of the optimal estimation error. Another direction of generalization is to study the high-SNR asymptotics of MMSE with a *mismatched* model in the setup of [88] or [89].

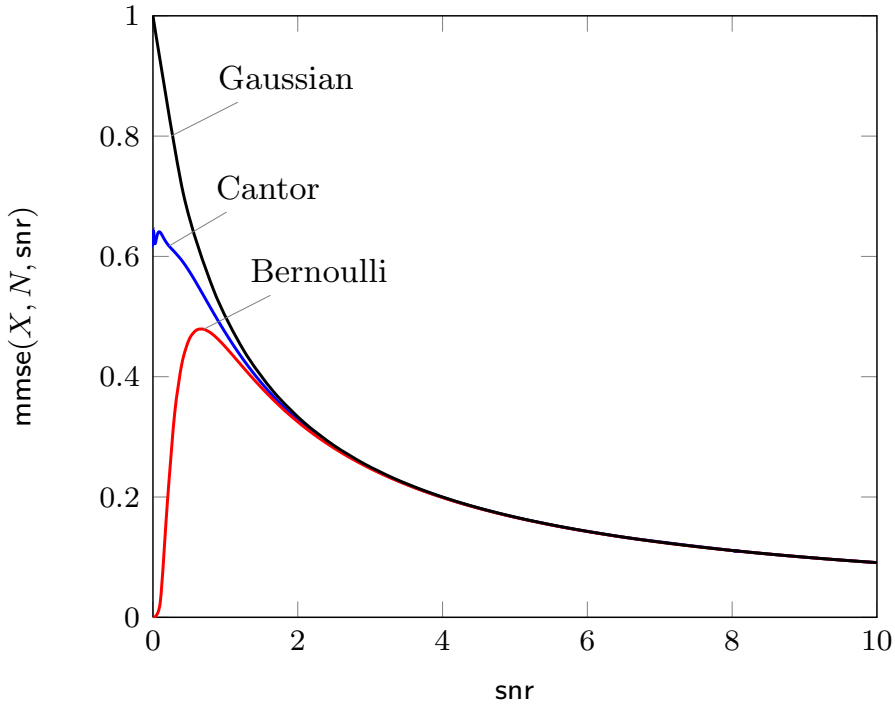

(a) Behavior of MMSE under various noise distributions.

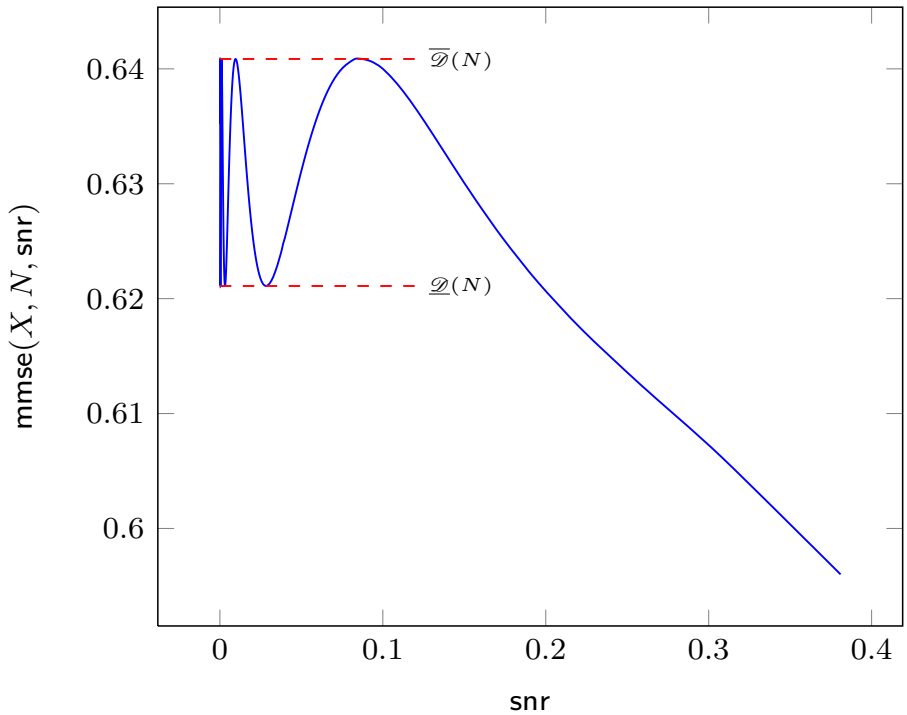

(b) Low-SNR plot for Cantor distributed N.

Figure 4.3: Plot of $\mathsf{mmse}(X, N, \mathsf{snr})$ against snr, where X is standard Gaussian and N is standard Gaussian, equiprobable Bernoulli or Cantor distributed (normalized to have unit variance).

Chapter 5

Noiseless compressed sensing

In this chapter we present the theory of almost lossless analog compression, a Shannon theoretic framework for noiseless compressed sensing. It also extends the classical almost lossless source coding theory to uncountable source and code alphabets. Section 5.2 introduces a unified framework for lossless analog compression. We prove coding theorems under various regularity conditions on the compressor or decompressor, which, in the memoryless case, involves Rényi information dimension of the source. These results provide new operational characterization of the information dimension in Shannon theory. Proofs are deferred till Sections 5.3 – 5.6. Some technical lemmas are proved in Appendix C. The material in this chapter has been presented in part in [27, 5].

5.1 Motivations

5.1.1 Noiseless compressed sensing versus data compression

The "bit" is the universal currency in lossless source coding theory [90], where Shannon entropy is the fundamental limit of compression rate for discrete memoryless sources. Sources are modeled by stochastic processes and redundancy is exploited as probability is concentrated on a set of exponentially small cardinality as blocklength grows. Therefore by encoding this subset, data compression is achieved if we tolerate a positive, albeit arbitrarily small, block error probability.

Compressed sensing, on the other hand, employs linear transformations to encode real vectors by real-valued measurements. The formulation of noiseless compressed sensing is reminiscent of the traditional lossless data compression in the following sense:

- Sources are sparse in the sense that each vector is supported on a set much smaller than the blocklength. This type of redundancy in terms of sparsity is exploited to achieve effective compression by taking fewer number of linear measurements.
- In contrast to lossy data compression, block error probability, instead of distortion, is the performance benchmark.

- The central problem is to determine how many compressed measurements are sufficient/necessary for recovery with vanishing block error probability as block-length tends to infinity.
- Random coding is employed to show the existence of "good" linear encoders. For instance, when the random sensing matrices are drawn from certain ensembles (*e.g.* standard Gaussian), the restricted isometry property (RIP) is satisfied with overwhelming probability and guarantees exact recovery.

On the other hand, there are also significantly different ingredients in compressed sensing in comparison with information theoretic setups:

- Sources are not modeled probabilistically, and the fundamental limits are on a worst-case basis rather than on average. Moreover, block error probability is with respect to the distribution of the encoding random matrices.
- Real-valued sparse vectors are encoded by real numbers instead of bits.
- The encoder is confined to be linear while generally in information-theoretical problems such as lossless source coding we have the freedom to choose the best possible coding scheme.

Departing from the conventional compressed sensing setup, in this chapter we study fundamental limits of lossless source coding for real-valued memoryless sources within an information theoretic setup:

- Sources are modeled by random processes. As explained in Section 1.3, this method is more flexible to describe source redundancy which encompasses, but is not limited to, sparsity. For example, a mixed discrete-continuous distribution of the form (1.2) is suitable for characterizing linearly sparse vectors.
- Block error probability is evaluated by averaging with respect to the source.
- While linear compression plays an important role in our development, our treatment encompasses weaker regularity conditions.

5.1.2 Regularities of coding schemes

Discrete sources have been the sole object in lossless data compression theory. The reason is at least twofold. First of all, non-discrete sources have infinite entropy, which implies that representation with arbitrarily small block error probability requires arbitrarily large rate. On the other hand, even if we consider encoding analog sources by real numbers, the result is still trivial, as $\mathbb{R}$ and $\mathbb{R}^n$ have the same cardinality. Therefore a single real number is capable of representing a real vector losslessly, yielding a universal compression scheme for any analog source with *zero* rate and *zero* error probability.

However, it is worth pointing out that the compression method proposed above is *not* robust because the bijection between $\mathbb{R}$ and $\mathbb{R}^n$ is highly irregular. In fact, neither the encoder nor the decoder can be *continuous* [91, Exercise 6(c), p. 385]. Therefore, using such a compression scheme, a slight change of source realization could result in a drastic change of encoder output and the decoder could suffer from a large distortion due to a tiny perturbation of the codeword. This disadvantage motivates us to study how to compress not only *losslessly* but also *gracefully* on Euclidean spaces. In fact

some authors have also noticed the importance of regularity in data compression. In [92] Montanari and Mossel observed that the optimal data compression scheme often exhibits the following inconvenience: codewords tend to depend chaotically on the data; hence changing a single source symbol leads to a radical change in the codeword. In [92], a source code is said to be *smooth* (resp. *robust*) if the encoder (resp. decoder) is Lipschitz (see Definition 10) with respect to the Hamming distance. The fundamental limits of smooth lossless compression are analyzed in [92] for binary sources via sparse graph codes. In this thesis we focus on real-valued sources with general distributions. Introducing a *topological structure* makes the nature of the problem quite different from traditional formulations in the discrete world, and calls for machinery from dimension theory and geometric measure theory.

The regularity constraints on the coding procedures considered in this chapter are the linearity of compressors and the robustness of decompressors, the latter of which is motivated by the great importance of robust reconstruction in compressed sensing (e.g., [93, 94, 95, 96]), as *noise resilience* is an indispensable property for decompressing sparse signals from real-valued measurements. More specifically, the following two robustness constraints are considered:

Definition 9 ((L,Δ)-stable)**.** Let (U, d_U) and (V, d_V) be metric spaces and $T \subset U$. $g : U \to V$ is called (L,Δ)-stable on T if for all $x, y \in T$

$$d_U(x,y) \leq \Delta \Rightarrow d_V(g(x), g(y)) \leq L\Delta. \tag{5.1}$$

And we say g is Δ-stable if g is $(1,\Delta)$-stable.

Definition 10 (Hölder and Lipschitz continuity)**.** Let (U, d_U) and (V, d_V) be metric spaces. A function $g : U \to V$ is called (L,γ)-*Hölder continuous* if there exists $L, \gamma \geq 0$ such that for any $x, y \in U$,

$$d_V(g(x), g(y)) \leq L d_U(x,y)^\gamma. \tag{5.2}$$

g is called L-*Lipschitz* if g is $(L,1)$-Hölder continuous. g is simply called *Lipschitz* (resp. β-*Hölder continuous*) if g is L-Lipschitz (resp. (L,β)-Hölder continuous) for some $L \geq 0$. The *Lipschitz constant* of g is defined as

$$\mathrm{Lip}(g) \triangleq \sup_{x \neq y} \frac{d_V(g(x), g(y))}{d_U(x,y)}. \tag{5.3}$$

Remark 9. The Δ-stability is a much weaker notion than Lipschitz continuity, because it is easy to see that g is L-Lipschitz if and only if g is (L,Δ)-stable for every $\Delta > 0$.

To appreciate the relevance of Definitions 9 and 10 to compressed sensing, let us consider the following robustness result [96, Theorem 3.2]:

Theorem 29. *Let $y = \mathbf{A}x_0 + e$, where $\mathbf{A} \in \mathbb{R}^{k\times n}$, $x_0 \in \mathbb{R}^n$ and $e, y \in \mathbb{R}^k$ with $\|x_0\|_0 = S$ and $\|e\|_2 \leq \Delta \in \mathbb{R}_+$. Suppose $\mathbf{A}$ satisfies $\delta_S \leq 0.307$, where δ_S is the*

S-restricted isometry constant of matrix $\mathbf{A}$, defined as the smallest $\delta > 0$ such that

$$(1-\delta)\|u\|_2^2 \leq \|\mathbf{A}u\|_2^2 \leq (1+\delta)\|u\|_2^2 \tag{5.4}$$

for all $u \in \mathbb{R}^n$ such that $\|u\|_0 \leq S$. Then the following ℓ_2-constrained ℓ_1-minimization problem

$$\begin{aligned} \min \quad & \|x\|_1 \\ \text{s.t.} \quad & \|\mathbf{A}x - y\|_2 \leq \Delta. \end{aligned} \tag{5.5}$$

has a unique solution, denoted by $\hat{x}_\Delta$, which satisfies

$$\|\hat{x}_\Delta - x_0\|_2 \leq \frac{\Delta}{0.307 - \delta_S}. \tag{5.6}$$

In particular, when $\Delta = 0$, x_0 is the unique solution to the following ℓ_1-minimization problem

$$\begin{aligned} \min \quad & \|x\|_1 \\ \text{s.t.} \quad & \mathbf{A}x = \mathbf{A}x_0. \end{aligned} \tag{5.7}$$

Theorem 29 states that if the sensing matrix $\mathbf{A}$ has a sufficiently small restricted isometry constant, the constrained ℓ_1-minimization decoder $\hat{x}_\Delta$ (5.5) yields a reconstruction error whose ℓ_2-norm is upper bounded *proportionally* to the ℓ_2-norm of the noise. This is equivalent to the (L, Δ)-stability of the decoder with $L = \frac{1}{0.307-\delta_S}$. The drawback of this result is that the decoder depends on the noise magnitude Δ. To handle the potential case of unbounded noise (e.g., drawn from a Gaussian distribution), it is desirable to have a robust decoder independent of Δ, which, in view of Remark 9, amounts to the Lipschitz continuity of the decoder. In fact, the unconstrained ℓ_1-minimization decoder in (5.7) is Lipschitz continuous. To see this, observe that (5.7) can be equivalently formulated as a linear programming problem. It is known that the optimal solution to any linear programming problem is Lipschitz continuous with respect to the parameters [97]. However, it is likely that the Lipschitz constant blows up as the ambient dimension grows. As we will show later in Theorem 32, for discrete-continuous mixtures that are most relevant for compressed sensing problems, it is possible to construct decoders with *bounded* Lipschitz constants.

5.2 Definitions and Main results

5.2.1 Lossless data compression

Let the source $\{X_i : i \in \mathbb{N}\}$ be a stochastic process on $(\mathcal{X}^{\mathbb{N}}, \mathcal{F}^{\otimes\mathbb{N}})$, with $\mathcal{X}$ denoting the source alphabet and $\mathcal{F}$ a σ-algebra over $\mathcal{X}$. Let $(\mathcal{Y}, \mathcal{G})$ be a measurable space, where $\mathcal{Y}$ is called the *code alphabet*. The main objective of lossless data compression is to find efficient representations for source realizations $x^n \in \mathcal{X}^n$ by $y^k \in \mathcal{Y}^k$.

Definition 11. A (n, k)-code for $\{X_i : i \in \mathbb{N}\}$ over the code space $(\mathcal{Y}, \mathcal{G})$ is a pair of mappings:

1. Encoder: $f_n : \mathcal{X}^n \to \mathcal{Y}^k$ that is measurable relative to $\mathcal{F}^n$ and $\mathcal{G}^k$.

2. Decoder: $g_n : \mathcal{Y}^k \to \mathcal{X}^n$ that is measurable relative to $\mathcal{G}^k$ and $\mathcal{F}^n$.

The block error probability is $\mathbb{P}\{g_n(f_n(X^n)) \neq X^n\}$.

The fundamental limit of lossless source coding is:

Definition 12 (Lossless Data Compression)**.** Let $\{X_i : i \in \mathbb{N}\}$ be a stochastic process on $(\mathcal{X}^{\mathbb{N}}, \mathcal{F}^{\otimes \mathbb{N}})$. Define $r(\epsilon)$ to be the infimum of $r > 0$ such that there exists a sequence of $(n, \lfloor rn \rfloor)$-codes over the code space $(\mathcal{Y}, \mathcal{G})$, such that

$$\mathbb{P}\{g_n(f_n(X^n)) \neq X^n\} \leq \epsilon \tag{5.8}$$

for all sufficiently large n.

According to the classical discrete almost-lossless source coding theorem, if $\mathcal{X}$ is countable and $\mathcal{Y}$ is finite, the minimum achievable rate for any i.i.d. process with distribution P is

$$r(\epsilon) = \begin{cases} \frac{\log|\mathcal{X}|}{\log|\mathcal{Y}|}, & \epsilon = 0 \\ \frac{H(P)}{\log|\mathcal{Y}|}, & 0 < \epsilon < 1, \\ 0, & \epsilon = 1. \end{cases} \tag{5.9}$$

Using codes over an infinite alphabet, any discrete source can be compressed with zero rate and zero block error probability. In other words, if both $\mathcal{X}$ and $\mathcal{Y}$ are countably infinite, then for all $0 \leq \epsilon \leq 1$,

$$r(\epsilon) = 0 \tag{5.10}$$

for any random process.

5.2.2 Lossless analog compression with regularity conditions

In this subsection we consider the problem of encoding analog sources with analog symbols, that is, $(\mathcal{X}, \mathcal{F}) = (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ and $(\mathcal{Y}, \mathcal{G}) = (\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ or $([0,1], \mathcal{B}_{[0,1]})$ if *bounded* encoders are required, where $\mathcal{B}_{[0,1]}$ denotes the Borel σ-algebra. As in the countably infinite case, zero rate is achievable even for zero block error probability, because the cardinality of $\mathbb{R}^n$ is the same for any n [98]. This conclusion holds even if we require the encoder/decoder to be Borel measurable, because according to Kuratowski's theorem [99, Remark (i), p. 451] every uncountable standard Borel space is isomorphic[1] to $([0,1], \mathcal{B}_{[0,1]})$. Therefore a single real number has the capability of encoding a real vector, or even a real sequence, with a coding scheme that is both *universal* and *deterministic*.

However, the rich topological structure of Euclidean spaces enables us to probe the problem further. If we seek the fundamental limits of not only *lossless* coding but *"graceful" lossless* coding, the result is not trivial anymore. In this spirit, our

[1] Two measurable spaces are isomorphic if there exists a measurable bijection whose inverse is also measurable.

various definitions share the basic information-theoretic setup where a random vector is encoded with a function $f_n : \mathbb{R}^n \to \mathbb{R}^{\lfloor Rn \rfloor}$ and decoded with $g_n : \mathbb{R}^{\lfloor Rn \rfloor} \to \mathbb{R}^n$ with $R \leq 1$ such that f_n and g_n satisfy certain regularity conditions and the probability of incorrect reconstruction vanishes as $n \to \infty$.

Regularity conditions of encoders and decoders are imposed for the sake of both less complexity and more robustness. For example, although a surjection g from $[0, 1]$ to $\mathbb{R}^n$ is capable of lossless encoding, its irregularity requires specifying *uncountably* many real numbers to determine this mapping. Moreover, smoothness of the encoding and decoding operation is crucial to guarantee noise resilience of the coding scheme.

Definition 13. Let $\{X_i : i \in \mathbb{N}\}$ be a stochastic process on $(\mathbb{R}^{\mathbb{N}}, \mathcal{B}_{\mathbb{R}}^{\otimes \mathbb{N}})$. Define the minimum ϵ-achievable rate to be the infimum of $R > 0$ such that there exists a sequence of $(n, \lfloor Rn \rfloor)$-codes (f_n, g_n) , such that

$$\mathbb{P}\left\{g_n(f_n(X^n)) \neq X^n\right\} \leq \epsilon \tag{5.11}$$

for all sufficiently large n, and the encoder f_n and decoder g_n are constrained according to Table 5.1. Except for linear encoding where $\mathcal{Y} = \mathbb{R}$, it is assumed that $\mathcal{Y} = [0, 1]$.

Table 5.1: Regularity conditions of encoder/decoders and corresponding minimum ϵ-achievable rates.

Encoder	Decoder	Minimum ϵ-achievable rate
Continuous	Continuous	$\mathsf{R}_0(\epsilon)$
Linear	Borel	$\mathsf{R}^*(\epsilon)$
Borel	Lipschitz	$\mathsf{R}(\epsilon)$
Linear	Lipschitz	$\hat{\mathsf{R}}(\epsilon)$
Borel	Δ-stable	$\overline{\mathsf{R}}(\epsilon, \Delta)$

Remark 10. Let $S_n = \{x^n : g_n(f_n(x^n)) = x^n\}$. Then f_n is the inverse of g_n on S_n. If g_n is L-Lipschitz continuous, f_n satisfies

$$\|f_n(x) - f_n(y)\| \geq \frac{1}{L} \|x - y\| , \tag{5.12}$$

which implies the injectivity of f_n, a necessary condition for decodability. Furthermore, not only does f_n assign different codewords to different source symbols, it also keeps them sufficiently separated proportionally to their distance. In other words, the encoder f_n respects the metric structure of the source alphabet.

Remark 11. In our framework a stable or Lipschitz continuous coding scheme implies *robustness* with respect to noise added at the input of the decompressor, which could result from quantization, finite wordlength or other inaccuracies. For example, suppose that the encoder output $y^k = f_n(x^n)$ is quantized by a q-bit uniform quantizer, resulting in $\tilde{y}^k$. Using a 2^{-q}-stable coding strategy (f_n, g_n), we can decode as

follows: denote the following nonempty set

$$\mathcal{D}(\tilde{y}^k) = \left\{ z^k \in C_n : \left\| z^k - \tilde{y}^k \right\|_\infty \leq 2^{-q} \right\}. \tag{5.13}$$

where $C_n = \{f_n(x^n) : x^n \in \mathbb{R}^n, g_n(f_n(x^n)) = x^n\}$. Pick any z^k in $\mathcal{D}(\tilde{y}^k)$ and output $\hat{x}^n = g_n(z^k)$ as the reconstruction. By the stability of g_n and the triangle inequality, we have

$$\|\hat{x}^n - x^n\|_\infty \leq 2^{-(q-1)}, \tag{5.14}$$

i.e., each component in the decoder output will suffer at most twice the inaccuracy of the decoder input. Similarly, an L-Lipschitz decoder with respect to the ℓ_∞ norm incurs an error no more than $L2^{-q}$.

5.2.3 Coding theorems

The following general result holds for any input process:

Theorem 30. *For any input process and any* $0 \leq \epsilon \leq 1$,

$$\mathsf{R}^*(\epsilon) \leq \mathsf{R}_\mathsf{B}(\epsilon) \leq \mathsf{R}(\epsilon) \leq \hat{\mathsf{R}}(\epsilon), \tag{5.15}$$

where $\mathsf{R}_\mathsf{B}(\epsilon)$ *is defined in (3.15).*

The right inequality in (5.15) follows from the definitions, since we have $\hat{\mathsf{R}}(\epsilon) \geq \max\{\mathsf{R}^*(\epsilon), \mathsf{R}(\epsilon)\}$. The left and middle inequalities, proved in Sections 5.3.3 and 5.4.2 respectively, bridge the lossless Minkowski dimension formulation introduced in Definition 7 with the lossless analog compression framework in Definition 13. Consequently, upper and lower bounds on $\mathsf{R}_\mathsf{B}(\epsilon)$ proved in Section 3.2 provide *achievability* bounds for lossless linear encoding and *converse* bounds for Lipschitz decoding, respectively. Moreover, we have the following surprising result

$$\mathsf{R}^*(\epsilon) \leq \mathsf{R}(\epsilon), \tag{5.16}$$

which states that robust reconstruction is always harder to pursue than linear compression, regardless of the input statistics.

Next we proceed to give results for each of the minimum ϵ-achievable rates introduced in Definition 13 for *memoryless* sources. First we consider the case where the encoder is restricted to be linear.

Theorem 31 (Linear encoding: general achievability)**.** *Suppose that the source is memoryless. Then*

$$\mathsf{R}^*(\epsilon) \leq \hat{d}(X) \tag{5.17}$$

for all $0 < \epsilon < 1$, *where* $\hat{d}(X)$ *is defined in (2.42). Moreover,*

1. *For all linear encoders (except possibly those in a set of zero Lebesgue measure on the space of real matrices), block error probability* ϵ *is achievable.*

2. *The decoder can be chosen to be β-Hölder continuous for all $0 < \beta < 1 - \frac{\hat{d}(X)}{R}$, where $R > \hat{d}(X)$ is the compression rate.*

Proof. See Section 5.3.3. □

For discrete-continuous mixtures, the following theorem states that linear encoders and Lipschitz decoders can be realized *simultaneously* with *bounded* Lipschitz constants. This result also serves as the foundation for our construction in noisy compressed sensing (see Theorem 49).

Theorem 32 (Linear encoding: discrete-continuous mixture)**.** *Suppose that the source is memoryless with a discrete-continuous mixed distribution. Then*

$$\mathsf{R}^*(\epsilon) = d(X) \tag{5.18}$$

for all $0 < \epsilon < 1$. Moreover, if the discrete part has finite entropy, then for any rate $R > d(X)$, the decompressor can be chosen to be Lipschitz continuous with respect to the ℓ_2-norm with an $O(1)$ Lipschitz-constant that only depends on $R - d(X)$. Consequently,

$$\hat{\mathsf{R}}(\epsilon) = \mathsf{R}^*(\epsilon) = d(X) \tag{5.19}$$

Proof. See Section 5.3.3. □

Combining previous theorems yields the following tight result: for discrete-continuous memoryless mixtures of the form in (1.4) such that the discrete component has finite entropy, we have

$$\mathsf{R}^*(\epsilon) = \mathsf{R}(\epsilon) = \mathsf{R}_{\mathsf{B}}(\epsilon) = \hat{\mathsf{R}}(\epsilon) = \gamma \tag{5.20}$$

for all $0 < \epsilon < 1$.

Remark 12. The achievability proof of $O(1)$ Lipschitz constants in Theorem 32 works only with the ℓ_2 norm, because it depends crucially on the inner product structure endowed by the ℓ_2 norm. The reasons is twofold: first, the Lipschitz constant of a linear mapping with respect to the ℓ_2 norm is given by its maximal singular value, whose behavior for random measurement matrices is well studied. Second, Kirszbraun's theorem states that any Lipschitz mapping between two Hilbert spaces can be extended to the whole space with the same Lipschitz constant [100, Theorem 1.31, p. 21]. This result fails for general Banach spaces, in particular, for $\mathbb{R}^n$ equipped with any ℓ_p-norm ($p \neq 2$) [100, p. 20]. Although by equivalence of norms on finite-dimensional spaces, it is always possible to extend to a Lipschitz function with a larger Lipschitz constant, which, however, will possibly blow up as the dimension increases. Nevertheless, (5.25) shows that even if we allow a sequence of decompressors with Lipschitz constants that diverges as $n \to \infty$, the compression rate is still lower bounded by $\overline{d}(X)$.

Remark 13. In the setup of Theorem 32, it is interesting to examine the robustness of the decoder and its dependence on the measurement rate quantitatively. The proof of Theorem 32 (see (5.189)) shows that for the mixed input distribution (1.4) and any $R > \gamma$, the Lipschitz constant of the decoder can be upper bounded by

$$L = \sqrt{R} \exp\left\{ \frac{1}{R-\gamma} \left[(H(P_{\mathrm{d}}) + R - \gamma)(1-\gamma) + \frac{R}{2} h\left(\frac{\gamma}{R}\right) \right] + \frac{1}{2} \right\}, \tag{5.21}$$

where $h(\alpha) \triangleq \alpha \log \frac{1}{\alpha} + (1-\alpha) \log \frac{1}{1-\alpha}$ is the binary entropy function. Note that (5.21) depends on the excessive rate $\delta \triangleq R - \gamma$ as follows: as $\delta \downarrow 0$,

- if $H(P_{\mathrm{d}}) > 0$ (e.g., simple signals (1.3)), then L depends exponentially on $\frac{1}{\delta}$;
- if $H(P_{\mathrm{d}}) = 0$ (e.g., sparse signals (1.2)), then (5.21) simplifies to

$$L = \sqrt{R} \exp\left\{ \frac{R}{2\delta} h\left(\frac{\gamma}{R}\right) + \frac{3}{2} - \gamma \right\}, \tag{5.22}$$

  which behaves as $\Theta\left(\frac{1}{\sqrt{\delta}}\right)$.

Remark 14. For the sparse signal model (1.2), the optimal threshold γ in (5.20) can be achieved by the ℓ_0-minimization decoder that seeks the sparsest solution compatible with the linear measurements. This is because whenever the decoder used in the proof of Theorem 32 (as described in Remark 16) succeeds, the ℓ_0-minimization decoder also succeeds. However, the ℓ_0-minimization decoder is both computationally expensive and unstable. The optimality of the ℓ_0-minimization decoder sparse signals has also been observed in [10, Section IV-A1] based on replica heuristics.

Theorem 33 (Linear encoding: achievability for self-similar sources)**.** *Suppose that the source is memoryless with a self-similar distribution that satisfies the open set condition. Then*

$$\mathsf{R}^*(\epsilon) \le d(X) \tag{5.23}$$

for all $0 < \epsilon < 1$.

Proof. See Section 5.3.3. □

In Theorems 31 – 32 it has been shown that block error probability ϵ is achievable for Lebesgue-a.e. linear encoder. Therefore, choosing any random matrix whose distribution is absolutely continuous (e.g. i.i.d. Gaussian random matrices) satisfies block error probability ϵ almost surely.

Now, we drop the restriction that the encoder is linear, allowing very general encoding rules. Let us first consider the case where both the encoder and decoder are constrained to be continuous. It turns out that zero rate is achievable in this case.

Theorem 34 (Continuous encoder and decoder)**.** *For general sources,*

$$\mathsf{R}_0(\epsilon) = 0 \tag{5.24}$$

for all $0 < \epsilon \le 1$.

Proof. Since $(\mathbb{R}^n, \mathcal{B}_{\mathbb{R}^n})$ and $([0,1], \mathcal{B}_{[0,1]})$ are Borel isomorphic, there exist Borel measurable $f : \mathbb{R}^n \to [0,1]$ and $g : [0,1] \to \mathbb{R}^n$, such that $g = f^{-1}$. By Lusin's theorem [101, Theorem 7.10], there exists a compact set $K \subset \mathbb{R}^n$ such that f restricted on K is continuous and $\mathbb{P}\{X^n \notin K\} < \epsilon$. Since K is compact and f is injective on K, f is a homeomorphism from K to $T \triangleq f(K)$. Hence $g : T \to \mathbb{R}^n$ is continuous. Since both K and T are closed, by the Tietze extension theorem [91], f and g can be extended to continuous $f' : \mathbb{R}^n \to [0,1]$ and $g' : [0,1] \to \mathbb{R}^n$ respectively. Using f' and g' as the new encoder and decoder, the error probability satisfies $\mathbb{P}\{g'(f'(X^n)) \neq X^n\} \leq \mathbb{P}\{X^n \notin K\} < \epsilon$. □

Employing similar arguments as in the proof of Theorem 34, we see that imposing additional continuity constraints on the encoder (resp. decoder) has almost no impact on the fundamental limit $\mathsf{R}(\epsilon)$ (resp. $\mathsf{R}^*(\epsilon)$). This is because a continuous encoder (resp. decoder) can be obtained at the price of an arbitrarily small increase of error probability, which can be chosen to vanish as n grows.

Theorems 35 - 37 deal with Lipschitz decoding in Euclidean spaces.

Theorem 35 (Lipschitz decoding: General converse)**.** *Suppose that the source is memoryless. If $\overline{d}(X) < \infty$, then*

$$\mathsf{R}(\epsilon) \geq \overline{d}(X) \tag{5.25}$$

for all $0 < \epsilon < 1$.

Proof. See Section 5.4.2. □

Theorem 36 (Lipschitz decoding: Achievability for discrete/continuous mixture)**.** *Suppose that the source is memoryless with a discrete-continuous mixed distribution. Then*

$$\mathsf{R}(\epsilon) = d(X) \tag{5.26}$$

for all $0 < \epsilon < 1$.

Proof. See Section 5.4.3. □

For sources with a singular distribution, in general there is no simple answer due to their fractal nature. For an important class of singular measures, namely self-similar measures generated from i.i.d. digits (e.g. generalized Cantor distribution), the information dimension turns out to be the fundamental limit for lossless compression with Lipschitz decoder.

Theorem 37 (Lipschitz decoding: Achievability for self-similar measures)**.** *Suppose that the source is memoryless and bounded, and its M-ary expansion consists of independent identically distributed digits. Then*

$$\mathsf{R}(\epsilon) = d(X) \tag{5.27}$$

for all $0 < \epsilon < 1$. Moreover, if the distribution of each bit is equiprobable on its support, then (5.27) holds for $\epsilon = 0$

Proof. See Section 5.4.4. □

Example 6. As an example, we consider the setup in Theorem 37 with $M = 3$ and $P = \{p, 0, q\}$, where $p+q = 1$. The associated invariant set is the middle third Cantor set C [18] and X is supported on C. The distribution of X, denoted by μ, is called the generalized Cantor distribution [35]. In the ternary expansion of X, each digit is independent and takes value 0 and 2 with probability p and q respectively. Then by Theorem 37, for any $0 < \epsilon < 1$, $\mathsf{R}(\epsilon) = \frac{h(p)}{\log 3}$. Furthermore, when $p = 1/2$, μ coincides with the 'uniform' distribution on C, i.e., the standard Cantor distribution. Hence we have a stronger result that $\mathsf{R}(0) = \log_3 2 \approx 0.63$, i.e., exact lossless compression can be achieved with a Lipschitz continuous decompressor at the rate of the information dimension.

We conclude this section by discussing the fundamental limit of stable decoding (respect to the ℓ_∞ norm), which is given by the following tight result.

Theorem 38 (Stable decoding)**.** *Let the underlying metric be the ℓ_∞ distance. Suppose that the source is memoryless. Then for all $0 < \epsilon < 1$,*

$$\limsup_{\Delta \downarrow 0} \overline{\mathsf{R}}(\epsilon, \Delta) = \overline{d}(X), \tag{5.28}$$

that is, the minimum ϵ-achievable rate such that for all sufficiently small Δ there exists a Δ-stable coding strategy is given by $\overline{d}(X)$.

Proof. See Section 5.6. □

5.3 Lossless Linear Compression

In this section we analyze lossless compression with linear encoders, which are the basic elements in compressed sensing. Capitalizing on the approach of Minkowski-dimension compression developed in Chapter 3, we obtain achievability results for linear compression. For memoryless sources with a discrete-continuous mixed distribution, we also establish a converse (Theorem 32) which shows that the information dimension is the fundamental limit of lossless linear encoding.

5.3.1 Minkowski-dimension compression and linear compression

The following theorem establishes the relationship between Minkowski-dimension compression and linear data compression.

Theorem 39 (Linear encoding: general achievability)**.** *For general sources,*

$$\mathsf{R}^*(\epsilon) \leq \mathsf{R}_{\mathsf{B}}(\epsilon), \tag{5.29}$$

Moreover,

1. *For all linear encoders (except possibly those in a set of zero Lebesgue measure on the space of real matrices), block error probability ϵ is achievable.*

2. *For all $\epsilon' > \epsilon$ and*
$$0 < \beta < 1 - \frac{\mathsf{R}_{\mathsf{B}}(\epsilon)}{R}, \tag{5.30}$$
where $R > \mathsf{R}_{\mathsf{B}}(\epsilon)$ is the compression rate, there exists a β-Hölder continuous decoder that achieves block error probability ϵ'.

Consequently, in view of Theorem 39, the results on $\mathsf{R}_{\mathsf{B}}(\epsilon)$ for memoryless sources in Theorems 14 – 16 yield the achievability results in Theorems 31 – 33 respectively. Hölder exponents of the decoder can be found by replacing $\mathsf{R}_{\mathsf{B}}(\epsilon)$ in (5.30) by its respective upper bound.

Remark 15. For discrete-continuous sources, the achievability in Theorem 32 can be shown directly without invoking the general result in Theorem 39. See Remark 18. From the converse proof of Theorem 32, we see that effective compression can be achieved with linear encoders, i.e., $\mathsf{R}^*(\epsilon) < 1$, only if the source distribution is *not* absolutely continuous with respect to Lebesgue measure.

Linear embedding of low dimensional subsets in Banach spaces was previously studied in [19, 21, 20] etc. in a non-probabilistic setting. For example, [21, Theorem 1.1] showed that: for a subset S of a Banach space with $\overline{\dim}_{\mathrm{B}} S < k/2$, there exists a bounded linear function that embeds S into $\mathbb{R}^k$. Here in a probabilistic setup, the embedding dimension can be improved by a factor of two, in view of the following finite-dimensional version of Theorem 39, whose proof is almost identical.

Theorem 40. *Let X^n be a random vector with $\overline{\dim}^{\epsilon}_{\mathrm{B}}(P_{X^n}) \le k$. Let $m > k$. Then for Lebesgue almost every $\mathbf{A} \in \mathbb{R}^{m\times n}$, there exists a $\left(1 - \frac{k}{m}\right)$-Hölder continuous function $g : \mathbb{R}^m \to \mathbb{R}^n$, such that $\mathbb{P}\left\{g(\mathbf{A}X^n) \neq X^n\right\} \ge 1 - \epsilon$.*

Remark 16. In Theorem 40, the decoder can be chosen as follows: by definition of $\overline{\dim}^{\epsilon}_{\mathrm{B}}$, there exists $U \subset \mathbb{R}^n$, such that $\overline{\dim}_{\mathrm{B}}(U) \le k$. Then if x^n is the *unique* solution to the linear equation $\mathbf{A}x^n = y^k$ in U, the decoder outputs $g(y^k) = x^n$; otherwise $g(y^k) = 0$.

Particularizing Theorem 40, we obtain a non-asymptotic result of lossless linear compression for k-sparse vectors (i.e., with no more than k non-zero components), which is relevant to compressed sensing.

Corollary 2. *Denote the collection of all k-sparse vectors in $\mathbb{R}^n$ by*
$$\Sigma_k = \{x^n \in \mathbb{R}^n : \|x^n\|_0 \le k\}. \tag{5.31}$$
Let μ be a σ-finite Borel measure on $\mathbb{R}^n$. Then given any $l \ge k+1$, for Lebesgue-a.e. $l \times n$ real matrix $\mathbf{H}$, there exists a Borel function $g : \mathbb{R}^l \to \mathbb{R}^n$, such that $g(\mathbf{H}x^n) = x^n$ for μ-a.e. $x^n \in \Sigma_k$. Moreover, when μ is finite, for any $\epsilon > 0$ and $0 < \beta < 1 - \frac{k}{l}$, there exists a $l \times n$ matrix $\mathbf{H}^*$ and $g^* : \mathbb{R}^l \to \mathbb{R}^n$, such that $\mu\{x^n \in \Sigma_k : g^*(\mathbf{H}^* x^n) \neq x^n\} \le \epsilon$ and g^* is β-Hölder continuous.

Remark 17. The assumption that the measure μ is σ-finite is essential, because the validity of Corollary 2 hinges upon Fubini's theorem, where σ-finiteness is an indispensable requirement. Consequently, if μ is the distribution of a k-sparse random vector with uniformly chosen support and Gaussian distributed non-zero entries, we conclude that all k-sparse vectors can be linearly compressed except for a subset of zero measure under μ. On the other hand, if μ is the counting measure on Σ_k, Corollary 2 no longer applies because that μ is not σ-finite. In fact, if $l < 2k$, no linear encoder from $\mathbb{R}^n$ to $\mathbb{R}^l$ works for *every* k-sparse vector, even if no regularity constraint is imposed on the decoder. This is because no $l \times n$ matrix acts injectively on Σ_k. To see this, define $\mathcal{C} - \mathcal{C} = \{x - y : x, y \in \mathcal{C}\}$ for $\mathcal{C} \subset \mathbb{R}^n$. Then for any $l \times n$ matrix $\mathbf{H}$,

$$(\Sigma_k - \Sigma_k) \cap \mathrm{Ker}(\mathbf{H}) = \Sigma_{2k} \cap \mathrm{Ker}(\mathbf{H}) \neq \{0\}. \tag{5.32}$$

Hence there exist two k-sparse vectors that have the same image under $\mathbf{H}$.

On the other hand, $l = 2k$ is sufficient to linearly compress all k-sparse vectors, because (5.32) holds for Lebesgue-a.e. $2k \times n$ matrix $\mathbf{H}$. To see this, choose $\mathbf{H}$ to be a random matrix with i.i.d. entries according to some continuous distribution (e.g., Gaussian). Then (5.32) holds if and only if all $2k \times 2k$ submatrices formed by $2k$ columns of $\mathbf{H}$ are invertible. This is an almost sure event, because the determinant of each of the $\binom{n}{2k}$ submatrices is an absolutely continuous random variable. The sufficiency of $l = 2k$ is a bit stronger than the linear embedding result mentioned in the discussion before Theorem 40, which gives $l = 2k+1$. For an explicit construction of such a $2k \times n$ matrix, we can choose $\mathbf{H}$ to be the cosine matrix $H_{ij} = \cos\left(\frac{(i-1)j\pi}{n}\right)$ (see Appendix C.1).

5.3.2 Auxiliary results

Let $0 < k < n$. Denote by $G(n, k)$ the Grassmannian manifold [35] consisting of all k-dimensional subspaces of $\mathbb{R}^n$. For $V \in G(n, k)$, the orthogonal projection from $\mathbb{R}^n$ to V defines a linear mapping $\mathrm{proj}_V : \mathbb{R}^n \to \mathbb{R}^n$ of rank k. The technique we use in the achievability proof of linear analog compression is to use the random orthogonal projection proj_V as the encoder, where V is distributed according to the invariant probability measure on $G(n, k)$, denoted by $\gamma_{n,k}$ [35]. The relationship between $\gamma_{n,k}$ and the Lebesgue measure on $\mathbb{R}^{k\times n}$ is shown in the following lemma.

Lemma 14 ([35, Exercise 3.6])**.** *Denote the rows of a $k \times n$ matrix $\mathbf{H}$ by $H_1, \ldots, H_k$, the row span of $\mathbf{H}$ by $\mathrm{Im}(\mathbf{H}^{\mathrm{T}})$, and the volume of the unit ℓ_2-ball in $\mathbb{R}^n$ by $\alpha(n)$. Set*

$$\mathcal{B}_L = \{\mathbf{H} \in \mathbb{R}^{k\times n} : \|H_i\|_2 \leq L, i = 1, \ldots, k\}. \tag{5.33}$$

Then for $\mathcal{S} \subset G(n, k)$ measurable, i.e., a collection of k-dimensional subspaces of $\mathbb{R}^n$,

$$\gamma_{n,k}(\mathcal{S}) = \alpha(n)^{-k}\mathrm{Leb}\{\mathbf{H} \in \mathcal{B}_1 : \mathrm{Im}(\mathbf{H}^{\mathrm{T}}) \in \mathcal{S}\}. \tag{5.34}$$

The following result states that a random projection of a given vector is not too small with high probability. It plays a central role in estimating the probability of "bad" linear encoders.

Lemma 15 ([35, Lemma 3.11])**.** *For any $x^n \in \mathbb{R}^n \backslash \{0\}$,*

$$\gamma_{n,k}\{V : \|\mathrm{proj}_V x^n\|_2 \leq \delta\} \leq \frac{2^n \delta^k}{\alpha(n)\,\|x^n\|_2^k} \tag{5.35}$$

To show the converse part of Theorem 32 we will invoke the Steinhaus Theorem as an auxiliary result:

Lemma 16 (Steinhaus [102])**.** *For any measurable set $\mathcal{C} \subset \mathbb{R}^n$ with positive Lebesgue measure, there exists an open ball centered at 0 contained in $\mathcal{C} - \mathcal{C}$.*

Lastly, we give a characterization of the fundamental limit of lossless linear encoding as follows.

Lemma 17. *$\mathsf{R}^*(\epsilon)$ is the infimum of $R > 0$ such that for sufficiently large n, there exists a Borel set $S^n \subset \mathbb{R}^n$ and a linear subspace $H^n \subset \mathbb{R}^n$ of dimension at least $\lceil (1-R)n \rceil$, such that*

$$(S^n - S^n) \cap H^n = \{0\}. \tag{5.36}$$

Proof. (Converse): Suppose R is ϵ-achievable with linear encoders, i.e. $R \geq \mathsf{R}^*(\epsilon)$. Since any linear function between two finite-dimensional linear spaces is uniquely represented by a matrix, by Definition 13, for sufficiently large n, there exists a matrix $\mathbf{H} \in \mathbb{R}^{\lfloor Rn \rfloor \times n}$ and $g : \mathbb{R}^{\lfloor Rn \rfloor} \to \mathbb{R}^n$ such that

$$\mathbb{P}\left\{g(\mathbf{H}X^n) = X^n\right\} \geq 1 - \epsilon. \tag{5.37}$$

Thus there exists $S^n \subset \mathbb{R}^n$ such that $\mathbb{P}\left\{X^n \in S^n\right\} \geq 1 - \epsilon$ and $\mathbf{H}$ as a linear mapping restricted to S^n is *injective*, which implies that for any $x^n \neq y^n$ in S^n, $\mathbf{H}x^n \neq \mathbf{H}y^n$, that is, $x^n - y^n \notin \mathrm{Ker}(\mathbf{H})$. Therefore,

$$(S^n - S^n) \cap \mathrm{Ker}(\mathbf{H}) = \{0\}. \tag{5.38}$$

Note that $\dim \mathrm{Ker}(\mathbf{H}) \geq n - \lfloor Rn \rfloor$, hence (5.36) holds.

(Achievability): Suppose that such a sequence of S^n and H^n exists then, by definition of $\mathsf{R}^*(\epsilon)$, $R \geq \mathsf{R}^*(\epsilon)$. □

5.3.3 Proofs

Proof of Theorem 39. We first show (5.29). Fix $0 < \delta' < \delta$ arbitrarily. Let $R = \mathsf{R}_{\mathsf{B}}(\epsilon) + \delta$, $k = \lfloor Rn \rfloor$ and $k' = \lfloor (\mathsf{R}_{\mathsf{B}}(\epsilon) + \delta')n \rfloor$. We show that there exists a matrix $\mathbf{H}_n \in \mathbb{R}^{n \times n}$ of rank k and $g_n : \mathrm{Im}(\mathbf{H}_n) \to \mathbb{R}^n$ Borel measurable such that

$$\mathbb{P}\left\{g_n(\mathbf{H}_n X^n) = X^n\right\} \geq 1 - \epsilon \tag{5.39}$$

for sufficiently large n.

By definition of $\mathsf{R}_{\mathsf{B}}(\epsilon)$, there exists $N_0 \in \mathbb{N}$ such that for all $n \geq N_0$, there exists a compact $U_n \subset \mathbb{R}^n$ such that $\mathbb{P}\{X^n \in U_n\} \geq 1 - \epsilon$ and $\overline{\dim}_{\mathrm{B}} U_n \leq k'$. Given an encoding matrix $\mathbf{H}$, define the decoder $g_{\mathbf{H}} : \mathrm{Im}(\mathbf{H}) \to \mathbb{R}^n$ as

$$g_{\mathbf{H}}(y^k) = \min \mathbf{H}^{-1}(\{y^k\}) \cap U_n, \tag{5.40}$$

where the min is taken componentwise[2]. Since $\mathbf{H}^{-1}(\{y^k\})$ is closed and U_n is compact, $\mathbf{H}^{-1}(\{y^k\}) \cap U_n$ is compact. Hence $g_{\mathbf{H}}$ is well-defined.

Next consider a random orthogonal projection matrix $\mathbf{\Phi} = \mathrm{proj}_{\boldsymbol{V}}$ independent of X^n, where $\boldsymbol{V} \in G(n,k)$ is a random k-dimensional subspace distributed according to the invariant measure $\gamma_{n,k}$ on $G(n,k)$. We show that for all $n \geq N_0$,

$$\mathbb{P}\{g_{\mathbf{\Phi}}(\mathbf{\Phi} X^n) \neq X^n\} \leq \epsilon, \tag{5.41}$$

which implies that there exists at least one realization of $\mathbf{\Phi}$ that satisfies (5.39). To that end, we define

$$p_{\mathrm{e}}(\mathbf{H}) = \mathbb{P}\{X^n \in U_n, g_{\mathbf{H}}(\mathbf{H} X^n) \neq X^n\} \tag{5.42}$$

and use the union bound:

$$\mathbb{P}\{g_{\mathbf{\Phi}}(\mathbf{\Phi} X^n) \neq X^n\} \leq \mathbb{P}\{X^n \notin U_n\} + \mathbb{E}\left[p_{\mathrm{e}}(\mathbf{\Phi})\right], \tag{5.43}$$

where the first term $\mathbb{P}\{X^n \notin U_n\} \leq \epsilon$. Next we show that the second term is *zero*. Let $\mathcal{U}(x^n) = U_n - x^n$. Then

$$\begin{aligned}
\mathbb{E}\left[p_{\mathrm{e}}(\mathbf{\Phi})\right] &= \int_{U_n} \mathbb{P}\{g_{\mathbf{\Phi}}(\mathbf{\Phi} x^n) \neq x^n\} \mu^n(\mathrm{d}x^n) && (5.44)\\
&\leq \int_{U_n} \mathbb{P}\{\exists y^n \in \mathcal{U}(x^n) : \mathbf{\Phi} y^n = 0\} \mu^n(\mathrm{d}x^n) && (5.45)\\
&= \int_{U_n} \mathbb{P}\{\mathrm{Ker}(\mathbf{\Phi}) \cap \mathcal{U}(x^n) \neq \{0\}\} \mu^n(\mathrm{d}x^n), && (5.46)
\end{aligned}$$

We show that for all $x^n \in U_n$, $\mathbb{P}\{\mathrm{Ker}(\mathbf{\Phi}) \cap \mathcal{U}(x^n) \neq \{0\}\} = 0$. To this end, let

$$0 < \beta < \frac{\delta - \delta'}{R} = \frac{k - k'}{k}. \tag{5.47}$$

Define

$$\begin{aligned}
T_j(x^n) &= \{y^n \in \mathcal{U}(x^n) : \|y^n\|_2 \geq 2^{-j\beta}\}, && (5.48)\\
Q_j &= \{\mathbf{H} : \forall y^n \in T_j(x^n), \|\mathbf{H} y^n\|_2 \geq 2^{-j}\}. && (5.49)
\end{aligned}$$

[2]Alternatively, we can use any other tie-breaking strategy as long as Borel measurability is satisfied.

Then

$$\bigcup_{j\geq 1} T_j(x^n) = \mathcal{U}(x^n)\backslash\{0\}. \tag{5.50}$$

Observe that $\mathbf{H} \in \bigcap_{j\geq J} Q_j$ implies that

$$\|\mathbf{H}y^n\|_2 \geq 2^{-(J+1)} \|y^n\|_2^{\frac{1}{\beta}}, \quad \forall y^n \in \mathcal{U}(x^n). \tag{5.51}$$

Therefore $\mathbf{H} \in Q_j$ for all but a finite number of j's if and only if

$$\|\mathbf{H}y^n\|_2 \geq C(\mathbf{H}, x^n) \|y^n\|_2^{\frac{1}{\beta}}, \quad \forall y^n \in \mathcal{U}(x^n) \tag{5.52}$$

for some $C(\mathbf{H}, x^n) > 0$.

Next we show that $\mathbf{\Phi} \in Q_j$ for all but a finite number of j's with probability one. Cover $\mathcal{U}(x^n)$ with $2^{-(j+1)}$-balls. The centers of those balls that intersect $T_j(x^n)$ are denoted by $\{w_{i,j}, i = 1, \ldots, M_j\}$. Pick $y_{i,j} \in T_j(x^n) \cap B(w_{i,j}, 2^{-(j+1)})$. Then $B(w_{i,j}, 2^{-(j+1)}) \subset B(y_{i,j}, 2^{-j})$, hence $\{B(y_{i,j}, 2^{-j}) : i = 1, \ldots, M_j\}$ cover $T_j(x^n)$. Suppose $\|\mathbf{H}y_{i,j}\|_2 \geq 2^{-(j-1)}$, then for any $z \in B(y_{i,j}, 2^{-j})$,

$$\begin{aligned}
\|\mathbf{H}z\|_2 &\geq \|\mathbf{H}y_{i,j}\|_2 - \|\mathbf{H}(y_{i,j} - z)\|_2 && (5.53)\\
&\geq \|\mathbf{H}y_{i,j}\|_2 - \|y_{i,j} - z\|_2 && (5.54)\\
&\geq 2^{-j}, && (5.55)
\end{aligned}$$

where (5.54) follows because $\mathbf{H}$ is an orthogonal projection. Thus $\|\mathbf{H}y_{i,j}\|_2 \geq 2^{-(j-1)}$ for all $i = 1, \ldots, M_j$ implies that $\mathbf{H} \in Q_j$. Therefore by the union bound,

$$\sum_{j\geq 1} \mathbb{P}\left\{\mathbf{\Phi} \in Q_j^{\mathrm{c}}\right\} \leq \sum_{j\geq 1} \sum_{i=1}^{M_j} \mathbb{P}\left\{\|\mathbf{\Phi}y_{i,j}\|_2 \leq 2^{-(j-1)}\right\}. \tag{5.56}$$

By Lemma 15,

$$\begin{aligned}
\mathbb{P}\left\{\|\mathbf{\Phi}y_{i,j}\|_2 \leq 2^{-(j-1)}\right\} &= \gamma_{n,k}(\{V : \|\mathrm{proj}_V y_{i,j}\|_2 \leq 2^{-(j-1)}\}) && (5.57)\\
&\leq \frac{2^{n-(j-1)k}}{\alpha(n)\|y_{i,j}\|_2^k} && (5.58)\\
&\leq \frac{2^{n+k(1-j+j\beta)}}{\alpha(n)}, && (5.59)
\end{aligned}$$

where (5.59) is due to $\|y_{i,j}\|_2 \geq 2^{-j\beta}$, because $y_{i,j} \in T_j(x^n)$. Since $\overline{\dim}_{\mathrm{B}} U_n < k'$, there is a constant C_1 such that $N_{U_n}(2^{-j-1}) \leq C_1 2^{jk'}$. Since $\mathcal{U}(x^n)$ is a translation of U_n, it follows that

$$M_j \leq C_1 2^{jk'}. \tag{5.60}$$

Thus

$$\sum_{j\geq 1}\mathbb{P}\left\{\mathbf{\Phi}\in Q_j^{\mathrm{c}}\right\}\leq C_1\alpha(n)^{-1}2^{n+k}\sum_{j\geq 1}2^{j(k'-(1-\beta)k)} \tag{5.61}$$

$$<\infty, \tag{5.62}$$

where

- (5.61): by substituting (5.59) and (5.60) into (5.56);
- (5.62): by (5.47).

Therefore by the Borel-Cantelli lemma, $\mathbf{\Phi}\in Q_j$ for all but a finite number of j's with probability one. Hence

$$\mathbb{P}\left\{\left\|\mathbf{\Phi}y^n\right\|_2\geq C(\mathbf{\Phi},x^n)\left\|y^n\right\|_2^{\frac{1}{\beta}},\ \forall y^n\in\mathcal{U}(x^n)\right\}=1, \tag{5.63}$$

which implies that for any $x^n\in U_n$,

$$\mathbb{P}\left\{\mathrm{Ker}(\mathbf{\Phi})\cap\mathcal{U}(x^n)=\{0\}\right\}=1. \tag{5.64}$$

In view of (5.46),

$$\mathbb{E}\left[p_{\mathrm{e}}(\mathbf{\Phi})\right]=\mathbb{P}\left\{X^n\in U_n, g_{\mathbf{\Phi}}(\mathbf{\Phi}X^n)\neq X^n\right\}=0, \tag{5.65}$$

whence (5.41) follows. This shows the ϵ-achievability of R. By the arbitrariness of $\delta>0$, (5.29) is proved.

Now we show that

$$\mathbb{P}\left\{g_{\mathbf{A}}(\mathbf{A}X^n)\neq X^n\right\}\leq\epsilon \tag{5.66}$$

holds for all $\mathbf{A}\in\mathbb{R}^{k\times n}$ except possibly on a set of zero Lebesgue measure, where $g_{\mathbf{A}}$ is the corresponding decoder for $\mathbf{A}$ defined in (5.40). Note that,

$$\left\{\mathbf{A}:\mathbb{P}\left\{g_{\mathbf{A}}(\mathbf{A}X^n)\neq X^n\right\}>\epsilon\right\}\subset\left\{\mathbf{A}:p_{\mathrm{e}}(\mathbf{A})>0\right\} \tag{5.67}$$

$$=\left\{\mathbf{A}:p_{\mathrm{e}}(\mathrm{proj}_{\mathrm{Im}(\mathbf{A}^{\mathrm{T}})})>0\right\}, \tag{5.68}$$

where

- (5.67): by (5.43).
- (5.68): by (5.46) and $\mathrm{Ker}(\mathbf{A})=\mathrm{Ker}(\mathrm{proj}_{\mathrm{Im}(\mathbf{A}^{\mathrm{T}})})$.

Define

$$\mathcal{S}\triangleq\left\{V\in G(n,k):p_{\mathrm{e}}(\mathrm{proj}_V)>0\right\}. \tag{5.69}$$

Recalling $\mathcal{B}_L$ defined in (5.33), we have

$$\mathrm{Leb}\left\{\mathbf{A}\in\mathcal{B}_1:p_{\mathrm{e}}(\mathrm{proj}_{\mathrm{Im}(\mathbf{A}^{\mathrm{T}})})>0\right\}=\mathrm{Leb}\left\{\mathbf{A}\in\mathcal{B}_1:\mathrm{Im}(\mathbf{A}^{\mathrm{T}})\in\mathcal{S}\right\} \tag{5.70}$$

$$=\alpha(n)^k\gamma_{n,k}(\mathcal{S}) \tag{5.71}$$

$$=0, \tag{5.72}$$

where

- (5.70): by (5.68) and since $\mathrm{Im}(\mathbf{A}^{\mathrm{T}}) \in G(n,k)$ holds Lebesgue-a.e.
- (5.71): by Lemma 14.
- (5.72): by (5.65).

Observe that (5.72) implies that for any L,

$$\mathrm{Leb}\big\{\mathbf{A} \in \mathcal{B}_L : p_{\mathrm{e}}(\mathrm{proj}_{\mathrm{Im}(\mathbf{A}^{\mathrm{T}})}) > 0\big\} = 0. \tag{5.73}$$

Since $\mathbb{R}^{n\times k} = \bigcup_{L\geq 1} \mathcal{B}_L$, in view of (5.68) and (5.73), we conclude that (5.66) holds Lebesgue-a.e.

Finally we show that for any $\epsilon' > \epsilon$, there exists a sequence of matrices and β-Hölder continuous decoders that achieves compression rate R and block error probability ϵ'. Since $\mathbf{\Phi} \in Q_j$ for all but a finite number of j's a.s., there exists a J_n (independent of x^n), such that

$$\mathbb{P}\left\{\mathbf{\Phi} \in \bigcap_{j\geq J_n} Q_j\right\} \geq 1 - \epsilon' + \epsilon. \tag{5.74}$$

Thus by (5.51), for any $x^n \in U_n$,

$$\mathbb{P}\left\{\|\mathbf{\Phi} y^n\|_2 \geq 2^{-(J_n+1)} \|y^n\|_2^{\frac{1}{\beta}}, \forall y^n \in \mathcal{U}(x^n)\right\} \geq 1 - \epsilon' + \epsilon. \tag{5.75}$$

Integrating (5.75) with respect to $\mu^n(\mathrm{d}x^n)$ on U_n and by Fubini's theorem, we have

$$\begin{aligned}
&\mathbb{P}\big\{X^n \in U_n, \|\mathbf{\Phi}(y^n - X^n)\|_2 \geq 2^{-J_n+1} \|y^n - X^n\|_2^{\frac{1}{\beta}}, \forall y^n \in U_n\big\} \\
&= \int \mathbb{P}\big\{X^n \in U_n, \|\mathbf{H}(y^n - X^n)\|_2 \geq 2^{-J_n+1} \|y^n - X^n\|_2^{\frac{1}{\beta}}, \forall y^n \in U_n\big\} P_{\mathbf{\Phi}}(\mathrm{d}\mathbf{H}) \qquad (5.76)\\
&\geq 1 - \epsilon'. \qquad (5.77)
\end{aligned}$$

Hence there exists $S_n \subset U_n$ and an orthogonal projection matrix $\mathbf{H}_n$ of rank k, such that $\mathbb{P}\{X^n \in S_n\} \geq 1 - \epsilon'$ and for all $x^n, y^n \in S_n$,

$$\|\mathbf{H}_n(y^n - x^n)\|_2 \geq 2^{-(J_n+1)} \|y^n - x^n\|_2^{\frac{1}{\beta}} \tag{5.78}$$

for all $x^n, y^n \in S_n$. Therefore[3] $\mathbf{H}_n^{-1}|_{\mathbf{H}_n(S_n)}$ is $(2^{\beta(J_n+1)}, \beta)$-Hölder continuous. By the extension theorem of Hölder continuous mappings [103], $\mathbf{H}_n^{-1}$ can be extended to $g_n : \mathbb{R}^k \to \mathbb{R}^n$ that is β-Hölder continuous. Then

$$\mathbb{P}\{g_n(\mathbf{H}_n X^n) \neq X^n\} \leq \epsilon'. \tag{5.79}$$

Recall from (5.47) that $0 < \beta < \frac{R - \mathsf{R}_{\mathsf{B}}(\epsilon) - \delta'}{R}$. By the arbitrariness of δ', (5.30) holds. □

[3] $f|_A$ denotes the restriction of f on the subset A.

Remark 18. Without recourse to the general result in Theorem 39, the achievability for discrete-continuous sources in Theorem 32 can be proved directly as follows. In (5.43), choose

$$U_n = \{x^n \in \mathbb{R}^n : |\mathrm{spt}(x^n)| \le (\rho + \delta/2)n\}, \tag{5.80}$$

and consider $\mathbf{\Phi}$ whose entries are i.i.d. standard Gaussian (or any other absolutely continuous distribution on $\mathbb{R}$). Using linear algebra, it is straightforward to show that the second term in (5.43) is zero. Thus the block error probability vanishes since $\mathbb{P}\{X^n \notin U_n\} \to 0$.

Proof of Corollary 2. It is sufficient to consider $l = k+1$. We show that for Lebesgue-a.e. $l \times n$ matrix $\mathbf{H}$, $\mu\{x^n \in \Sigma_k : g_{\mathbf{H}}(\mathbf{H}x^n) \neq x^n\} = 0$. Since $g_{\mathbf{H}}(\mathbf{H}x^n) = x^n$ if and only if $g_{\mathbf{H}}(\mathbf{H}\alpha x^n) = \alpha x^n$ for any $\alpha > 0$, it is sufficient to show

$$\mu\{x^n \in S_k : g_{\mathbf{H}}(\mathbf{H}x^n) \neq x^n\} = 0, \tag{5.81}$$

where $S_k \triangleq \Sigma_k \cap B_2^n(0,1)$.

Consider the $l \times n$ random projection matrix $\mathbf{\Phi}$ defined in the proof of Theorem 39. Denote by $P_{\mathbf{\Phi}}$ the distribution of $\mathbf{\Phi}$ and by $\mu \times P_{\mathbf{\Phi}}$ the product measure on $\mathbb{R}^n \times \mathbb{R}^{l\times n}$. Fix $0 < \beta < 1 - \frac{k}{l}$. Observe that Σ_k is the union of $\binom{n}{k}$ k-dimensional subspaces of $\mathbb{R}^n$, and by Lemma 8, $\dim_{\mathrm{B}} S_k = k$. Following the same arguments leading to (5.77), we have

$$\mu \times P_{\mathbf{\Phi}}\{x^n \in S_k, \mathbf{H} : \|\mathbf{H}(y^n - x^n)\|_2 \ge C(\mathbf{H}, x^n)\, \|y^n - x^n\|_2^{\frac{1}{\beta}},\ \forall y^n \in S_k\} = \mu(S_k). \tag{5.82}$$

By Fubini's theorem and Lemma 14, (5.81) holds for Lebesgue-a.e. $\mathbf{H}$.

When μ is finite, a β-Hölder continuous decoder that achieves block error probability ϵ can be obtained using the same construction procedure as in the proof of Theorem 39. □

Finally we complete the proof of Theorem 32 by proving the converse.

Converse proof of Theorem 32. Let the distribution of X be defined as in (3.70). We show that for any $\epsilon < 1$, $\mathsf{R}^*(\epsilon) \ge \rho$. Since $\mathsf{R}^*(\epsilon) \ge 0$, assume $\rho > 0$. Fix an arbitrary $0 < \delta < \rho$. Suppose $R = \rho - \delta$ is ϵ-achievable. Let $k = \lfloor(\rho - \delta)n\rfloor$ and $k' = \lfloor(\rho - \delta/2)n\rfloor$. By Lemma 17, for sufficiently large n, there exist a Borel set S^n and a linear subspace $H^n \subset \mathbb{R}^n$, such that $\mathbb{P}\{X^n \in S^n\} \ge 1 - \epsilon$, $S^n - S^n \cap H^n = \{0\}$ and $\dim H^n \ge n - k$.

If $\rho = 1$, $\mu = \mu_{\mathrm{c}}$ is absolutely continuous with respect to Lebesgue measure. Therefore S^n has positive Lebesgue measure. By Lemma 16, $S^n - S^n$ contains an open ball in $\mathbb{R}^n$. Hence $S^n - S^n \cap H^n = \{0\}$ cannot hold for any subspace H^n with positive dimension. This proves $\mathsf{R}^*(\epsilon) \ge 1$. Next we assume that $0 < \rho < 1$.

Let

$$T_n = \{x^n \in \mathbb{R}^n : |\mathrm{spt}(x^n)| \ge k'\} \tag{5.83}$$

and $G_n = S^n \cap T_n$. By (3.72), for sufficiently large n, $\mathbb{P}\{X^n \notin T_n\} \le (1-\epsilon)/2$, hence $\mathbb{P}\{X^n \in G_n\} \ge (1-\epsilon)/2$.

Next we decompose G_n according to the generalized support of x^n:

$$G_n = \bigcup_{\substack{U \subset \{1,\ldots,n\} \\ |U| \geq k'}} C_U, \tag{5.84}$$

where we have denoted the disjoint subsets

$$C_U = \{x^n \in S^n : \mathrm{spt}(x^n) = U\}. \tag{5.85}$$

Then

$$\sum_{\substack{U \subset \{1,\ldots,n\} \\ |U| \leq k'}} \mathbb{P}\{X^n \in C_U\} = \mathbb{P}\{X^n \in G_n\} \tag{5.86}$$

$$\geq (1-\epsilon)/2 \ > 0. \tag{5.87}$$

So there exists $U \subset \{1,\ldots,n\}$ such that $|U| \leq k'$ and $\mathbb{P}\{X^n \in C_U\} > 0$.

Next we decompose each C_U according to $x^n_{U^c}$ which can only take countably many values. Let $j = |U|$. For $y \in \mathcal{A}^{n-j}$, let

$$B_y = \{x^n : x^n \in C_U, x^n_{U^c} = y\}, \tag{5.88}$$

$$D_y = \{x^n_U : x^n \in B_y\}. \tag{5.89}$$

Then C_U can be written as a disjoint union of B_y:

$$C_U = \bigcup_{y \in \mathcal{A}^{n-j}} B_y. \tag{5.90}$$

Since $\mathbb{P}\{X^n \in C_U\} > 0$, there exists $y \in \mathcal{A}^{n-j}$ such that $\mathbb{P}\{X^n \in B_y\} > 0$.

Note that

$$\mathbb{P}\{X^n \in B_y\} = \mu^n(B_y) \tag{5.91}$$

$$= \rho^j (1-\rho)^{n-j} \mu_{\mathrm{c}}^j(D_y) \prod_{i=1}^{n-j} \mu_{\mathrm{d}}(\{y_i\}) \tag{5.92}$$

$$> 0. \tag{5.93}$$

Therefore $\mu_{\mathrm{c}}^j(D_y) > 0$. Since μ_{c}^j is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}^j$, D_y has positive Lebesgue measure. By Lemma 16, $D_y - D_y$ contains an open ball $B(0,r)$ for some $r > 0$. Therefore we have

$$K \subset B_y - B_y \subset S^n - S^n, \tag{5.94}$$

where $K = \{x^n \in \mathbb{R}^n : x^n_{U^c} = 0, x^n_U \in B(0,r)\}$. Hence K contains j linear independent vectors, denoted by $\{a_1,\ldots,a_j\}$. Let $\{b_1,\ldots,b_m\}$ be a basis for H^n, where $m \geq n-k$ by assumption. Since $j = |U| \geq k' > k$, we conclude that $\{a_1,\ldots,a_j,b_1,\ldots,b_m\}$ are

linearly *dependent*. Therefore

$$\sum_{l=1}^{m} \beta_l b_l = \sum_{l=1}^{j} \alpha_l a_l \neq 0 \tag{5.95}$$

where $\alpha_i \neq 0$ and $\beta_l \neq 0$ for some i and l. If we choose those nonzero coefficients sufficiently small, then $\sum_{l=1}^{j} \alpha_l a_l \in K$ and $\sum_{l=1}^{m} \beta_l b_l \in H^n$ since H^n is a linear subspace. This contradicts $S^n - S^n \cap H^n = \{0\}$. Thus $\mathsf{R}^*(\epsilon) \geq \rho - \delta$, and $\mathsf{R}^*(\epsilon) \geq \rho$ follows from the arbitrariness of δ.

□

5.4 Lossless Lipschitz Decompression

In this section we study the fundamental limit of lossless compression with Lipschitz decoders. To facilitate the discussion, we first introduce several important concepts from geometric measure theory. Then we proceed to give proofs of Theorems 35 – 37.

5.4.1 Geometric measure theory

Geometric measure theory [104, 35] is an area of analysis studying the geometric properties of sets (typically in Euclidean spaces) through measure theoretic methods. One of the core concepts in this theory is *rectifiability*, a notion of smoothness or regularity of sets and measures. Basically a set is rectifiable if it is the image of a subset of a Euclidean space under some Lipschitz function. Rectifiable sets admit a smooth analog coding strategy. Therefore, lossless compression with Lipschitz decoders boils down to finding high-probability subsets of source realizations that are rectifiable. In contrast, the goal of conventional almost-lossless data compression is to show concentration of probability on sets of small cardinality. This characterization enables us to use results from geometric measure theory to study Lipschitz coding schemes.

Definition 14 (Hausdorff measure and dimension)**.** Let $s > 0$ and $A \subset \mathbb{R}^n$. Define

$$\mathcal{H}_\delta^s(A) = \inf\left\{\sum_i \operatorname{diam}(E_i)^s : A \subset \bigcup_i E_i, \operatorname{diam}(E_i) \leq \delta\right\}, \tag{5.96}$$

where $\operatorname{diam}(E_i) = \sup\{\|x-y\|_2 : x, y \in E_i\}$. Define the *s-dimensional Hausdorff measure* on $\mathbb{R}^n$ by

$$\mathcal{H}^s(A) = \lim_{\delta \downarrow 0} \mathcal{H}_\delta^s(A), \tag{5.97}$$

The *Hausdorff dimension* of A is defined by

$$\dim_{\mathrm{H}}(A) = \inf\{s : \mathcal{H}^s(A) < \infty\}. \tag{5.98}$$

Hausdorff measure generalizes both the counting measure and Lebesgue measure and provides a non-trivial way to measure low-dimensional sets in a high-dimensional space. When $s = n$, $\mathcal{H}^n$ is just a rescaled version of the usual n-dimensional Lebesgue measure [35, 4.3]; when $s = 0$, $\mathcal{H}^0$ reduces to the counting measure. For $0 < s < n$, $\mathcal{H}^s$ gives a non-trivial measure for sets of Hausdorff dimension s in $\mathbb{R}^n$, because if $\dim_{\mathrm{H}} A < s$, $\mathcal{H}^s(A) = 0$; if $\dim_{\mathrm{H}} A > s$, $\mathcal{H}^s(A) = \infty$. As an example, consider $n = 1$ and $s = \log_3 2$. Let C be the middle-third Cantor set in the unit interval, which has zero Lebesgue measure. Then $\dim_{\mathrm{H}} C = s$ and $\mathcal{H}^s(C) = 1$ [18, 2.3].

Definition 15 (Rectifiable sets, [104, 3.2.14])**.** $E \subset \mathbb{R}^n$ is called *m-rectifiable* if there exists a Lipschitz mapping from some bounded set in $\mathbb{R}^m$ onto E.

Definition 16 (Rectifiable measures, [35, Definition 16.6])**.** Let μ be a measure on $\mathbb{R}^n$. μ is called *m-rectifiable* if $\mu \ll \mathcal{H}^m$ and there exists a μ-a.s. set $E \subset \mathbb{R}^n$ that is m-rectifiable.

Several useful facts about rectifiability are presented as follows:

Lemma 18 ([104])**.**

1. *An l-rectifiable set is also m-rectifiable for $m \geq l$.*
2. *The Cartesian product of an m-rectifiable set and an l-rectifiable set is $(m+l)$-rectifiable.*
3. *The finite union of m-rectifiable sets is m-rectifiable.*
4. *Countable sets are 0-rectifiable.*

Using the notion of rectifiability, we give a sufficient condition for the ϵ-achievability of Lipschitz decompression:

Lemma 19. *$\mathsf{R}(\epsilon) \leq R$ if there exists a sequence of $\lfloor Rn \rfloor$-rectifiable sets $S^n \subset \mathbb{R}^n$ with*

$$\mathbb{P}\left\{X^n \in S^n\right\} \geq 1 - \epsilon \tag{5.99}$$

for all sufficiently large n.

Proof. See Appendix C.4. □

Definition 17 (k-dimensional density, [35, Definition 6.8])**.** Let μ be a measure on $\mathbb{R}^n$. The k-dimensional *upper and lower densities* of μ at x are defined as

$$\overline{D}_k(\mu, x) = \limsup_{r \downarrow 0} \frac{\mu(B(x,r))}{r^k}, \tag{5.100}$$

$$\underline{D}_k(\mu, x) = \liminf_{r \downarrow 0} \frac{\mu(B(x,r))}{r^k}, \tag{5.101}$$

If $\overline{D}_k(\mu, x) = \underline{D}_k(\mu, x)$, the common value is called the k-dimensional density of μ at x, denoted by $D_k(\mu, x)$.

The following important result in geometric measure theory gives a density characterization of rectifiability for Borel measures.

Theorem 41 ([105, Theorem 5.6], Preiss Theorem)**.** *A σ-finite Borel measure on $\mathbb{R}^n$ is m-rectifiable if and only if $0 < D_m(\mu, x) < \infty$ for μ-a.e. $x \in \mathbb{R}^n$.*

Recalling the expression for information dimension $d(P_X)$ in (2.7), we see that for the information dimension of a measure to be equal to m it requires that the exponent of the *average* measure of ϵ-balls equals m, whereas m-rectifiability of a measure requires that the measure of *almost every* ϵ-ball scales as $O(\epsilon^m)$, a much stronger condition than the existence of information dimension. Obviously, if a probability measure μ is m-rectifiable, then $d(\mu) = m$.

5.4.2 Converse

In view of the lossless Minkowski dimension compression results developed in Section 3.2, the general converse in Theorem 35 is rather straightforward to show. We prove the following non-asymptotic converse of Lipschitz decoding, which implies Theorem 35 in particular:

Theorem 42. *For any random vector X^n, if there exists a Borel measurable $f : \mathbb{R}^n \to \mathbb{R}^k$ and a Lipschitz continuous $g : \mathbb{R}^k \to \mathbb{R}^n$ such that $\mathbb{P}\{g(f(X^n)) \neq X^n\} \leq \epsilon$, then*

$$k \geq \overline{\dim}_{\mathrm{B}}^{\epsilon}(P_{X^n}) \tag{5.102}$$

$$\geq \overline{d}(X^n) - \epsilon n. \tag{5.103}$$

In view of (5.102) and the definition of $\mathsf{R}_{\mathsf{B}}(\epsilon)$ in (3.14), we immediately have the following general inequality

$$\mathsf{R}(\epsilon) \geq \mathsf{R}_{\mathsf{B}}(\epsilon). \tag{5.104}$$

If the source is memoryless and $\overline{d}(X) < \infty$, then it follows from Theorem 14 that $\mathsf{R}(\epsilon) \geq \overline{d}(X)$. This completes the proof of Theorem 35.

Note that using the weaker result in (5.103) yields the following converse for general input processes:

$$\mathsf{R}(\epsilon) \geq \limsup_{n\to\infty} \frac{\overline{d}(X^n)}{n} - \epsilon, \tag{5.105}$$

which, for memoryless inputs, becomes the following *weak* converse:

$$\mathsf{R}(\epsilon) \geq \overline{d}(X) - \epsilon. \tag{5.106}$$

To get rid of ϵ in (5.106), it is necessary to use the *strong* converse in Theorem 14.

Proof of Theorem 42. To prove (5.102), denote $C = \{f(x^n) \in \mathbb{R}^n : g(f(x^n)) = x^n\} \subset \mathbb{R}^k$. Then

$$k \geq \overline{\dim}_{\mathrm{B}}(C) \tag{5.107}$$
$$\geq \overline{\dim}_{\mathrm{B}}(g(C)) \tag{5.108}$$
$$\geq \overline{\dim}_{\mathrm{B}}^{\epsilon}(P_{X^n}), \tag{5.109}$$

where

- (5.107): Minkowski dimension never exceeds the ambient dimension;
- (5.108): Minkowski dimension never increases under Lipschitz mappings [35, Exercise 7.6, p.108] (see Lemma 20 below for a proof);
- (5.109): by $\mathbb{P}\{X^n \in g(C)\} \geq 1 - \epsilon$ and (3.4).

It remains to show (5.103). By definition of $\overline{\dim}_{\mathrm{B}}^{\epsilon}$, for any $\delta > 0$, there exists E such that $P_{X^n}(E) \geq 1 - \epsilon$ and $\overline{\dim}_{\mathrm{B}}(E) \geq \overline{\dim}_{\mathrm{B}}^{\epsilon}(P_{X^n}) - \delta$. Since P_{X^n} can be written as a convex combination of $P_{X^n|X^n \in E}$ and $P_{X^n|X^n \notin E}$, applying Theorem 5 yields

$$\overline{d}(X^n) \leq \overline{d}(P_{X^n|X^n \in E})P_{X^n}(E) + \overline{d}(P_{X^n|X^n \notin E})(1 - P_{X^n}(E)) \tag{5.110}$$
$$\leq \overline{\dim}_{\mathrm{B}}^{\epsilon}(P_{X^n}) - \delta + \epsilon n, \tag{5.111}$$

where (5.111) holds because the information dimension of any distribution is upper bounded by the Minkowski dimension of its support [18]. Then the desired (5.103) follows from the arbitrariness of δ. □

Lemma 20. *Let $T \subset \mathbb{R}^k$ be bounded and $f : T \to \mathbb{R}^n$ be a Lipschitz mapping. Then $\overline{\dim}_{\mathrm{B}}(f(T)) \leq \overline{\dim}_{\mathrm{B}}(T)$.*

Proof. Note that

$$\overline{\dim}_{\mathrm{B}} T = \limsup_{\delta \downarrow 0} \frac{\log N_T(\delta)}{\log \frac{1}{\delta}}. \tag{5.112}$$

By definition of $N_T(\delta)$, there exists $\{x_i : i = 1, \ldots, N_T(\delta)\}$, such that T is covered by the union of $B(x_i, \delta), i = 1, \ldots, N_T(\delta)$. Setting $L = \mathrm{Lip}(f)$, we have

$$f(T) \subset \bigcup_{i=1}^{N_T(\delta)} f(B(x_i, \delta)) \subset \bigcup_{i=1}^{N_T(\delta)} B(f(x_i), L\delta), \tag{5.113}$$

which implies that $N_{f(T)}(L\delta) \leq N_T(\delta)$. Dividing both sides by $\log \frac{1}{\delta}$ and sending $\delta \to 0$ yield $\overline{\dim}_{\mathrm{B}}(f(T)) \leq \overline{\dim}_{\mathrm{B}}(T)$. □

5.4.3 Achievability for finite mixture

We first prove a general achievability result for finite mixtures, a corollary of which applies to discrete-continuous mixed distributions in Theorem 36.

Theorem 43 (Achievability of Finite Mixtures)**.** *Let the distribution μ of X be a mixture of finitely many Borel probability measures on $\mathbb{R}$, i.e.,*

$$\mu = \sum_{i=1}^{N} \rho_i \mu_i, \tag{5.114}$$

where $\{\rho_1, \ldots, \rho_N\}$ is a probability mass function. If R_i is ϵ_i-achievable with Lipschitz decoders for μ_i, $i = 1, \ldots, N$, then R is ϵ-achievable for μ with Lipschitz decoders, where

$$\epsilon = \sum_{i=1}^{N} \epsilon_i, \tag{5.115}$$

$$R = \sum_{i=1}^{N} \rho_i R_i. \tag{5.116}$$

Proof. By induction, it is sufficient to show the result for $N = 2$. Denote $\rho_1 = \rho = 1 - \rho_2$. Let $\{W_i\}$ be a sequence of i.i.d. binary random variables with $\mathbb{P}\{W_i = 1\} = \rho$. Let $\{X_i\}$ be a i.i.d. sequence of real-valued random variables, such that the distribution of each X_i conditioned on the events $\{W_i = 1\}$ and $\{W_i = 0\}$ are μ_1 and μ_2 respectively. Then $\{X_i : i \in \mathbb{N}\}$ is a memoryless process with common distribution μ. Since the claim of the theorem depends only on the *probability law* of the source, we base our calculation of block error probability on this specific construction.

Fix $\delta > 0$. Since R_1 and R_2 are achievable for μ_1 and μ_2 respectively, by Lemma 19, there exists N_1 such that for all $n > N_1$, there exists $S_1^n, S_2^n \subset \mathbb{R}^n$, with $\mu_1^n(S_1^n) \geq 1 - \epsilon_1$, $\mu_2^n(S_2^n) \geq 1 - \epsilon_2$ and S_1^n is $\lfloor R_1 n \rfloor$-rectifiable and S_2^n is $\lfloor R_2 n \rfloor$-rectifiable. Let

$$\mathcal{W}_n = \{w^n \in \{0,1\}^n : (\rho - \delta)n \leq \|w^n\|_0 \leq (\rho + \delta)n\}, \tag{5.117}$$

By WLLN, $\mathbb{P}\{W^n \in \mathcal{W}_n\} \to 1$. Hence for any $\epsilon' > 0$, there exists N_2, such that for all $n > N_2$,

$$\mathbb{P}\{W^n \in \mathcal{W}_n\} \geq 1 - \epsilon'. \tag{5.118}$$

Let

$$n > \max\left\{\frac{N_1}{\rho - \delta}, \frac{N_1}{1 - \rho - \delta}, N_2\right\}. \tag{5.119}$$

Define

$$\begin{aligned} S^n = \{x^n \in \mathbb{R}^n : \exists m, T, (\rho - \delta)n \leq m \leq (\rho + \delta)n, \\ T \subset \{1, \ldots, n\}, |T| = m, x_T \in S_1^m, x_{T^c} \in S_2^{n-m}\}. \end{aligned} \tag{5.120}$$

Next we show that S^n is $\lfloor (R+2\delta)n \rfloor$-rectifiable. For all $(\rho-\delta)n \le m \le (\rho+\delta)n$, it follows from (5.119) that

$$N_1 \le (\rho-\delta)n \le m, \tag{5.121}$$
$$N_1 \le (1-\rho-\delta)n \le n-m. \tag{5.122}$$

and

$$\begin{aligned} R_1 m + R_2(n-m) &\le R_1(\rho+\delta)n + R_2(1-\rho+\delta)n && (5.123)\\ &= ((R_1+R_2)\delta + R)n && (5.124)\\ &\le (2\delta + R)n && (5.125) \end{aligned}$$

where $R = \rho R_1 + (1-\rho)R_2$ according to (5.116). Observe that S^n is a finite union of subsets, each of which is a Cartesian product of a $\lfloor R_1 m \rfloor$-rectifiable set in $\mathbb{R}^m$ and a $\lfloor R_2(n-m) \rfloor$-rectifiable set in $\mathbb{R}^{n-m}$. In view of Lemma 18 and (5.125), S^n is $\lfloor (R+2\delta)n \rfloor$-rectifiable,

Now we calculate the measure of S^n under μ^n:

$$\begin{aligned} \mathbb{P}\{X^n \in S^n\} &\ge \mathbb{P}\{X^n \in S^n, W^n \in \mathcal{W}_n\} && (5.126)\\ &= \sum_{m=\lceil(\rho-\delta)n\rceil}^{\lfloor(\rho+\delta)n\rfloor} \sum_{\|w^n\|_0=m} \mathbb{P}\{W^n = w^n\} \\ &\qquad \times \mathbb{P}\{X^n \in S^n \mid W^n = w^n\} && (5.127)\\ &\ge \sum_{m=\lceil(\rho-\delta)n\rceil}^{\lfloor(\rho+\delta)n\rfloor} \sum_{|T|=m} \sum_{\mathrm{supp}(w^n)=T} \mathbb{P}\{W^n = w^n\} \\ &\qquad \times \mathbb{P}\left\{X_T \in S_1^m, X_{T^c} \in S_2^{n-m} \mid W^n = w^n\right\} && (5.128)\\ &= \sum_{m=\lceil(\rho-\delta)n\rceil}^{\lfloor(\rho+\delta)n\rfloor} \sum_{\|w^n\|_0=m} \mathbb{P}\{W^n = w^n\} \\ &\qquad \times \mu_1^m(S_1^m)\mu_2^{n-m}(S_2^{n-m}) && (5.129)\\ &\ge \sum_{m=\lceil(\rho-\delta)n\rceil}^{\lfloor(\rho+\delta)n\rfloor} \sum_{\|w^n\|_0=m} (1-\epsilon_1)(1-\epsilon_2)\,\mathbb{P}\{W^n = w^n\} && (5.130)\\ &\ge (1-\epsilon)\,\mathbb{P}\{W^n \in \mathcal{W}_n\} && (5.131)\\ &\ge 1-\epsilon-\epsilon', && (5.132) \end{aligned}$$

where

- (5.128): by construction of S_n in (5.120).
- (5.130): by (5.121) and (5.122).
- (5.131): $\epsilon = \epsilon_1 + \epsilon_2$ according to (5.115).
- (5.132): by (5.118).

In view of Lemma 19, $R + 2\delta$ is $(\epsilon + \epsilon')$-achievable for X. By the arbitrariness of δ and ϵ', the ϵ-achievability of R follows. □

Proof of Theorem 36. Let the distribution of X be $\mu = (1-\rho)\mu_{\mathrm{d}} + \rho\mu_{\mathrm{c}}$ as defined in (3.70), where μ_{d} is discrete and μ_{c} is absolutely continuous. By Lemma 18, countable sets are 0-rectifiable. For any $0 < \epsilon < 1$, there exists $A > 0$ such that $\mu_{\mathrm{c}}([-A, A]) > 1 - \epsilon$. By definition, $[-A, A]$ is 1-rectifiable. Therefore by Lemma 19 and Theorem 43, ρ is an ϵ-achievable rate for X. The converse follows from (5.104) and Theorem 15. □

5.4.4 Achievability for singular distributions

In this subsection we prove Theorem 37 for memoryless sources, using isomorphism results in ergodic theory. The proof outline is as follows: a classical result in ergodic theory states that Bernoulli shifts are isomorphic if they have the same entropy. Moreover, the homomorphism can be chosen to be *finitary*, that is, each coordinate only depends on finitely many coordinates. This finitary homomorphism naturally induces a Lipschitz decoder in our setup; however, the caveat is that the Lipschitz continuity is with respect to an *ultrametric* (Definition 18) that is *not* equivalent to the usual Euclidean distance. Nonetheless, by an arbitrarily small increase in the compression rate, the decoder can be modified to be Lipschitz with respect to the Euclidean distance. Before proceeding to the proof, we first present some necessary results of ultrametric spaces and finitary coding in ergodic theory.

Definition 18. Let (X, d) be a metric space. d is called an *ultrametric* if

$$d(x, z) \leq \max\{d(x, y), d(y, z)\} \tag{5.133}$$

for all $x, y, z \in X$.

A canonical class of ultrametric spaces is the ultrametric Cantor space [106]: let $\mathcal{X} = \{0, \ldots, M-1\}^{\mathbb{Z}_+}$ denote the set of all one-sided M-ary sequences $x = (x_0, \ldots)$. To endow $\mathcal{X}$ with an ultrametric, define

$$d_\alpha(x, y) = \begin{cases} 0 & x = y, \\ \alpha^{-\min\{n \in \mathbb{N} : x_n \neq y_n\}} & x \neq y. \end{cases} \tag{5.134}$$

Then for every $\alpha > 1$, d_α is an ultrametric on $\mathcal{X}$. In a similar fashion, we define an ultrametric on $[0, 1]^k$ by considering the M-ary expansion of real vectors. Similar to the binary expansion defined in (2.46), for $x^k \in [0, 1]^k$, $i \in \mathbb{Z}_+$, $M \in \mathbb{N}$ and $M \geq 2$, define

$$(x^k)_{M,i} = \lfloor M^i x^k \rfloor - M \lfloor M^{i-1} x^k \rfloor \in \{0, \ldots, M-1\}^k. \tag{5.135}$$

then

$$x^k = \sum_{i \in \mathbb{Z}_+} (x^k)_{M,i} M^{-i}. \tag{5.136}$$

Denoting for brevity

$$(x^k) = ((x^k)_{M,0}, (x^k)_{M,1}, \ldots), \tag{5.137}$$

(5.134) induces an ultrametric on $[0,1]^k$:

$$\hat{d}(x^k, y^k) = d_M((x^k), (y^k)). \tag{5.138}$$

It is important to note that $\hat{d}$ is *not* equivalent to the ℓ_∞ distance (or any ℓ_p distance), since we only have

$$\hat{d}(x^k, y^k) \geq \frac{1}{M} \left\| x^k - y^k \right\|_\infty . \tag{5.139}$$

To see the impossibility of the other direction of (5.139), consider $x = 1/M$ and $y = \sum_{i=2}^{l} M^{-i}$. As $l \to \infty$, $|x - y| \to 0$ but $\hat{d}(x, y)$ remains $1/M$. Therefore, a Lipschitz function with respect to $\hat{d}$ is *not* necessarily Lipschitz under $\|\cdot\|_\infty$. However, the following lemma bridges the gap if the dimension of the domain and the Lipschitz constant are allowed to increase.

Lemma 21. *Let $\hat{d}$ be the ultrametric on $[0,1]^k$ defined in (5.138). $W \subset [0,1]^k$ and $g : (W, \hat{d}) \to (\mathbb{R}^n, \|\cdot\|_\infty)$ is Lipschitz. Then there exists $W' \subset [0,1]^{k+1}$ and $g' : (W', \|\cdot\|_\infty) \to (\mathbb{R}^n, \|\cdot\|_\infty)$ such that $g(W) = g'(W')$ and g' is Lipschitz.*

Proof. See Appendix C.5. □

Next we recall several results on finitary coding of Bernoulli shifts. Kolmogorov-Ornstein theory studies whether two processes with the same entropy rate are isomorphic. In [107] Kean and Smorodinsky showed that two *double-sided* Bernoulli shifts of the same entropy are *finitarily isomorphic*. For the single-sided case, Del Junco [108] showed that there is a *finitary homomorphism* between two *single-sided* Bernoulli shifts of the same entropy, which is a finitary improvement of Sinai's theorem [109]. We will see how a finitary homomorphism of the digits is related to a real-valued Lipschitz function, and how to apply Del Junco's ergodic-theoretic result to our problem.

Definition 19 (Finitary homomorphisms)**.** Let $\mathcal{C}$ and $\mathcal{D}$ be finite sets. Let σ and τ denote the left shift operators on the product spaces $\mathcal{X} = \mathcal{C}^{\mathbb{Z}_+}$ and $\mathcal{Y} = \mathcal{D}^{\mathbb{Z}_+}$ respectively. Let μ and ν be measures on $\mathcal{X}$ and $\mathcal{Y}$ (with product σ-algebras). A *homomorphism* $\phi : (\mathcal{X}, \mu, \sigma) \to (\mathcal{Y}, \nu, \tau)$ is a measure preserving mapping that commutes with the shift operator, i.e., $\nu = \mu \circ \phi^{-1}$ and $\phi \circ \sigma = \tau \circ \phi$ μ-a.e. ϕ is said to be *finitary* if there exist sets of zero measure $A \subset \mathcal{X}$ and $B \subset \mathcal{Y}$ such that $\phi : \mathcal{X} \backslash A \to \mathcal{Y} \backslash B$ is continuous (with respect to the product topology).

Informally, finitariness means that for almost every $x \in \mathcal{X}$, $\phi(x)_0$ is determined by finitely many coordinates in x. The following lemma characterizes this intuition in precise terms:

Lemma 22 ([110, Conditions 5.1, p. 281])**.** *Let $\mathcal{X} = \mathcal{C}^{\mathbb{Z}_+}$ and $\mathcal{Y} = \mathcal{D}^{\mathbb{Z}_+}$. Let $\phi : (\mathcal{X}, \mu, \sigma) \to (\mathcal{Y}, \nu, \tau)$ be a homomorphism. Then the following statements are equivalent:*

1. *ϕ is finitary.*

2. *For μ-a.e. $x \in \mathcal{X}$, there exists $j(x) \in \mathbb{N}$, such that for any $x' \in \mathcal{X}$, $(x')_0^{j(x)} = (x)_0^{j(x)}$ implies that $(\phi(x))_0 = (\phi(x'))_0$.*

3. *For each $j \in \mathcal{D}$, the inverse image $\phi^{-1}\{y \in \mathcal{Y} : y_0 = j\}$ of each time-0 cylinder set in $\mathcal{Y}$ is, up to a set of measure 0, a countable union of cylinder sets in $\mathcal{X}$.*

Theorem 44 ([108, Theorem 1])**.** *Let P and Q be probability distributions on finite sets $\mathcal{C}$ and $\mathcal{D}$. Let $(\mathcal{X}, \mu) = (\mathcal{C}^{\mathbb{Z}_+}, P^{\mathbb{Z}_+})$ and $(\mathcal{Y}, \nu) = (\mathcal{D}^{\mathbb{Z}_+}, Q^{\mathbb{Z}_+})$. If P and Q each have at least three non-zero components and $H(P) = H(Q)$, then there is a finitary homomorphism $\phi : (\mathcal{X}, \mu) \to (\mathcal{Y}, \nu)$.*

We now use Lemmas 21 – 22 and Theorem 44 to prove Theorem 37.

Proof of Theorem 37. Without loss of generality, assume that the random variable satisfies $0 \le X \le 1$. Denote by P the common distribution of the M-ary digits of X. By Lemma 1,

$$d(X) = \frac{H(P)}{\log M}. \tag{5.140}$$

Fix n. Let $d = d(X)$, $k = \lceil dn \rceil$ and $\mathcal{C} = \{0, \dots, M-1\}^n$. Let Q be a probability measure on $\mathcal{D} = \{0, \dots, M-1\}^k$ such that $H(P^n) = H(Q)$. Such a Q always exists because $\log |\mathcal{D}| = k \log M \ge nH(P)$. Let $\mu = (P^n)^{\mathbb{Z}_+}$ and $\nu = Q^{\mathbb{Z}_+}$ denote the product measure on $\mathcal{C}^{\mathbb{Z}_+}$ and $\mathcal{D}^{\mathbb{Z}_+}$ respectively. Since μ and ν has the same entropy rate, by Theorem 44, there exists a finitary homomorphism $\phi : (\mathcal{D}^{\mathbb{Z}_+}, \nu, \sigma) \to (\mathcal{C}^{\mathbb{Z}_+}, \mu, \tau)$. By the characterization of finitariness in Lemma 22, for any $u \in \mathcal{D}^{\mathbb{Z}_+}$, there exists $j(u) \in \mathbb{N}$ such that $\phi(u)_0$ is determined only by $u_0, \dots, u_{j(u)}$. Denote the closed ultrametric ball

$$B_u = \{v \in \mathcal{D}^{\mathbb{Z}_+} : \hat{d}(u, v) \le M^{-(j(u)+1)}\}, \tag{5.141}$$

where $\hat{d}$ is defined in (5.138). Then for any $v \in B_u$, $\phi(u)_0 = \phi(v)_0$. Note that $\{B_u : u \in \mathcal{D}^{\mathbb{Z}_+}\}$ forms a *countable* cover of $\mathcal{D}^{\mathbb{Z}_+}$. This is because B_u is just a cylinder set in $\mathcal{D}^{\mathbb{Z}_+}$ with base $(u)_0^{j(u)}$, and the total number of cylinders is countable. Furthermore, since intersecting ultrametric balls are contained in each other [111], there exists a sequence $\{u^{(i)}\}$ in $\mathcal{D}^{\mathbb{Z}_+}$, such that $\bigcup_{i \in \mathbb{N}} B_{u^{(i)}}$ *partitions* $\mathcal{D}^{\mathbb{Z}_+}$. Therefore, for all $0 < \epsilon < 1$, there exists N, such that $\nu(E) \ge 1 - \epsilon$, where

$$E = \bigcup_{i=1}^{N} B_{u^{(i)}}. \tag{5.142}$$

For $x \in [0,1]^k$, recall the M-ary expansion of x defined in (5.137), denoted by $(x) \in \mathcal{D}^{\mathbb{Z}_+}$. Let

$$F = \phi(E) \subset \mathcal{C}^{\mathbb{Z}_+} \tag{5.143}$$

$$W = \{x \in [0,1]^k : (x) \in E\} \tag{5.144}$$

$$S = \{x \in [0,1]^n : (x) \in F\}. \tag{5.145}$$

Since ϕ is measure-preserving, $\mu = \nu \circ \phi^{-1}$, therefore

$$\mathbb{P}\{X^n \in S\} = \mu(F) = \nu(\phi^{-1}(F)) \geq \nu(E) \geq 1 - \epsilon. \tag{5.146}$$

Next we use ϕ to construct a real-valued Lipschitz mapping g. Define $g : [0,1]^k \to [0,1]^n$ by

$$g(x) = \sum_{i \in \mathbb{Z}_+} \phi((x))_i M^{-i}. \tag{5.147}$$

Since ϕ commutes with the shift operator, for all $z \in \mathcal{D}^{\mathbb{Z}_+}$, $\phi(z)_i = (\tau^i \phi(z))_0 = \phi(\tau^i z)_0$. Also, for $x \in [0,1]^k$, $\tau^i(x) = (M^i x)$. Therefore

$$g(x) = \sum_{i \in \mathbb{Z}_+} \phi((M^i x))_0 M^{-i}. \tag{5.148}$$

Next we proceed to show that $g : (W, \hat{d}) \to ([0,1]^n, \|\cdot\|_\infty)$ is Lipschitz. In view of (5.139) and (5.134), it is sufficient to show that $\phi : (E, d_M) \to (\mathcal{C}^{\mathbb{Z}_+}, d_M)$ is Lipschitz. Let

$$J = \max_{1 \leq i \leq N} j(u^{(i)}). \tag{5.149}$$

First observe that ϕ is M^J-Lipschitz on each ultrametric ball $B_{u^{(i)}}$ in E. To see this, consider distinct points $v, w \in B_{u^{(i)}}$. Let $d_M(v,w) = M^{-m}$. Then $m \geq j(u^{(i)}) + 1$. Since $(v)_0^{m-1} = (w)_0^{m-1}$, $\phi(v)$ and $\phi(w)$ coincide on their first $m - j(u^{(i)}) - 1$ digits. Therefore

$$d_M(\phi(v), \phi(w)) \leq M^{-m+j(u^{(i)})} \tag{5.150}$$

$$\leq M^{j(u^{(i)})} d_M(v,w) \tag{5.151}$$

$$\leq M^J d_M(v,w). \tag{5.152}$$

Since every closed ultrametric ball is also open [111, Proposition 18.4], $B_{u^{(1)}}, \ldots, B_{u^{(N)}}$ are disjoint, therefore ϕ is L-Lipschitz on E for some $L > 0$. Then for any $y, z \in W$,

$$\|g(y) - g(z)\|_\infty \leq M \hat{d}(g(y), g(z)) \tag{5.153}$$

$$= M d_M(\phi((y)), \phi((z))) \tag{5.154}$$

$$\leq ML\, d_M((y),(z)) \tag{5.155}$$

$$= ML\, \hat{d}(y,z). \tag{5.156}$$

where

- (5.153): by (5.139);
- (5.154) and (5.156): by (5.138);
- (5.155): by (5.152).

Hence $g : (W, \hat{d}) \to ([0,1]^n, \|\cdot\|_\infty)$ is Lipschitz.

By Lemma 21, there exists a subset $W' \subset [0,1]^{k+1}$ and a Lipschitz mapping $g' : (W', \|\cdot\|_\infty) \to (\mathbb{R}^n, \|\cdot\|_\infty)$ such that $g'(W') = g(W) = S$. By Kirszbraun's theorem [104, 2.10.43], we extend g' to a Lipschitz function $g_n : [0,1]^{k+1} \to \mathbb{R}^n$. Then $S = g_n(W)$. Since g_n is continuous and $[0,1]^{k+1}$ is compact, by Lemma 52, there exists a Borel function $f_n : \mathbb{R}^n \to [0,1]^{k+1}$, such that $g_n = f_n^{-1}$ on S.

To summarize, we have obtained a Borel function $f_n : \mathbb{R}^n \to [0,1]^{k+1}$ and a Lipschitz function $g_n : [0,1]^{k+1} \to \mathbb{R}^n$, where $k = \lceil dn \rceil$, such that $\mathbb{P}\{g_n(f_n(X^n)) = X^n\} \geq \mathbb{P}\{X^n \in S\} \geq 1 - \epsilon$. Therefore we conclude that $\mathsf{R}(\epsilon) \leq d$. The converse follows from Theorem 35.

Lastly, we show that

$$\mathsf{R}(0) = d = \log_M m \tag{5.157}$$

for the special case when P is equiprobable on its support, where $m = |\mathrm{supp}(P)|$. Recalling the construction of self-similar measures in Section 2.3, we first note that the distribution of X is a self-similar measure that is generated by the IFS $(S_0, \ldots, S_{M-1})$, where

$$S_i(x) = \frac{x+i}{M}, \quad i = 0, \ldots, M-1. \tag{5.158}$$

This IFS satisfies the open set condition, since $(0,1) \supset \bigcup_{i=0}^{M-1} S_i((0,1))$ and the union is disjoint. Denote by $E \subset [0,1]$ the invariant set of the reduced IFS $\{S_i : i \in \mathrm{supp}(P)\}$. By [25, Corollary 4.1], the distribution of X, denoted by μ_X, is in fact the normalized d-dimensional Hausdorff measure on E, i.e., $\mu_X(\cdot) = \mathcal{H}^d(\cdot \cap E)/\mathcal{H}^d(E)$. Therefore $\mu_{X^n}(\cdot) = \mathcal{H}^{dn}(\cdot \cap E^n)/\mathcal{H}^{dn}(E^n)$. By [18, Exercise 9.11], there exists a constant $c_1 > 0$, such that for all $x^n \in E^n$,

$$\underline{D}_{dn}(\mathcal{H}^{dn}, x^n) \geq c_1, \tag{5.159}$$

that is, E^n has positive lower dn-density everywhere. By [112, Theorem 4.1(1)], for any $k > dn$, there exists $F \subset E^n$ such that $\mathcal{H}^{dn}(F) = 0$ and $E^n \backslash F$ is k-rectifiable. Therefore $\mathbb{P}\{X^n \in F\} = 0$. By Lemma 19, the rate d is 0-achievable. □

Remark 19. The essential idea for constructing the Lipschitz mapping in [112, Theorem 4.1(1)] is the following: cover E^n and $[0,1]^k$ with δ-balls respectively, denoted by $\{B_1, \ldots, B_M\}$ and $\{\tilde{B}_1, \ldots, \tilde{B}_N\}$. Let g_1 be the mapping such that each B_i is mapped to the center of some $\tilde{B}_j$. By (5.159), g_1 can be shown to be Lipschitz. Then for each i and j cover $B_i \cap E^n$ and $\tilde{B}_j$ by δ^2-balls and repeat the same procedure to construct g_2. In this fashion, a sequence of Lipschitz mappings $\{g_m\}$ with *the same Lipschitz constant* is obtained, which converges to a Lipschitz mapping whose image is E^n.

5.5 Lossless linear compression with Lipschitz decompressors

In this section we prove the second half of Theorem 32, which shows that for discrete-continuous mixture with finite entropy, it is possible to achieve linear compression and Lipschitz decompression with bounded Lipschitz constants simultaneously. To this end, we need the following large-deviations result on Gaussian random matrices.

Lemma 23. *Let $\sigma_{\min}(\boldsymbol{B}_n)$ denote the smallest singular value of the $k_n \times m_n$ matrix $\boldsymbol{B}_n$ consisting of i.i.d. Gaussian entries with zero mean and variance $\frac{1}{n}$. For any $t > 0$, denote*

$$F_{k_n,m_n}(t) \triangleq \mathbb{P}\left\{\sigma_{\min}(\boldsymbol{B}_n) \le t\right\}. \tag{5.160}$$

Suppose that $\frac{m_n}{k_n} \xrightarrow{n\to\infty} \alpha \in (0,1)$ and $\frac{k_n}{n} \xrightarrow{n\to\infty} R > 0$. Then

$$\liminf_{n\to\infty} \frac{1}{n} \log \frac{1}{F_{k_n,m_n}(t)} \ge R(1-\alpha) \log \frac{1}{t\sqrt{\mathrm{e}R}} - \frac{Rh(\alpha)}{2}. \tag{5.161}$$

Proof. For brevity let $\boldsymbol{H}_n = \sqrt{n}\boldsymbol{B}_n$ and suppress the dependence of k_n and m_n on n. Then $\boldsymbol{H}_n^{\mathrm{T}}\boldsymbol{H}_n$ is an $l \times l$ Gaussian Wishart matrix. The minimum eigenvalue of $\boldsymbol{H}_n^{\mathrm{T}}\boldsymbol{H}_n$ has a density, which admits the following upper bound [113, p. 553].

$$f_{\lambda_{\min}(\boldsymbol{H}_n^{\mathrm{T}}\boldsymbol{H}_n)}(x) \le E_{k,m} x^{\frac{k-m-1}{2}} \mathrm{e}^{-\frac{x}{2}}, \quad x \ge 0, \tag{5.162}$$

where

$$E_{k,m} \triangleq \sqrt{\pi} 2^{-\frac{k-m+1}{2}} \frac{\Gamma\left(\frac{k+1}{2}\right)}{\Gamma\left(\frac{m}{2}\right)\Gamma\left(\frac{k-m+1}{2}\right)\Gamma\left(\frac{k-m+2}{2}\right)}. \tag{5.163}$$

Then

$$\mathbb{P}\left\{\sigma_{\min}(\boldsymbol{B}_n) \le t\right\} = \mathbb{P}\left\{\lambda_{\min}(\boldsymbol{H}_n^{\mathrm{T}}\boldsymbol{H}_n) \le nt^2\right\} \tag{5.164}$$

$$\le E_{k,m} \int_0^{nt^2} x^{\frac{k-m-1}{2}} \mathrm{e}^{-\frac{x}{2}} \mathrm{d}x \tag{5.165}$$

$$\le \frac{2n^{\frac{k-m+1}{2}} E_{k,m}}{k-m+1} t^{k-m+1}. \tag{5.166}$$

Applying Stirling's approximation to (5.166) yields (5.161). □

The next lemma provides a probabilistic lower bound on the operator norm of the inverse of a Gaussian random matrix restricted on an affine subspace. The point of this result is that the lower bound depends only on the dimension of the subspace.

Lemma 24. *Let $\boldsymbol{A}$ be a $k \times n$ random matrix with i.i.d. Gaussian entries with zero mean and variance $\frac{1}{n}$. Let $k > l$. Then for any l-dimensional affine subspace U of*

$\mathbb{R}^n$,

$$\mathbb{P}\left\{\min_{x\in U\setminus\{0\}}\frac{\|\boldsymbol{A}x\|}{\|x\|}\le t\right\}\ge F_{k,l+1}(t). \tag{5.167}$$

Proof. By definition, there exists $v\in\mathbb{R}^n$ and an m-dimensional linear subspace V such that $U=v+V$. First assume that $v\notin V$. Then $0\notin U$. Let $\{v_0,\dots,v_l\}$ be an orthonormal basis for $V'=\operatorname{span}(v,V)$. Set $\Psi=[v_0,\dots,v_l]$. Then

$$\min_{x\in U}\frac{\|\boldsymbol{A}x\|}{\|x\|}=\min_{x\in V'}\frac{\|\boldsymbol{A}x\|}{\|x\|} \tag{5.168}$$

$$=\min_{y\in\mathbb{R}^k}\frac{\|\boldsymbol{A}\Psi y\|}{\|y\|} \tag{5.169}$$

$$=\sigma_{\min}(\boldsymbol{A}\Psi). \tag{5.170}$$

Then (5.167) holds with equality since $\boldsymbol{A}\Psi$ is an $k\times(l+1)$ random matrix with i.i.d. normal entries of zero mean and variance $\frac{1}{n}$. If $v\in V$, then (5.167) holds with equality and $l+1$ replaced by l. The proof is then complete because $m\mapsto F_{k,m}(t)$ is decreasing. □

Lemma 25. *Let T be a union of N affine subspaces of $\mathbb{R}^n$ with dimension not exceeding l. Let $\mathbb{P}\{X^n\in T\}\ge 1-\epsilon$. Let $\boldsymbol{A}$ be defined in Lemma 24 independent of X^n. Then*

$$\mathbb{P}\left\{X^n\in T,\inf_{y\in T\setminus\{X\}}\frac{\|\boldsymbol{A}(y-X)\|}{\|y-X\|}\ge t\right\}\ge 1-\epsilon', \tag{5.171}$$

where

$$\epsilon'=\epsilon+NF_{k,l+1}(t). \tag{5.172}$$

Moreover, there exists a subset $E\subset\mathbb{R}^{k\times n}$ with $\mathbb{P}\{\boldsymbol{A}\in E\}\ge 1-\sqrt{\epsilon'}$, such that for any $\mathbf{K}\in E$, there exists a Lipschitz continuous function $g_{\mathbf{K}}:\mathbb{R}^k\to\mathbb{R}^n$ with $\operatorname{Lip}(g_{\mathbf{K}})\le t^{-1}$ and

$$\mathbb{P}\{g_{\mathbf{A}}(\mathbf{A}X^n)\ne X^n\}\ge 1-\sqrt{\epsilon'}. \tag{5.173}$$

Proof. By the independence of X^n and $\boldsymbol{A}$,

$$\mathbb{P}\left\{X^n\in T,\inf_{y\in T\setminus\{X\}}\frac{\|\boldsymbol{A}(y-X)\|}{\|y-X\|}\ge t\right\}=\int_T P_{X^n}(\mathrm{d}x)\mathbb{P}\left\{\inf_{z\in(T-x)\setminus\{0\}}\frac{\|\boldsymbol{A}z\|}{\|z\|}\ge t\right\} \tag{5.174}$$

$$\ge\mathbb{P}\{X^n\in T\}(1-NF_{k,l+1}) \tag{5.175}$$

$$\ge 1-\epsilon'. \tag{5.176}$$

where (5.175) follows by applying Lemma 24 to each affine subspace in $T-x$ and the union bound. To prove (5.173), denote by $p(\mathbf{K})$ the probability in the left-hand side of (5.171) conditioned on the random matrix $\boldsymbol{A}$ being equal to $\mathbf{K}$. By Fubini's

theorem and Markov inequality,

$$\mathbb{P}\left\{p(\boldsymbol{A}) \geq 1-\sqrt{\epsilon'}\right\} \geq 1-\sqrt{\epsilon'}. \tag{5.177}$$

Put $E = \{\mathbf{K} : p(\mathbf{K}) \geq 1-\sqrt{\epsilon'}\}$. For each $\mathbf{K} \in E$, define

$$U_{\mathbf{K}} = \left\{x \in T : \inf_{y \in T \setminus \{x\}} \frac{\|\mathbf{K}(y-x)\|}{\|y-x\|} \geq t\right\} \subset T. \tag{5.178}$$

Then for any $x, y \in U_{\mathbf{K}}$, we have

$$\|\mathbf{K}(x-y)\| \geq t\,\|x-y\|\,, \tag{5.179}$$

which implies that $\mathbf{K}|_{U_{\mathbf{K}}}$, the linear mapping $\mathbf{K}$ restricted on the set $U_{\mathbf{K}}$, is injective. Moreover, its inverse $g_{\mathbf{K}} : \mathbf{K}(U_{\mathbf{K}}) \to U_{\mathbf{K}}$ is t^{-1}-Lipschitz. By Kirszbraun's theorem [104, 2.10.43], $g_{\mathbf{K}}$ can be extended to a Lipschitz function on the whole space $\mathbb{R}^k$ with the same Lipschitz constant. For those $\mathbf{K} \notin E$, set $g_{\mathbf{K}} \equiv 0$. Since $\mathbb{P}\{X^n \in U_{\mathbf{K}}\} \geq 1-\sqrt{\epsilon'}$ for all $\mathbf{K} \in E$, we have

$$\mathbb{P}\{g_{\mathbf{K}}(\mathbf{K}X^n) \neq X^n\} \geq \mathbb{P}\{X^n \in U_{\mathbf{A}}, \mathbf{A} \in E\} \geq 1-\sqrt{\epsilon'}. \tag{5.180}$$

This completes the proof of the lemma. □

Proof of the second half of Theorem 32. It remains to establish the achievability part of (5.19): $\hat{\mathsf{R}}(\epsilon) \leq \gamma$. Fix $\delta, \delta' > 0$ arbitrarily small. In view of Lemma 25, to prove the achievability of $R = \gamma + \delta + \delta'$, it suffices to show that, with high probability, X^n lies in the union of exponentially many affine subspaces whose dimensions do not exceed nR.

To this end, let $\mathsf{t}(x^n)$ denote the discrete part of x^n, i.e., the vector formed by those $x_i \in \mathcal{A}$ in increasing order of i. For $k \geq 1$, define

$$\mathsf{T}_k = \left\{z^k \in \mathcal{A}^k : \frac{1}{k}\sum_{i=1}^{k} \log \frac{1}{P_{\mathrm{d}}(z_i)} \leq H(P_{\mathrm{d}}) + \delta'\right\}. \tag{5.181}$$

Since $H(P_{\mathrm{d}}) < \infty$, we have $|\mathsf{T}_k| \leq \exp((H(P_{\mathrm{d}}) + \delta')k)$. Moreover, $P_{\mathrm{d}}^k(\mathsf{T}_k) \geq 1-\epsilon$ for all sufficiently large k.

Recall the generalized support $\mathrm{spt}(x^n)$ of x^n defined in (3.73). Let

$$C_n = \left\{x^n \in \mathbb{R}^n : \left||\mathrm{spt}(x^n)| - \gamma n\right| \leq \delta n,\ \mathsf{t}(x^n) \in \mathsf{T}_{n-|\mathrm{spt}(x^n)|}\right\} \tag{5.182}$$

$$= \bigcup_{\substack{S \subset \{1,\ldots,n\} \\ ||S|-\gamma n| \leq \delta n}} \bigcup_{z \in \mathsf{T}_{n-|S|}} \{x^n \in \mathbb{R}^n : \mathrm{spt}(x^n) = S,\ \mathsf{t}(x^n) = z\}. \tag{5.183}$$

Note that each of the subsets in the right-hand side of (5.183) is an affine subspace of dimension no more than $(\gamma+\delta)n$. Therefore C_n consists of N_n affine subspaces,

with $N_n \leq \sum_{|k-\gamma n|\leq \delta n} |\mathsf{T}_{n-k}|$, hence

$$\limsup_{n\to\infty} \frac{1}{n} \log N_n \leq (H(P_{\mathrm{d}}) + \delta')(1-\gamma-\delta)n. \tag{5.184}$$

Moreover, for sufficiently large n, we have

$$\begin{aligned}
\mathbb{P}\{X^n \in C_n\} &= \sum_{||S|-\gamma n|\leq \delta n} \mathbb{P}\{X^n \in C_n, \mathrm{spt}(X^n) = S\} && (5.185)\\
&= \sum_{||S|-\gamma n|\leq \delta n} \mathbb{P}\{\mathrm{spt}(X^n) = S\} P_{\mathrm{d}}^{n-|S|}(\mathsf{T}_{n-|S|}) && (5.186)\\
&\geq \mathbb{P}\left\{\left||\mathrm{spt}(X^n)| - \gamma n\right| \leq \delta n\right\}(1-\epsilon) && (5.187)\\
&\geq 1 - 2\epsilon. && (5.188)
\end{aligned}$$

where (5.187) is due to (3.72). To apply Lemma 24, it remains to select a t sufficiently small but fixed, such that $N_n F_{k,m+1}(t) = o(1)$ as $n \to \infty$. This is always possible, in view of (5.184) and Lemma 23, as long as

$$(1-\alpha)R \log \frac{1}{t\sqrt{\mathrm{e}R}} - \frac{Rh(\alpha)}{2} > (H(P_{\mathrm{d}}) + \delta')(1-\gamma-\delta), \tag{5.189}$$

where $\alpha = \frac{\gamma+\delta}{R}$ and $R = \gamma + \delta + \delta'$. By the arbitrariness of δ and δ', the proof of $\hat{\mathsf{R}}(\epsilon) \leq \gamma$ is complete. □

5.6 Lossless stable decompression

First we present a key result used in the converse proof of Theorem 38.

Lemma 26. *Let $S \subset \mathbb{R}^n$ and $T \subset [0,1]^k$. Let $g : T \to S$ be bijective and 2^{-m}-stable with respect to the ℓ_∞ distance. Then for any $\delta > 0$, there exists $q : [S]_m \to [0,1]^l$, where $l = k + \lceil \delta n \rceil$ such that*

$$\|q(x) - q(y)\|_\infty \geq \min\{2^{-m}, 2^{-\frac{1}{\delta}}\} \tag{5.190}$$

holds for all distinct x and y.

Proof. See Appendix C.6. □

The proof of Lemma 26 uses a graph coloring argument and the following result from graph theory.

Lemma 27 (Greedy Graph Coloring, [114, Theorem 13.1.5])**.** *Let G be a graph. Denote by $\Delta(\mathsf{G})$ the maximum degree of vertices in G and by $\chi(\mathsf{G})$ the coloring number of G. Then the greedy algorithm produces a $(\Delta(\mathsf{G})+1)$-coloring of the vertices of G and hence*

$$\chi(\mathsf{G}) \leq \Delta(\mathsf{G}) + 1. \tag{5.191}$$

Now we prove the converse coding theorem for stable analog compression.

Proof of Theorem 38 (Converse). Fix $0 \leq \epsilon < 1$, $\delta > 0$ and $m \in \mathbb{N}$ arbitrarily. Let $R = \overline{\mathsf{R}}(\epsilon, 2^{-m}) + \delta$ and $k = \lfloor Rn \rfloor$. By definition, there exist $f_n : \mathbb{R}^n \to [0,1]^k$ and $g_n : [0,1]^k \to \mathbb{R}^n$, such that g_n is 2^{-m}-stable and

$$\mathbb{P}\left\{g_n(f_n(X^n)) \neq X^n\right\} \leq \epsilon. \tag{5.192}$$

Let $S^n = \{x^n \in \mathbb{R}^n : g_n(f_n(x^n)) = x^n\}$, then

$$\begin{aligned} 1 - \epsilon &\leq \mathbb{P}\left\{X^n \in S^n\right\} && (5.193)\\ &\leq \mathbb{P}\left\{[X^n]_m \in [S^n]_m\right\}. && (5.194) \end{aligned}$$

holds for any n. Since X^n is an i.i.d. random vector, $[X^n]_m$ is a random vector consisting of n i.i.d. copies of the source $[X]_m$. By the strong converse of the lossless source coding theorem for discrete memoryless sources, we have

$$\liminf_{n\to\infty} \frac{|[S^n]_m|}{n} \geq H([X]_m). \tag{5.195}$$

Next we upper bound the cardinality of $[S^n]_m$ through a volume argument. Since $g_n : f(S^n) \to S^n$ is bijective and 2^{-m}-stable, by Lemma 26, there exists a mapping $q : [S^n]_m \to [0,1]^l$, where

$$l = k + \lceil \delta n \rceil, \tag{5.196}$$

such that

$$\|q(x) - q(y)\|_\infty \geq \min\{2^{-m}, 2^{-\frac{1}{\delta}}\}, \tag{5.197}$$

which implies that points in the image set $q([S^n]_m)$ are separated by at least 2^{-m} for all $m > \frac{1}{\delta}$. Therefore $B_\infty\left(q(x), 2^{-(m+1)}\right)$ are *disjoint* for all $x \in [S^n]_m$. Since $B_\infty(z, r)$ is just an l-dimensional cube of side length $2r$, we have

$$\mathrm{Leb}\left(B_\infty^l\left(q(x), 2^{-(m+1)}\right)\right) = 2^{-ml}. \tag{5.198}$$

Therefore

$$|[S^n]_m| \leq \frac{\mathrm{Leb}\left([0,1]^k\right)}{2^{-ml}} = 2^{ml}, \tag{5.199}$$

hence

$$\begin{aligned} \frac{H([X]_m)}{m} &\leq \frac{1}{m} \liminf_{n\to\infty} \frac{\log |[S^n]_m|}{n} && (5.200)\\ &\leq \frac{1}{m} \limsup_{n\to\infty} \frac{\log |[S^n]_m|}{n} && (5.201)\\ &\leq R + \delta && (5.202)\\ &= \overline{\mathsf{R}}(\epsilon, 2^{-m}) + 2\delta, && (5.203) \end{aligned}$$

holds for all $m > \frac{1}{\delta}$, where (5.200) is due to (5.195). Recalling Lemma 1, we have

$$\begin{aligned}\overline{d}(X) &= \limsup_{m\to\infty} \frac{H([X]_m)}{m} && (5.204)\\ &\le \limsup_{m\to\infty} \overline{\mathsf{R}}(\epsilon, 2^{-m}) && (5.205)\\ &\le \limsup_{\Delta\to\infty} \overline{\mathsf{R}}(\epsilon, \Delta). && (5.206)\end{aligned}$$

where (5.205) is because of (5.203) and the arbitrariness of δ.

□

The achievability of stable compression is shown through an explicit construction.

Proof of Theorem 38 (Achievability). Fix an arbitrary $\Delta > 0$. There exists $m \in \mathbb{N}$, such that

$$2^{-(m+1)} \le \Delta < 2^{-m}. \tag{5.207}$$

Consider the memoryless source $\{[X_i]_{m+1} : i \in \mathbb{N}\}$ with entropy $H([X]_{m+1})$. for all $\delta > 0$, there exists $G^n \subset 2^{-(m+1)}\mathbb{Z}^n$ such that

$$\begin{aligned}|G^n| &\le 2^{n[H([X]_{m+1})+\delta]} && (5.208)\\ \lim_{n\to\infty} \mathbb{P}\left\{[X^n]_{m+1} \in G^n\right\} &= 1. && (5.209)\end{aligned}$$

Define

$$S^n = \left\{x^n \in \mathbb{R}^n : [x^n]_{m+1} \in G\right\}. \tag{5.210}$$

Then

$$\lim_{n\to\infty} \mathbb{P}\left\{X^n \in S^n\right\} = 1. \tag{5.211}$$

Note that S^n is a disjoint union of at most $|G^n|$ hypercubes in $\mathbb{R}^n$. Let $k \in \mathbb{N}$. We construct $f_n : \mathbb{R}^n \to [0,1]^k$ as follows. Uniformly divide $[0,1]^k$ into 2^{km} hypercubes. Choose $2^{k(m-1)}$ cubes out of them as the codebook, such that any pair of cubes are separated by 2^{-m} in ℓ_∞ distance. Denote the codebook by C_n. In view of (5.208), letting

$$\frac{k}{n} = \frac{H([X]_{m+1}) + \delta}{m-1}. \tag{5.212}$$

is sufficient to ensure that the number of cubes in C_n is no less than $|G^n|$. Construct f_n by mapping each hypercube in C_n to one in S_n by an arbitrary bijective mapping. Denote the inverse of f_n by g_n. Now we proceed to show that g_n is Δ-stable. Consider any $x, y \in C_n$ such that $\|x - y\|_\infty \le \Delta$. Since $\Delta < 2^{-m}$, x and y reside in the same hypercube in C_n, so will $g(x)$ and $g(y)$ be in S^n. Therefore $\|g(x) - g(y)\|_\infty \le 2^{-(m+1)} \le \Delta$. In view of (5.212), the Δ-stability of g_n implies that

$$\overline{\mathsf{R}}(\epsilon, \Delta) \le \frac{H([X]_{m+1}) + \delta}{m-1} \tag{5.213}$$

holds for all $0 < \epsilon < 1$, where $m = \left\lfloor \log_2 \frac{1}{\Delta} \right\rfloor$ by (5.207). Sending $\Delta \to 0$, we have

$$\limsup_{\Delta \downarrow 0} \overline{\mathsf{R}}(\epsilon, \Delta) \leq \limsup_{m \to \infty} \frac{H([X]_{m+1}) + \delta}{m - 1} = \overline{d}(X) \tag{5.214}$$

for all $0 < \epsilon < 1$. $\square$

Chapter 6

Noisy compressed sensing

This chapter extends the analog compression framework in Chapter 5 to the case where the encoder output is contaminated by additive noise. Noise sensitivity, the ratio between the reconstruction error and the noise variance, is used to gauge the robustness of the decoding procedure. Focusing on optimal decoders (the MMSE estimator), in Section 6.2 we formulate the fundamental tradeoff between measurement rate and reconstruction fidelity under three classes of encoders, namely optimal nonlinear, optimal linear and random linear encoders. In Section 6.5, optimal phase transition thresholds of noise sensitivity are determined as a functional of the input distribution. In particular, we show that for discrete-continuous mixtures, Gaussian sensing matrices incur no penalty on the phase transition threshold with respect to optimal nonlinear encoding. Our results also provide a rigorous justification of previous results based on replica heuristics [7] in the weak-noise regime. The material in this chapter has been presented in part in [115, 116].

6.1 Setup

The basic setup of noisy compressed sensing is a *joint source-channel coding* problem as shown in Fig. 6.1, where we assume that

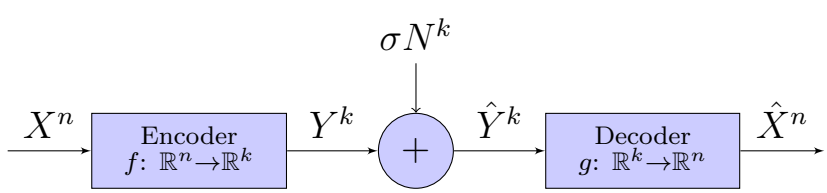

Figure 6.1: Noisy compressed sensing setup.

- The source is memoryless, where X^n consists of i.i.d. copies of a real-valued random variable X with unit variance.
- The channel is memoryless with additive Gaussian noise σN^k where $N^k \sim \mathcal{N}(0, \mathbf{I}_k)$.

- The encoder satisfies a unit average power constraint:

$$\frac{1}{k}\mathbb{E}[\|f(X^n)\|_2^2] \leq 1. \tag{6.1}$$

- The reconstruction error is gauged by the per-symbol MSE distortion:

$$d(x^n, \hat{x}^n) = \frac{1}{n}\|\hat{x}^n - x^n\|_2^2. \tag{6.2}$$

In this setup, the fundamental question is: For a given noise variance and measurement rate, what is the smallest possible reconstruction error? For a given encoder f, the corresponding optimal decoder g is the *MMSE estimator* of the input X^n given the channel output $\hat{Y}^k = f(X^n) + \sigma N^k$. Therefore the optimal distortion achieved by encoder f is

$$\mathsf{mmse}(X^n | f(X^n) + \sigma N^k). \tag{6.3}$$

6.2 Distortion-rate tradeoff

For a fixed noise variance, we define three distortion-rate functions that corresponds to *optimal* encoding, *optimal linear* encoding and *random linear* encoding respectively.

6.2.1 Optimal encoder

Definition 20. The minimal distortion achieved by the optimal encoding scheme is given by:

$$D^*(X, R, \sigma^2) \triangleq \limsup_{n\to\infty} \frac{1}{n} \inf_f \left\{ \mathsf{mmse}(X^n | f(X^n) + \sigma N^k) : \mathbb{E}[\|f(X^n)\|_2^2] \leq Rn \right\}. \tag{6.4}$$

The asymptotic optimization problem in (6.4) can be solved by applying Shannon's joint source-channel coding separation theorem for memoryless channels and sources [90], which states that the lowest rate that achieves distortion D is given by

$$R = \frac{R_X(D)}{C(\sigma^2)}, \tag{6.5}$$

where $R_X(\cdot)$ is the rate-distortion function of X defined in (2.64) and $C(\sigma^2) = \frac{1}{2}\log(1 + \sigma^{-2})$ is the AWGN channel capacity. By the monotonicity of the rate-distortion function, we have

$$D^*(X, R, \sigma^2) = R_X^{-1}\left(\frac{R}{2}\log(1 + \sigma^{-2})\right). \tag{6.6}$$

In general, optimal encoders are nonlinear. In fact, Shannon's separation theorem states that the composition of an optimal lossy source encoder and an optimal channel encoder is asymptotically optimal when blocklength $n \to \infty$. Such a construction results in an encoder that is *finitely-valued*, hence nonlinear. For fixed n and k, linear encoders are in general suboptimal.

6.2.2 Optimal linear encoder

To analyze the fundamental limit of noisy compressed sensing, we restrict the encoder f to be a linear mapping, denoted by a matrix $\mathbf{H} \in \mathbb{R}^{k \times n}$. Since X^n are i.i.d. with zero mean and unit variance, the input power constraint (6.1) simplifies to

$$\mathbb{E}[\|\mathbf{H}X^n\|_2^2] = \mathbb{E}[X^{n\mathrm{T}}\mathbf{H}^{\mathrm{T}}\mathbf{H}X^n] = \mathsf{Tr}(\mathbf{H}^{\mathrm{T}}\mathbf{H}) = \|\mathbf{H}\|_{\mathrm{F}}^2 \leq k, \tag{6.7}$$

where $\|\cdot\|_{\mathrm{F}}$ denotes the Frobenius norm.

Definition 21. Define the optimal distortion achievable by linear encoders as:

$$D_{\mathrm{L}}^*(X, R, \sigma^2) \triangleq \limsup_{n\to\infty} \frac{1}{n} \inf_{\mathbf{H}} \left\{ \mathsf{mmse}(X^n | \mathbf{H}X^n + \sigma N^k) : \|\mathbf{H}\|_{\mathrm{F}}^2 \leq Rn \right\}. \tag{6.8}$$

A general formula for $D_{\mathrm{L}}^*(X, R, \sigma^2)$ is yet unknown. One example where we compute it explicitly in Section 6.4 is the Gaussian source.

6.2.3 Random linear encoder

We consider the ensemble performance of random linear encoders and relax the power constraint in (6.7) to hold on average:

$$\mathbb{E}\left[\|\boldsymbol{A}\|_{\mathrm{F}}^2\right] \leq k. \tag{6.9}$$

In particular, we focus on the following ensemble of random sensing matrices, for which (6.9) holds with equality.:

Definition 22. Let $\boldsymbol{A}_n$ be a $k \times n$ random matrix with i.i.d. entries of zero mean and variance $\frac{1}{n}$. The minimal expected distortion achieved by this ensemble of linear encoders is given by:

$$D_{\mathrm{L}}(X, R, \sigma^2) \triangleq \limsup_{n\to\infty} \frac{1}{n} \mathsf{mmse}(X^n | (\boldsymbol{A}X^n + \sigma N^k, \boldsymbol{A})) \tag{6.10}$$

$$= \limsup_{n\to\infty} \mathsf{mmse}(X_1 | (\boldsymbol{A}X^n + \sigma N^k, \boldsymbol{A})) \tag{6.11}$$

where (6.11) follows from symmetry and $\mathsf{mmse}(\cdot|\cdot)$ is defined in (4.2).

Based on the statistical-physics approach in [117, 118], the decoupling principle results in [118] were imported into the compressed sensing setting in [7] to *postulate*

the following formula for $D_{\mathrm{L}}(X, R, \sigma^2)$. Note that this result is based replica heuristics lacking a rigorous justification.

Claim 1 ([7, Corollary 1, p.5])**.**

$$D_{\mathrm{L}}(X, R, \sigma^2) = \mathsf{mmse}(X, \eta R\, \sigma^{-2}), \tag{6.12}$$

where $0 < \eta < 1$ *satisfies the following equation [7, (12) – (13), pp. 4 – 5]:*[1]

$$\frac{1}{\eta} = 1 + \frac{1}{\sigma^2}\mathsf{mmse}(X, \eta R\, \sigma^{-2}). \tag{6.13}$$

When (6.13) has more than one solution, η *is chosen to be the one that minimizes the free energy*

$$I(X; \sqrt{\eta R \sigma^{-2}} X + N) + \frac{R}{2}(\eta - 1 - \log \eta). \tag{6.14}$$

Note that $D_{\mathrm{L}}(X, R, \sigma^2)$ defined in Definition 22 depends implicitly on the entry-wise distribution of the random measurement matrix $\boldsymbol{A}$, while Claim 1 asserts that the the formula (6.12) for $D_{\mathrm{L}}(X, R, \sigma^2)$ holds *universally*, as long as its entries are i.i.d. with variance $\frac{1}{n}$. Consequently, it is possible to employ random sparse measurement matrix so that each encoding operation involves only a few signal components, for example,

$$A_{ij} \sim \frac{p}{2}\delta_{\frac{-1}{\sqrt{pn}}} + (1-p)\delta_0 + \frac{p}{2}\delta_{\frac{1}{\sqrt{pn}}} \tag{6.15}$$

for some $0 < p < 1$.

Remark 20. The solutions to (6.13) are precisely the stationary points of the free energy (6.14) as a function of η. In fact it is possible for (6.12) to have arbitrarily many solutions. To see this, let $\beta = \eta R\, \sigma^{-2}$. Then (6.12) becomes

$$\beta \mathsf{mmse}(X, \beta) = R - \sigma^2 \beta. \tag{6.16}$$

Consider a Cantor distributed X, for which $\beta\mathsf{mmse}(X, \beta)$ oscillates logarithmically in β between the lower and upper MMSE dimension of X (c.f. Fig. 4.1). In Fig. 6.2 we plot both sides of (6.16) against $\log_3 \beta$. We see that as σ^2 decreases, $\beta \mapsto R - \sigma^2\beta$ becomes flatter and number of solutions blows up.

6.2.4 Properties

Theorem 45. *1. For fixed* σ^2*,* $D^*(X, R, \sigma^2)$ *and* $D^*_{\mathrm{L}}(X, R, \sigma^2)$ *are both decreasing, convex and continuous in* R *on* $(0, \infty)$*.*

2. For fixed R*,* $D^*(X, R, \sigma^2)$ *and* $D^*_{\mathrm{L}}(X, R, \sigma^2)$ *are both increasing, convex and continuous in* $\frac{1}{\sigma^2}$ *on* $(0, \infty)$*.*

[1] In the notation of [7, (12)], γ and $\epsilon\mu$ correspond to $R\sigma^{-2}$ and R in our formulation.

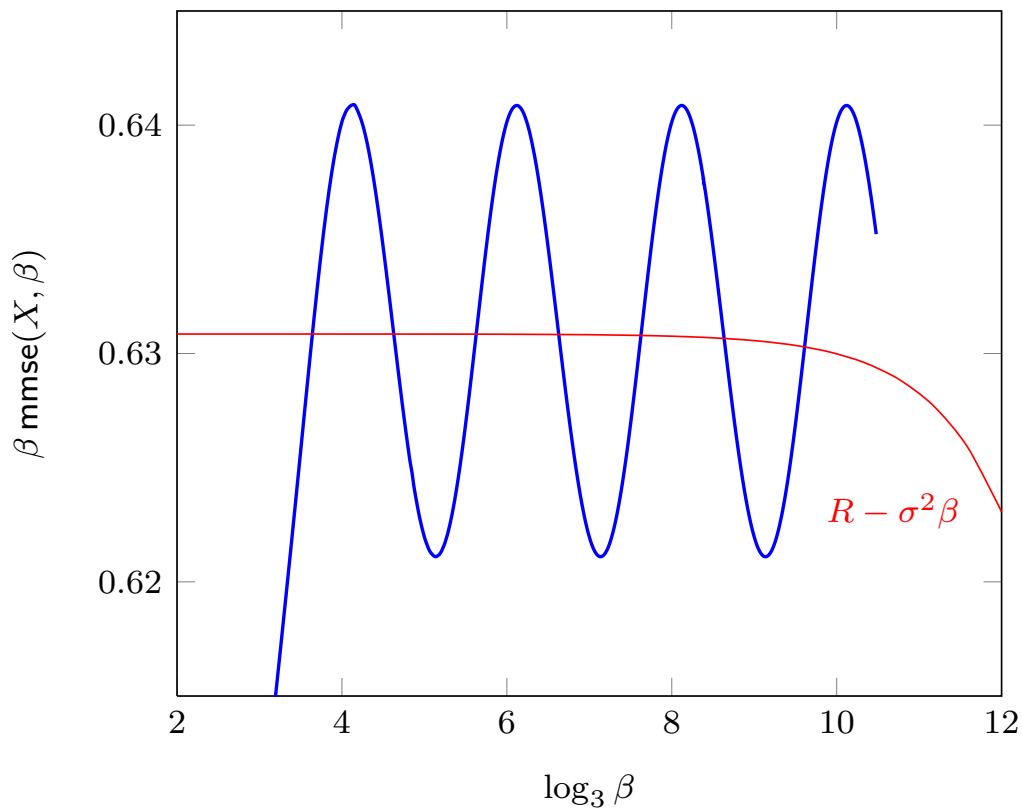

Figure 6.2: For sufficiently small σ^2, there are arbitrarily many solutions to (6.16).

3.

$$D^*(X, R, \sigma^2) \leq D^*_{\mathrm{L}}(X, R, \sigma^2) \leq D_{\mathrm{L}}(X, R, \sigma^2) \leq 1. \tag{6.17}$$

Proof. 1. Fix σ^2. Monotonicity with respect to the measurement rate R is straightforward from the definition of D^* and D^*_{L}. Convexity follows from time-sharing between two encoding schemes. Finally, convexity on the real line implies continuity.

2. Fix R. For any n and any encoder $f : \mathbb{R}^n \to \mathbb{R}^k$, $\sigma^2 \mapsto \mathsf{mmse}(X^n | f(X^n) + \sigma N^k)$ is increasing. This is a result of the infinite divisibility of the Gaussian distribution as well as the data processing lemma of MMSE [87]. Consequently, $\sigma^2 \mapsto D^*(X, R, \sigma^2)$ is also increasing. Since D^* can be equivalently defined as

$$D^*(X, R, \sigma^2) = \limsup_{n\to\infty} \frac{1}{n} \inf_f \left\{ \mathsf{mmse}(X^n | f(X^n) + \sigma N^k) : \mathbb{E}[\|f(X^n)\|_2^2] \leq \frac{k}{\sigma^2} \right\}, \tag{6.18}$$

convexity in $\frac{1}{\sigma^2}$ follows from time-sharing. The results on D^*_{L} follows analogously.

3. The leftmost inequality in (6.17) follows directly from the definition, while the rightmost inequality follows because we can always discard the measurements

and use the mean as an estimate. Although the best sensing matrix will beat the average behavior of any ensemble, the middle inequality in (6.17) is not quite trivial because the power constraint in (6.1) is not imposed on each matrix in the ensemble. We show that for any fixed $\epsilon > 0$,

$$D_{\mathrm{L}}^*(X, R, \sigma^2) \leq D_{\mathrm{L}}(X, R, (1+\epsilon)^2\sigma^2). \tag{6.19}$$

By the continuity of $\sigma^{-2} \mapsto D_{\mathrm{L}}^*(X, R, \sigma^2)$ proved in Theorem 45, $\sigma^2 \mapsto D_{\mathrm{L}}^*(X, R, \sigma^2)$ is also continuous. Therefore sending $\epsilon \downarrow 0$ in (6.19) yields the second inequality in (6.17). To show (6.19), recall that $\boldsymbol{A}$ consists of i.i.d. entries with zero mean and variance $\frac{1}{n}$. Since $k = nR$, $\frac{1}{k}\|\boldsymbol{A}\|_{\mathrm{F}}^2 \xrightarrow{\mathbb{P}} 1$ as $n \to \infty$, by the weak law of large numbers. Therefore $\mathbb{P}\{\boldsymbol{A} \in E_n\} \to 1$ where

$$E_n \triangleq \left\{\mathbf{A} : \|\mathbf{A}\|_{\mathrm{F}}^2 \leq k(1+\epsilon)^2\right\}. \tag{6.20}$$

Therefore

$$\begin{aligned}
& D_{\mathrm{L}}(X, R, (1+\epsilon)\sigma^2) \\
= & \limsup_{n\to\infty} \frac{1}{n}\mathsf{mmse}\left(X^n | \boldsymbol{A}X^n + (1+\epsilon)^2\sigma^2 N^k, \boldsymbol{A}\right) && (6.21)\\
= & \limsup_{n\to\infty} \frac{1}{n}\mathsf{mmse}\left(X^n \middle| \frac{\boldsymbol{A}}{1+\epsilon}X^n + N^k, \boldsymbol{A}\right) && (6.22)\\
\geq & \limsup_{n\to\infty} \frac{\mathbb{P}\{\boldsymbol{A} \in E_n\}}{n}\mathbb{E}\left[\mathsf{mmse}\left(X^n \middle| \frac{\boldsymbol{A}}{1+\epsilon}X^n + N^k, \boldsymbol{A}\right) \middle| \boldsymbol{A} \in E_n\right] && (6.23)\\
= & D_{\mathrm{L}}^*(X, R, \sigma^2), && (6.24)
\end{aligned}$$

where (6.23) holds because $\frac{\mathbf{A}}{1+\epsilon}$ satisfies the power constraint for any $\mathbf{A} \in E_n$. □

Remark 21. Alternatively, the convexity properties of D^* can be derived from (6.6). Since $R_X(\cdot)$ is decreasing and concave, $R_X^{-1}(\cdot)$ is decreasing and convex, which, composed with the concave mapping $\sigma^{-2} \mapsto \frac{R}{2}\log(1+\sigma^{-2})$, gives a convex function $\sigma^{-2} \mapsto D^*(X, R, \sigma^2)$ [119, p. 84]. The convexity of $R \mapsto D^*(X, R, \sigma^2)$ can be similarly proved.

Remark 22. Note that the time-sharing proofs of Theorem 45.1 and 45.2 do not work for D_{L}, because time-sharing between two random linear encoders results in a block-diagonal matrix with diagonal submatrices each filled with i.i.d. entries. This ensemble is outside the scope of the random matrices with i.i.d. entries consider in Definition 22. Therefore, proving convexity of $R \mapsto D_{\mathrm{L}}(X, R, \sigma^2)$ amounts to showing that the replacing all zeroes in the block-diagonal matrix with independent entries of the same distribution always helps with the estimation. This is certainly not true for individual matrices.

6.3 Phase transition of noise sensitivity

The main objective of noisy compressed sensing is the robust recovery of the input vector, obtaining a reconstruction error that is proportional to the noise variance. To quantify robustness, we analyze *noise sensitivity*: the ratio between the mean-square error and the noise variance, at a given R and σ^2. As a succinct characterization of robustness, we focus particular attention on the worst-case noise sensitivity:

Definition 23. The *worst-case noise sensitivity* of optimal encoding is defined as

$$\zeta^*(X, R) = \sup_{\sigma^2>0} \frac{D^*(X, R, \sigma^2)}{\sigma^2}. \tag{6.25}$$

For linear encoding, $\zeta_{\mathrm{L}}^*(X, R)$ and $\zeta_{\mathrm{L}}(X, R)$ are analogously defined with D^* in (6.25) replaced by D_{L}^* and D_{L}.

The phase transition threshold of the noise sensitivity is defined as the minimal measurement rate R such that the noise sensitivity is bounded for all noise variance [6, 115]:

Definition 24. Define

$$\mathcal{R}^*(X) \triangleq \inf \{R > 0 : \zeta^*(X, R) < \infty\}. \tag{6.26}$$

For linear encoding, $\mathcal{R}_{\mathrm{L}}^*(X)$ and $\mathcal{R}_{\mathrm{L}}(X)$ are analogously defined with ζ^* in (6.26) replaced by ζ_{L}^* and ζ_{L}.

By (6.17), the phase transition thresholds in Definition 24 are ordered naturally as

$$0 \le \mathcal{R}^*(X) \le \mathcal{R}_{\mathrm{L}}^*(X) \le \mathcal{R}_{\mathrm{L}}(X) \le 1, \tag{6.27}$$

where the rightmost inequality is shown below (after Theorem 46).

Remark 23. In view of the convexity properties in Theorem 45.2, the three worst-case sensitivities in Definition 23 are all (extended real-valued) convex functions of R.

Remark 24. Alternatively, we can consider the *asymptotic noise sensitivity* by replacing the supremum in (6.25) with the limit as $\sigma^2 \to 0$, denoted by $\xi^*, \xi_{\mathrm{L}}^*$ and ξ_{L} respectively. Asymptotic noise sensitivity characterizes the convergence rate of the reconstruction error as the noise variance vanishes. Since $D^*(X, R, \sigma^2)$ is always bounded above by $\mathsf{var}X = 1$, we have

$$\zeta^*(X, R) < \infty \Leftrightarrow \xi^*(X, R) < \infty. \tag{6.28}$$

Therefore $\mathcal{R}^*(X)$ can be equivalently defined as the infimum of all rates $R > 0$, such that

$$D^*(X, R, \sigma^2) = O\left(\sigma^2\right), \quad \sigma^2 \to 0. \tag{6.29}$$

This equivalence also applies to D_{L}^* and D_{L}. It should be noted that although finite worst-case noise sensitivity is equivalent to finite asymptotic noise sensitivity, the supremum in (6.28) need not be achieved as $\sigma^2 \to 0$. An example is given by the Gaussian input analyzed in Section 6.4.

6.4 Least-favorable input: Gaussian distribution

In this section we compute the distortion-rate tradeoffs for the Gaussian input distribution, which simultaneously maximizes all three distortion-rate functions subject to the variance constraint.

Theorem 46. *Let $X_{\mathsf{G}} \sim \mathcal{N}(0,1)$. Then for any R, σ^2 and X of unit variance,*

$$D^*(X, R, \sigma^2) \le D^*(X_{\mathsf{G}}, R, \sigma^2) = \frac{1}{(1+\sigma^{-2})^R}. \tag{6.30}$$

$$D_{\mathrm{L}}^*(X, R, \sigma^2) \le D_{\mathrm{L}}^*(X_{\mathsf{G}}, R, \sigma^2) = \begin{cases} 1 - \dfrac{R}{1+\sigma^2} & R \le 1, \\ \dfrac{1}{1+R\,\sigma^{-2}} & R > 1, \end{cases} \tag{6.31}$$

$$D_{\mathrm{L}}(X, R, \sigma^2) \le D_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2) = \frac{1}{2}\left(1 - R - \sigma^2 + \sqrt{(1-R)^2 + 2(1+R)\sigma^2 + \sigma^4}\right) \tag{6.32}$$

Proof. Since the Gaussian distribution maximizes the rate-distortion function pointwise under the variance constraint [120, Theorem 4.3.3], the inequality in (6.30) follows from (6.6). For linear encoding, linear estimators are optimal for Gaussian inputs since the channel output and the input are jointly Gaussian, but suboptimal for non-Gaussian inputs. Moreover, the linear MMSE depends only on the input variance. Therefore the inequalities in (6.31) and (6.32) follow. The distortion-rate functions of X_{G} are computed in Appendix D. □

Remark 25. While Theorem 46 asserts that Gaussian input is the least favorable among distributions with a given variance, it is not clear what the *least favorable sparse* input distribution is, i.e., the maximizer of D^* (D_{L}^* or D_{L}) in the set $\{P_X : \mathbb{P}\{X = 0\} \ge \gamma\}$.

The Gaussian distortion-rate tradeoffs in (6.30) – (6.32) are plotted in Figs. 6.3 and 6.4. We see that linear encoders are optimal for lossy encoding of Gaussian sources in Gaussian channels if and only if $R = 1$, i.e.,

$$D^*(X_{\mathsf{G}}, 1, \sigma^2) = D_{\mathrm{L}}^*(X_{\mathsf{G}}, 1, \sigma^2), \tag{6.33}$$

which is a well-known fact [121, 122]. As a result of (6.32), the rightmost inequality in (6.27) follows.

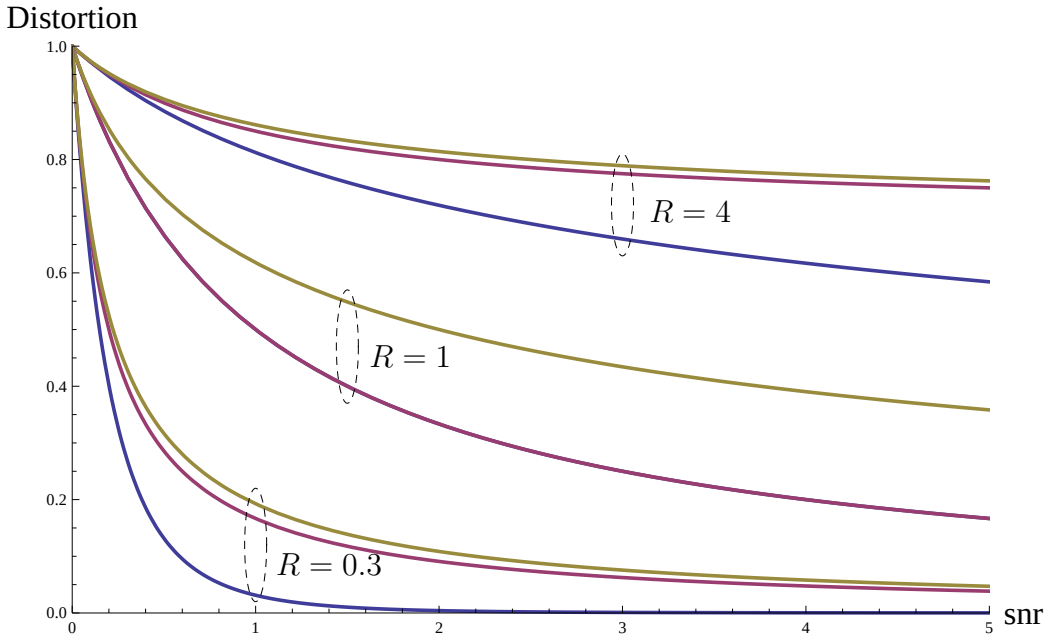

Figure 6.3: $D^*(X_{\mathsf{G}}, R, \sigma^2), D^*_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2), D_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2)$ against $\mathsf{snr} = \sigma^{-2}$.

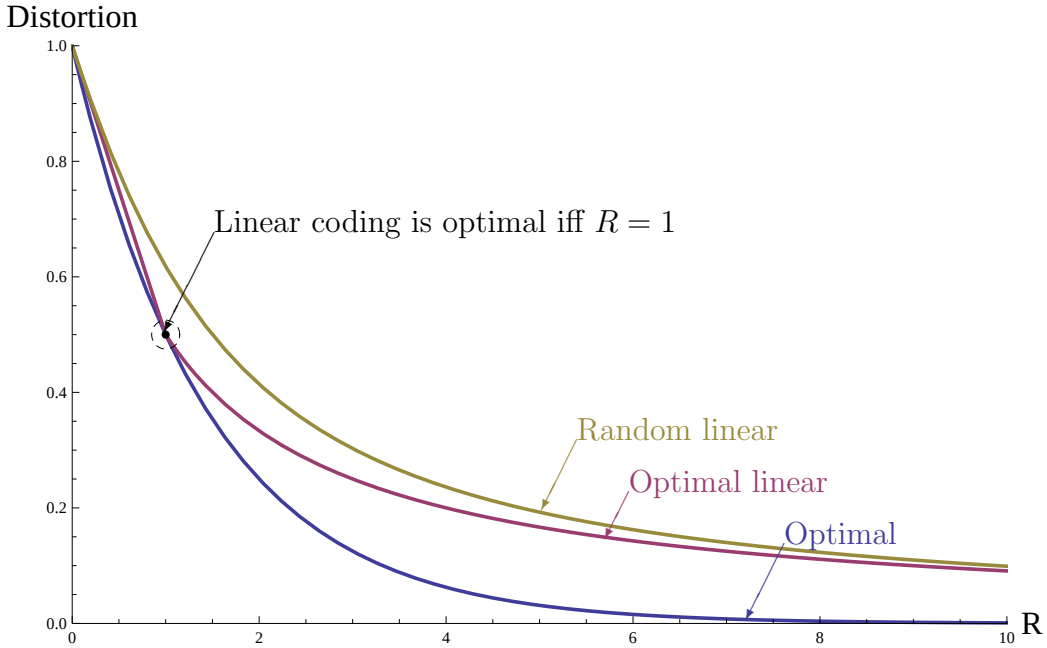

Figure 6.4: $D^*(X_{\mathsf{G}}, R, \sigma^2), D^*_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2), D_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2)$ against R when $\sigma^2 = 1$.

Next we analyze the high-SNR asymptotics of (6.30) – (6.32). Clearly, $D^*(X_{\mathsf{G}}, R, \sigma^2)$ is the smallest among the three, which vanishes polynomially in σ^2 according to

$$D^*(X_{\mathsf{G}}, R, \sigma^2) = \sigma^{2R} + O(\sigma^{2R+2}), \quad \sigma^2 \to 0 \tag{6.34}$$

regardless of how small $R > 0$ is. For linear encoding, we have

$$D_{\mathrm{L}}^*(X, R, \sigma^2) = \begin{cases} 1 - R + R\sigma^2 + O(\sigma^4) & 0 \le R < 1, \\ \sigma^2 + O(\sigma^4) & R = 1, \\ \frac{\sigma^2}{R} + O(\sigma^4) & R > 1. \end{cases} \tag{6.35}$$

$$D_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2) = \begin{cases} 1 - R + \frac{R}{1-R}\sigma^2 + O(\sigma^4) & 0 \le R < 1, \\ \sigma - \frac{\sigma^2}{2} + O(\sigma^3) & R = 1, \\ \frac{\sigma^2}{R-1} + O(\sigma^4) & R > 1. \end{cases} \tag{6.36}$$

The weak-noise behavior of D_{L}^* and D_{L} are compared in different regimes of measurement rates:

- $0 \le R < 1$: both D_{L}^* and D_{L} converge to $1 - R > 0$. This is an intuitive result, because even in the absence of noise, the orthogonal projection of the input vector onto the nullspace of the sensing matrix cannot be recovered, which contributes a total mean-square error of $(1 - R)n$; Moreover, D_{L} has strictly worse second-order asymptotics than D_{L}^*, especially when R is close to 1.
- $R = 1$: $D_{\mathrm{L}} = \sigma(1+o(1))$ is much worse than $D_{\mathrm{L}}^* = \sigma^2(1+o(1))$, which is achieved by choosing the encoding matrix to be identity. In fact, with nonnegligible probability, the optimal estimator that attains (6.32) blows up the noise power when inverting the random matrix;
- $R > 1$: both D_{L}^* and D_{L} behave according to $\Theta(\sigma^2)$, but the scaling constant of D_{L}^* is strictly worse, especially when R is close to 1.

The foregoing high-SNR analysis shows that the average performance of random sensing matrices with i.i.d. entries is much worse than that of optimal sensing matrices, except if $R \ll 1$ or $R \gg 1$. Although this conclusion stems from the high-SNR asymptotics, we test it with several numerical results. Fig. 6.3 ($R = 0.3$ and 5) and Fig. 6.4 ($\sigma^2 = 1$) illustrate that the superiority of optimal sensing matrices carries over to the regime of non-vanishing σ^2. However, as we will see, randomly selected matrices are as good as the optimal matrices (and in fact, optimal nonlinear encoders) as far as the phase transition threshold of the noise sensitivity is concerned.

From (6.34) and (6.36), we observe that both D_{L}^* and D_{L} exhibit a sharp phase transition near the critical rate $R = 1$:

$$\lim_{\sigma^2 \to 0} D_{\mathrm{L}}^*(X, R, \sigma^2) = \lim_{\sigma^2 \to 0} D_{\mathrm{L}}(X, R, \sigma^2) \tag{6.37}$$

$$= (1 - R)^+. \tag{6.38}$$

where $x^+ \triangleq \max\{0, x\}$. Moreover, from (6.30) – (6.32) we obtain the worst-case and asymptotic noise sensitivity functions for the Gaussian input as follows:

$$\zeta^*(X_{\mathsf{G}}, R) = \begin{cases} \frac{(R-1)^{R-1}}{R^R} & R \geq 1 \\ \infty & R < 1 \end{cases}, \quad \xi^*(X_{\mathsf{G}}, R) = \begin{cases} 0 & R > 1 \\ 1 & R = 1 \\ \infty & R < 1 \end{cases} \tag{6.39}$$

$$\zeta_{\mathrm{L}}^*(X_{\mathsf{G}}, R) = \xi_{\mathrm{L}}^*(X_{\mathsf{G}}, R) = \begin{cases} \frac{1}{R} & R \geq 1 \\ \infty & R < 1 \end{cases} \tag{6.40}$$

and

$$\zeta_{\mathrm{L}}(X_{\mathsf{G}}, R) = \xi_{\mathrm{L}}(X_{\mathsf{G}}, R) = \begin{cases} \frac{1}{R-1} & R > 1 \\ \infty & R \leq 1 \end{cases} \tag{6.41}$$

The worst-case noise sensitivity functions are plotted in Fig. 6.5 against the measurement rate R. Note that (6.39) provides an example for Remark 24: for Gaussian input and $R > 1$, the asymptotic noise sensitivity for optimal coding is zero, while the worst-case noise sensitivity is always strictly positive.

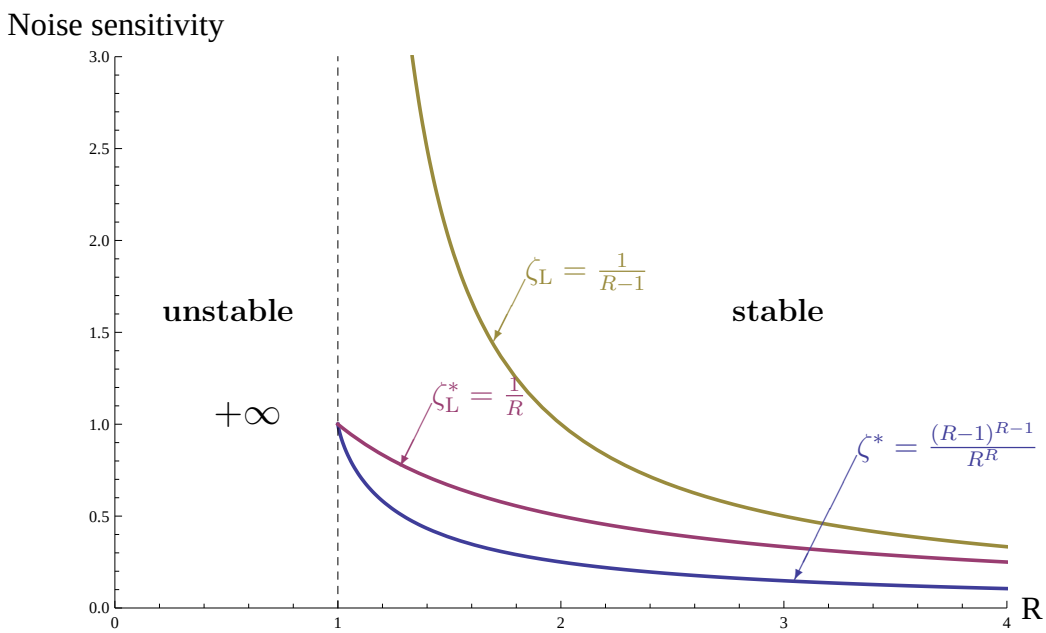

Figure 6.5: Worst-case noise sensitivity $\zeta^*, \zeta_{\mathrm{L}}^*$ and ζ_{L} for the Gaussian input, which all become infinity when $R < 1$ (the unstable regime).

In view of (6.39) – (6.41), the phase-transition thresholds in the Gaussian signal case are:

$$\mathcal{R}^*(X_{\mathsf{G}}) = \mathcal{R}_{\mathrm{L}}^*(X_{\mathsf{G}}) = \mathcal{R}_{\mathrm{L}}(X_{\mathsf{G}}) = 1. \tag{6.42}$$

The equality of the three phase-transition thresholds is in fact very general. In the next section, we formulate and prove the existence of the phase transition thresholds

of all three distortion-rate functions for discrete-continuous mixtures, which turn out to be equal to the information dimension of the input distribution.

6.5 Non-Gaussian inputs

This section contains our main results, which show that the phase transition thresholds are equal to the information dimension of the input, under rather general conditions. Therefore, the optimality of random sensing matrices in terms of the worst-case sensitivity observed in Section 6.4 carries over well beyond the Gaussian case. Proofs are deferred to Section 6.6.

The phase transition threshold for *optimal encoding* is given by the upper *information dimension* of the input:

Theorem 47. *For any X that satisfies (2.10),*

$$\mathcal{R}^*(X) = \overline{d}(X). \tag{6.43}$$

Moreover, if P_X is a discrete-continuous mixture as in (1.4), then as $\sigma \to 0$,

$$D^*(X,R) = \frac{\exp\left(2H(P_{\mathrm{d}})\frac{1-\gamma}{\gamma} - 2D(P_{\mathrm{c}})\right)}{(1-\gamma)^{2(1-\gamma)}\gamma^{\gamma}}\sigma^{\frac{2R}{\gamma}}(1+o(1)) \tag{6.44}$$

where $D(\cdot)$ denotes the non-Gaussianness of a probability measure defined in (2.68). Consequently the asymptotic noise sensitivity of optimal encoding is

$$\xi^*(X,R) = \begin{cases} \infty & R < \gamma \\ \frac{\exp\left(2H(P_{\mathrm{d}})\frac{1-\gamma}{\gamma} - 2D(P_{\mathrm{c}})\right)}{(1-\gamma)^{2(1-\gamma)}\gamma^{\gamma}} & R = \gamma \\ 0 & R > \gamma. \end{cases} \tag{6.45}$$

Assuming the validity of the replica calculations of D_{L} in Claim 1, it can be shown that the phase transition threshold for random linear encoding is always sandwiched between the lower and the upper *MMSE dimension* of the input. The relationship between the MMSE dimension and the information dimension in Theorem 20 plays a key role in analyzing the minimizer of the free energy (6.14).[2]

Theorem 48. *Assume that Claim 1 holds for X. Then for any i.i.d. random measurement matrix $\boldsymbol{A}$ whose entries have zero mean and variance $\frac{1}{n}$,*

$$\underline{\mathscr{D}}(X) \le \mathcal{R}_{\mathrm{L}}(X) \le \overline{\mathscr{D}}(X). \tag{6.46}$$

[2]It can be shown that in the limit of $\sigma^2 \to 0$, the minimizer of (6.14) when $R > \overline{\mathscr{D}}(X)$ and $R < \underline{\mathscr{D}}(X)$ corresponds to the largest and the smallest root of the fixed-point equation (6.12) respectively.

Therefore if $\mathscr{D}(X)$ exists, we have

$$\mathcal{R}_{\mathrm{L}}(X) = \mathscr{D}(X) = d(X), \tag{6.47}$$

and in addition,

$$D_{\mathrm{L}}(X, R, \sigma^2) = \frac{d(X)}{R - d(X)}\sigma^2(1 + o(1)). \tag{6.48}$$

Remark 26. In fact, the proof of Theorem 48 shows that the converse part (left inequality) of (6.46) holds as long as there is no residual error in the weak-noise limit, that is, if $D_{\mathrm{L}}(X, R, \sigma^2) = o(1)$ as $\sigma^2 \to 0$, then $R \geq \underline{\mathscr{D}}(X)$ must hold. Therefore the converse in Theorem 48 holds even if we weaken the right-hand side in (6.29) from $O(\sigma^2)$ to $o(1)$.

The general result in Theorem 48 holds for any input distribution but relies on the conjectured validity of Claim 1. For the special case of discrete-continuous mixtures as in (1.4), in view of Theorem 27, (6.47) predicts that the phase transition threshold for random linear encoding is γ. As the next result shows, this conclusion can indeed be rigorously established for random matrix encoders with i.i.d. *Gaussian* entries. Since, in view of Theorem 47, nonlinear encoding achieves the same threshold, (random) linear encoding suffices for robust reconstruction as long as the input distribution contains no singular component.

Theorem 49. *Assume that X has a discrete-continuous mixed distribution as in (1.4), where the discrete component P^{d} has finite entropy. Then*

$$\mathcal{R}^*(X) = \mathcal{R}^*_{\mathrm{L}}(X) = \mathcal{R}_{\mathrm{L}}(X) = \gamma. \tag{6.49}$$

Moreover, (6.49) holds for any non-Gaussian noise distribution with finite non-Gaussianness.

Remark 27. The achievability proof of $\mathcal{R}_{\mathrm{L}}(X)$ is a direct application of Theorem 32, where the Lipschitz decompressor in the noiseless case is used as a suboptimal estimator in the noisy case. The outline of the argument is as follows: suppose that we have obtained a sequence of linear encoders and $L(R)$-Lipschitz decoders $\{(\mathbf{A}_n, g_n)\}$ with rate R and error probability $\epsilon_n \to 0$ as $n \to \infty$. Then

$$\mathbb{E}\left[\left\|g_n(\mathbf{A}_n X^n + \sigma N^k) - X^n\right\|^2\right] \leq L(R)^2\sigma^2\mathbb{E}\left[\left\|N^k\right\|^2\right] + \epsilon_n = kL(R)^2\sigma^2\mathsf{var}N + \epsilon_n, \tag{6.50}$$

which implies robust construction is achievable at rate R and the worst-case noise sensitivity is upper bounded by $L(R)^2R$.

Notice that the above Lipschitz-based noisy achievability approach applies to any noise with finite variance, without requiring that the noise be additive, memoryless or that it have a density. In contrast, replica-based results rely crucially on the fact that the additive noise is memoryless Gaussian. Of course, in order for the converse (via

$\mathcal{R}^*$) to hold, the non-Gaussian noise needs to have finite non-Gaussianness. The disadvantage of this approach is that currently it lacks an explicit construction because the extendability of Lipschitz functions (Kirszbraun's theorem) is only an existence result which relies on Hausdorff maximal principle [100, Theorem 1.31, p. 21].

In view of Remark 13, it it interesting to examine whether the Lipschitz noiseless decoder allows us to achieve the optimal asymptotic noise sensitivity $\frac{1}{R-\gamma}$ predicted by the replica method in (6.48). If the discrete component of the input distribution has non-zero entropy (e.g., for simple signals (1.3)), the Lipschitz constant in (5.21) depends exponentially on $\frac{1}{R-\gamma}$. However, for sparse input (1.2), the Lipschitz constant in (5.21) behaves as $\frac{1}{\sqrt{R-\delta}}$, which, upon squaring, attains the optimal noise sensitivity.

Remark 28. We emphasize the following "universality" aspects of Theorem 49:

- Gaussian random sensing matrices achieve the optimal transition threshold for any discrete-continuous mixture;
- The fundamental limit depends on the input statistics only through the weight on the analog component, regardless of the specific discrete and continuous components. In the conventional sparsity model (1.2) where P_X is the mixture of an absolutely continuous distribution and a mass of $1-\gamma$ at 0, the fundamental limit is γ;
- The suboptimal estimator used in the achievability proof comes from the noiseless Lipschitz decoder, which does not depend on the noise distribution, let alone its variance;
- The conclusion holds for non-Gaussian noise as long as it has finite non-Gaussianness.

Remark 29. Assume the validity of Claim 1. Combining Theorem 47, Theorem 48 and (6.27) gives an *operational* proof for the inequality $\overline{d}(X) \leq \overline{\mathscr{D}}(X)$, which has been proven analytically in Theorem 20.

6.6 Proofs

Proof of Theorem 47. The proof of (6.26) is based on the low-distortion asymptotics of $R_X(D)$. Particularizing Theorem 8 to the MSE distortion ($r=2$), we have

$$\limsup_{D\downarrow 0} \frac{R_X(D)}{\frac{1}{2}\log\frac{1}{D}} = \overline{d}(X), \tag{6.51}$$

Converse: Fix $R > \mathcal{R}^*(X)$. By definition, there exits $a > 0$ such that $D^*(X,R,\sigma^2) \leq a\sigma^2$ for all $\sigma^2 > 0$. By (6.6),

$$\frac{R}{\frac{1}{2}\log(1+\sigma^{-2})} \geq R_X(a\sigma^2). \tag{6.52}$$

Dividing both sides by $\frac{1}{2}\log\frac{1}{a\sigma^2}$ and taking $\limsup_{\sigma^2\to 0}$ yield $R > \overline{d}(X)$ in view of (6.51). By the arbitrariness of R, we have $\mathcal{R}^*(X) > \overline{d}(X)$.

Achievability: Fix $\delta > 0$ arbitrarily and let $R = \overline{d}(X) + 2\delta$. We show that $R \le \mathcal{R}^*(X)$, i.e., worst-case noise sensitivity is finite. By Remark 24, this is equivalent to achieving (6.29). By (6.51), there exists $D_0 > 0$ such that for all $D < D_0$,

$$R_X(D) \le \frac{\overline{d}(X) + \delta}{2} \log \frac{1}{D}. \tag{6.53}$$

By Theorem 45, $D^*(X, R, \sigma^2) \downarrow 0$ as $\sigma^2 \downarrow 0$. Therefore there exists $\sigma_0^2 > 0$, such that $D^*(X, R, \sigma^2) < D_0$ for all $\sigma^2 < \sigma_0^2$. In view of (6.6) and (6.53), we have

$$\frac{d + 2\delta}{2} \log \frac{1}{\sigma^2} = R_X(D^*(X, R, \sigma^2)) \le \frac{\overline{d}(X) + \delta}{2} \log \frac{1}{D^*(X, R, \sigma^2)}, \tag{6.54}$$

i.e.,

$$D^*(X, R, \sigma^2) \le \sigma^{2\frac{\overline{d}(X)+2\delta}{\overline{d}(X)+\delta}} \tag{6.55}$$

holds for all $\sigma^2 < \sigma_0^2$. This obviously implies the desired (6.29).

We finish the proof by proving (6.44) and (6.45). The low-distortion asymptotic expansion of rate-distortion functions of discrete-continuous mixtures is found in [42, Corollary 1], which improves (2.64):[3]

$$R_X(D) = \frac{\gamma}{2} \log \frac{\gamma}{D} + h(\gamma) + (1 - \gamma) H(P_{\mathrm{d}}) + \gamma h(P_{\mathrm{c}}) + o(1), \tag{6.56}$$

where P_X is given by (1.4). Plugging into (6.6) gives (6.44), which implies (6.45) as a direct consequence. □

Proof of Theorem 48. Achievability: We show that $\mathcal{R}_{\mathrm{L}}(X) \le \overline{\mathscr{D}}(X)$. Fix $\delta > 0$ arbitrarily and let $R = \overline{\mathscr{D}}(X) + 2\delta$. Set $\gamma = R\,\sigma^{-2}$ and $\beta = \eta\gamma$. Define

$$u(\beta) = \beta\,\mathsf{mmse}(X, \beta) - R\left(1 - \frac{\beta}{\gamma}\right) \tag{6.57}$$

$$f(\beta) = I(X; \sqrt{\beta} X + NG) - \frac{R}{2} \log \beta \tag{6.58}$$

$$g(\beta) = f(\beta) + \frac{R\beta}{2\gamma}, \tag{6.59}$$

which satisfy the following properties:

1. Since $\mathsf{mmse}(X, \cdot)$ is smooth on $(0, \infty)$ [66, Proposition 7], u, f and g are all smooth functions on $(0, \infty)$. In particular, by the I-MMSE relationship [65],

$$\dot{f}(\beta) = \frac{\beta\,\mathsf{mmse}(X, \beta) - R}{2\beta}. \tag{6.60}$$

[3] In fact $h(\gamma) + (1-\gamma)H(P_{\mathrm{d}}) + \gamma h(P_{\mathrm{c}})$ is the γ-dimensional entropy of (1.4) defined by Rényi [24, Equation (4) and Theorem 3].

2. For all $0 \le \beta \le \gamma$,

$$f(\beta) \le g(\beta) \le f(\beta) + \frac{R}{2}. \tag{6.61}$$

3. Recalling the scaling law of mutual information in (2.66), we have

$$\limsup_{\beta\to\infty} \frac{f(\beta)}{\log \beta} = \frac{\overline{d}(X) - \overline{\mathscr{D}}(X) - \delta}{2} \le -\frac{\delta}{2}, \tag{6.62}$$

where the last inequality follows from the sandwich bound between information dimension and MMSE dimension in (4.46).

Note that the solutions to (6.13) are precisely the roots of u, at least one of which belongs to $[0, \gamma]$ because $u(0) = -R < 0$, $u(\gamma) = \gamma \mathsf{mmse}(X, \gamma) > 0$ and u is continuous. According to Claim 1,

$$D_{\mathrm{L}}(X, R, \sigma^2) = \mathsf{mmse}(X, \beta_\gamma), \tag{6.63}$$

where β_γ is the root of u in $[0, \gamma]$ that maximizes $g(\beta)$.

Proving the achievability of R amounts to showing that

$$\limsup_{\sigma\to 0} \frac{D_{\mathrm{L}}(X, R, \sigma^2)}{\sigma^2} < \infty, \tag{6.64}$$

which, in view of (6.63), is equivalent to

$$\liminf_{\gamma\to\infty} \frac{\beta_\gamma}{\gamma} > 0. \tag{6.65}$$

By the definition of $\overline{\mathscr{D}}(X)$ and (6.62), there exists $B > 0$ such that for all $\beta > B$,

$$\beta\, \mathsf{mmse}(X, \beta) < R - \delta \tag{6.66}$$

and

$$f(\beta) \le -\frac{\delta}{4} \log \beta. \tag{6.67}$$

In the sequel we focus on sufficiently large γ. Specifically, we assume that

$$\gamma > \frac{R}{\delta} \max\left\{B, \mathrm{e}^{-\frac{4}{\delta}\left(K - \frac{R}{2}\right)}\right\}, \tag{6.68}$$

where $K \triangleq \min_{\beta\in[0,B]} g(\beta)$ is finite by the continuity of g.

Let

$$\beta_0 = \frac{\delta\gamma}{R}. \tag{6.69}$$

Then $\beta_0 > B$ by (6.68). By (6.66), $u(\beta_0) = \beta_0\, \mathsf{mmse}(X, \beta_0) - R + \delta < 0$. Since $u(\gamma) > 0$, by the continuity of u and the intermediate value theorem, there exists

$\beta_0 \leq \beta^* \leq \gamma$, such that $u(\beta^*) = 0$. By (6.66),

$$\dot{f}(\beta) \leq -\frac{\delta}{2\beta} < 0, \quad \forall \beta > B. \tag{6.70}$$

Hence f strictly decreases on (B, ∞). Denote the root of u that minimizes $f(\beta)$ by β'_γ, which must lie beyond β^*. Consequently, we have

$$B < \frac{\delta\gamma}{R} = \beta_0 \leq \beta^* \leq \beta'_\gamma. \tag{6.71}$$

Next argue that β_γ cannot differ from β'_γ by a constant factor. In particular, we show that

$$\beta_\gamma \geq \mathrm{e}^{-\frac{R}{\delta}} \beta'_\gamma, \tag{6.72}$$

which, combined with (6.71), implies that

$$\frac{\beta_\gamma}{\gamma} \geq \frac{\delta}{R} \mathrm{e}^{-\frac{R}{\delta}} \tag{6.73}$$

for all γ that satisfy (6.68). This yields the desired (6.65). We now complete the proof by showing (6.72). First, we argue that that $\beta_\gamma > B$. This is because

$$\begin{aligned}
g(\beta_\gamma) &\leq g(\beta'_\gamma) && (6.74)\\
&= f(\beta'_\gamma) + \frac{R\beta'_\gamma}{2\gamma} && (6.75)\\
&\leq f(\beta_0) + \frac{R}{2} && (6.76)\\
&\leq -\frac{\delta}{4} \log \frac{\delta\gamma}{R} + \frac{R}{2} && (6.77)\\
&< K && (6.78)\\
&= \min_{\beta \in [0,B]} g(\beta). && (6.79)
\end{aligned}$$

where

- (6.74): by definition, β_γ and β'_γ are both roots of u and β'_γ minimizes g among all roots;
- (6.76): by (6.71) and the fact that f is strictly decreasing on (B, ∞);
- (6.77): by (6.67);
- (6.79): by (6.68).

Now we prove (6.72) by contradiction. Suppose $\beta_\gamma < \mathrm{e}^{-\frac{R}{\delta}}\beta'_\gamma$. Then

$$g(\beta'_\gamma) - g(\beta_\gamma) = \frac{R}{2\gamma}(\beta'_\gamma - \beta_\gamma) + f(\beta'_\gamma) - f(\beta_\gamma) \tag{6.80}$$

$$\leq \frac{R}{2} + \int_{\beta_\gamma}^{\beta'_\gamma} \dot{f}(b)\mathrm{d}b \tag{6.81}$$

$$\leq \frac{R}{2} - \frac{\delta}{2}\log\frac{\beta'_\gamma}{\beta_\gamma} \tag{6.82}$$

$$< 0, \tag{6.83}$$

contradicting (6.74), where (6.82) is due to (6.70).

Converse: We show that $\mathcal{R}_{\mathrm{L}}(X) \geq \underline{\mathscr{D}}(X)$. Recall that $\mathcal{R}_{\mathrm{L}}(X)$ is the minimum rate that guarantees that the reconstruction error $D_{\mathrm{L}}(X, R, \sigma^2)$ vanishes according to $O(\sigma^2)$ as $\sigma^2 \to 0$. In fact, we will show a stronger result: as long as $D_{\mathrm{L}}(X, R, \sigma^2) = o(1)$ as $\sigma^2 \to 0$, we have $R \geq \underline{\mathscr{D}}(X)$. By (6.63), $D_{\mathrm{L}}(X, R, \sigma^2) = o(1)$ if and only if $\beta_\gamma \to \infty$. Since $u(\beta_\gamma) = 0$, we have

$$R \geq \limsup_{\gamma\to\infty} R\left(1 - \frac{\beta_\gamma}{\gamma}\right) \tag{6.84}$$

$$= \limsup_{\gamma\to\infty} \beta_\gamma \,\mathsf{mmse}(X, \beta_\gamma) \tag{6.85}$$

$$\geq \liminf_{\beta\to\infty} \beta \,\mathsf{mmse}(X, \beta) \tag{6.86}$$

$$= \underline{\mathscr{D}}(X). \tag{6.87}$$

Asymptotic noise sensitivity: Finally, we prove (6.48). Assume that $\mathscr{D}(X)$ exists, i.e., $\mathscr{D}(X) = d(X)$, in view of (4.46). By definition of $\mathscr{D}(X)$, we have

$$\mathsf{mmse}(X, \beta) = \frac{\mathscr{D}(X)}{\beta} + o\left(\frac{1}{\beta}\right), \quad \beta \to \infty. \tag{6.88}$$

As we saw in the achievability proof, whenever $R > \mathscr{D}(X)$, (6.65) holds, i.e., $\eta_\gamma = \Omega(1)$ as $\gamma \to \infty$. Therefore, as $\gamma \to \infty$, we have

$$\frac{1}{\eta_\gamma} = 1 + \frac{\gamma}{R}\mathsf{mmse}(X, \eta_\gamma\gamma) = 1 + \frac{\mathscr{D}(X)}{\eta_\gamma R} + o(1), \tag{6.89}$$

i.e.,

$$\eta_\gamma = 1 - \frac{\mathscr{D}(X)}{R} + o(1). \tag{6.90}$$

By Claim 1,

$$D_{\mathrm{L}}(X, R, \sigma^2) = \mathsf{mmse}(X, \eta_\gamma \gamma) \tag{6.91}$$

$$= \frac{1-\eta_\gamma}{\eta_\gamma}\sigma^2 \tag{6.92}$$

$$= \frac{\mathscr{D}(X)}{R - \mathscr{D}(X)}\sigma^2(1+o(1)). \tag{6.93}$$

□

Remark 30. Note that β_γ is a *subsequence* parametrized by γ, which may take only a restricted subset of values. In fact, even if we impose the requirement that $D_{\mathrm{L}}(X, R, \sigma^2) = O(\sigma^2)$, it is still possible that the limit in (6.85) lies strictly between $\underline{\mathscr{D}}(X)$ and $\overline{\mathscr{D}}(X)$. For example, if X is Cantor distributed, it can be shown that the limit in (6.85) approaches the information dimension $d(X) = \log_3 2$.

Proof of Theorem 49. Let $R > \gamma$. We show that the worst-case noise sensitivity $\zeta^*(X, R)$ under Gaussian random sensing matrices is finite. We construct a suboptimal estimator based on the Lipschitz decoder in Theorem 32.[4] Let $\boldsymbol{A}_n$ be a $k \times n$ Gaussian sensing matrix and $g_{\boldsymbol{A}_n}$ the corresponding $L(R)$-Lipschitz decoder, such that $k = Rn$ and $\mathbb{P}\{E_n\} = o(1)$ where $E_n = \{g_{\boldsymbol{A}_n}(\boldsymbol{A}_n X^n) \neq X^n\}$ denotes the error event. Without loss of generality, we assume that $g_{\boldsymbol{A}_n}(0) = 0$. Fix $\tau > 0$. Then

$$\begin{aligned}
&\mathbb{E}\left[\left\|g_{\boldsymbol{A}_n}(\boldsymbol{A}_n X^n + \sigma N^k) - X^n\right\|^2\right] \\
&\le \mathbb{E}\left[\left\|g_{\boldsymbol{A}_n}(\boldsymbol{A}_n X^n + \sigma N^k) - X^n\right\|^2 \mathbf{1}_{\{E_n^{\mathrm{c}}\}}\right] \\
&\quad + 2L(R)^2 \mathbb{E}\left[\left\|\boldsymbol{A}_n X^n + \sigma N^k\right\|^2 \mathbf{1}_{\{E_n\}}\right] + 2\mathbb{E}\left[\left\|X^n\right\|^2 \mathbf{1}_{\{E_n\}}\right] \qquad (6.94) \\
&\le kL(R)^2\sigma^2 + \tau n(2L(R)^2+1)\mathbb{P}\{E_n\} + 2\mathbb{E}\left[\left\|X^n\right\|^2 \mathbf{1}_{\left\{\|X^n\|^2 > \tau n\right\}}\right] \\
&\quad + 2L(R)^2 \mathbb{E}\left[\left\|\boldsymbol{A}_n X^n + \sigma N^k\right\|^2 \mathbf{1}_{\left\{\left\|\boldsymbol{A}_n X^n + \sigma N^k\right\|^2 > \tau n\right\}}\right] \qquad (6.95)
\end{aligned}$$

Dividing both sides of (6.95) by n and sending $n \to \infty$, we have: for any $\tau > 0$,

$$\begin{aligned}
&\limsup_{n\to\infty} \frac{1}{n}\mathbb{E}\left[\left\|g_{\boldsymbol{A}_n}(\boldsymbol{A}_n X^n + \sigma N^k) - X^n\right\|^2\right] \\
&\le RL(R)^2\sigma^2 + 2\sup_n \frac{1}{n}\mathbb{E}\left[\left\|X^n\right\|^2 \mathbf{1}_{\left\{\|X^n\|^2 > \tau n\right\}}\right] \\
&\quad + 2L(R)^2 \sup_n \frac{1}{n}\mathbb{E}\left[\left\|\boldsymbol{A}_n X^n + \sigma N^k\right\|^2 \mathbf{1}_{\left\{\left\|\boldsymbol{A}_n X^n + \sigma N^k\right\|^2 > \tau n\right\}}\right]. \qquad (6.96)
\end{aligned}$$

[4] Since we assume that $\mathsf{var} X = 1$, the finite-entropy condition of Theorem 32 is satisfied automatically.

Since $\frac{1}{n}\|X^n\|^2 \xrightarrow{L^2} 1$ and $\frac{1}{n}\left\|\boldsymbol{A}_n X^n + \sigma N^k\right\|^2 \xrightarrow{L^2} R(1+\sigma^2)$, which implies uniform integrability, the last two terms on the right-hand side of (6.96) vanish as $\tau \to \infty$. This completes the proof of $\zeta^*(X, R) \leq RL^2(R)$. $\square$

Chapter 7

Suboptimality of practical reconstruction algorithms

In this section we compare the optimal phase transition thresholds derived in Chapters 5 and 6 to the suboptimal thresholds obtained by popular low-complexity algorithms, analyzed in [15, 4, 6]. The material in this chapter has been presented in part in [116].

We focus on the following three input distributions considered in [4, p. 18915], indexed by $\chi = \pm, +$ and $\square$ respectively, which all belong to the family of input distributions of the mixture form in (1.4):

- $\pm$: sparse signals (1.2);
- $+$: sparse non-negative signals (1.2) with the absolutely continuous component P_{c} supported on $\mathbb{R}_+$.
- $\square$: simple signals (1.3) with $p = \frac{1}{2}$ [15, Section 5.2, p. 540] where P_{c} is some absolutely continuous distribution supported on the unit interval.

7.1 Noiseless compressed sensing

In the noiseless case, we consider the phase transition of error probability. The phase transition thresholds for linear programming (LP) decoders and the approximate message passing (AMP) decoder [4, 123] have been obtained in [3, 4]. Greedy reconstruction algorithms have been analyzed in [124], which only provide upper bounds (achievability results) on the transition threshold of measurement rate. In this section we focus our comparison on algorithms whose phase transition thresholds are known exactly.

The following LP decoders are tailored to the three input distributions $\chi = \pm, +$ and $\square$ respectively (see Equations (P1), (LP) and (Feas) in [16, Section I]):

$$g_{\pm}(y) = \arg\min\{\|x\|_1 : x \in \mathbb{R}^n, \boldsymbol{A}x = y\}, \tag{7.1}$$

$$g_{+}(y) = \arg\min\{\|x\|_1 : x \in \mathbb{R}^n_+, \boldsymbol{A}x = y\}, \tag{7.2}$$

$$g_{\square}(y) = \{x : x \in [0,1]^n, \boldsymbol{A}x = y\}. \tag{7.3}$$

For sparse signals, (7.1) – (7.2) are based on ℓ_1-minimization (also known as Basis Pursuit) [12], while for simple signals, the decoder (7.3) solves an LP feasibility problem. Note that the above decoders are universal in the sense that it require no knowledge of the input distributions other than which subclass they belong. In general the decoders in (7.1) – (7.3) output a list of vectors upon receiving the measurement. The reconstruction is successful if and only if the output list contains only the true vector. The error probability is thus defined as $\mathbb{P}\left\{g_\chi(\boldsymbol{A}X^n) \neq \{X^n\}\right\}$, evaluated with respect to the product measure $(P_X)^n \times P_{\boldsymbol{A}}$.

The phase transition thresholds of the reconstruction error probability for decoders (7.1) – (7.3) are derived in [3] using combinatorial geometry. For sparse signals and ℓ_1-minimization decoders (7.1) – (7.2), the expressions of the corresponding thresholds $\mathsf{R}_\pm(\gamma)$ and $\mathsf{R}_+(\gamma)$ are quite involved, given implicitly in [3, Definition 2.3]. As observed in [4, Finding 1], $\mathsf{R}_\pm(\gamma)$ and $\mathsf{R}_+(\gamma)$ agree numerically with the solutions of following equation, respectively:[1]

$$\gamma = \max_{z\geq 0} \frac{R - 2((1+z^2)\Phi(-z) - z\varphi(z))}{1 + z^2 - 2((1+z^2)\Phi(-z) - z\varphi(z))}, \tag{7.4}$$

$$\gamma = \max_{z\geq 0} \frac{R - ((1+z^2)\Phi(-z) - z\varphi(z))}{1 + z^2 - ((1+z^2)\Phi(-z) - z\varphi(z))}, \tag{7.5}$$

For simple signals, the phase transition threshold is proved to be [15, Theorem 1.1]

$$\mathsf{R}_\square(\gamma) = \frac{\gamma+1}{2}. \tag{7.6}$$

Moreover, substantial *numerical* evidence in [4, 123] suggests that the phrase transition thresholds for the AMP decoder coincide with the LP thresholds for all three input distributions. The suboptimal thresholds obtained from (7.4) – (7.6) are plotted in Fig. 7.1 along with the optimal threshold obtained from Theorem 32 which is γ.[2] In the gray area below the diagonal in the (γ, R)-phase diagram, any sequence of sensing matrices and decompressors will fail to reconstruct the true signal with probability that tends to one. Moreover, we observe that the LP or AMP decoder is severely suboptimal unless γ is close to one.

In the highly sparse regime which is most relevant to compressed sensing problems, it follows from [16, Theorem 3] that for sparse signals ($\chi = \pm$ and $+$),

$$\mathsf{R}_\chi(\gamma) = 2\gamma \log_{\mathrm{e}} \frac{1}{\gamma}(1 + o(1)), \text{ as } \gamma \to 0, \tag{7.7}$$

[1] In the series of papers (e.g. [15, 16, 4, 6]), the phase diagram are parameterized by (ρ, δ), where $\delta = R$ is the measurement rate and $\rho = \frac{\gamma}{R}$ is the ratio between the sparsity and rate. In this thesis, the more straightforward parameterization (γ, R) is used instead. The ratio $\frac{\gamma}{\mathsf{R}_\chi(\gamma)}$ is denoted by $\rho(\gamma;\chi)$ in [4].

[2] A similar comparison between the suboptimal threshold $\mathsf{R}_\pm(\gamma)$ and the optimal threshold γ has been provided in [10, Fig. 2(a)] based on replica-heuristic calculation.

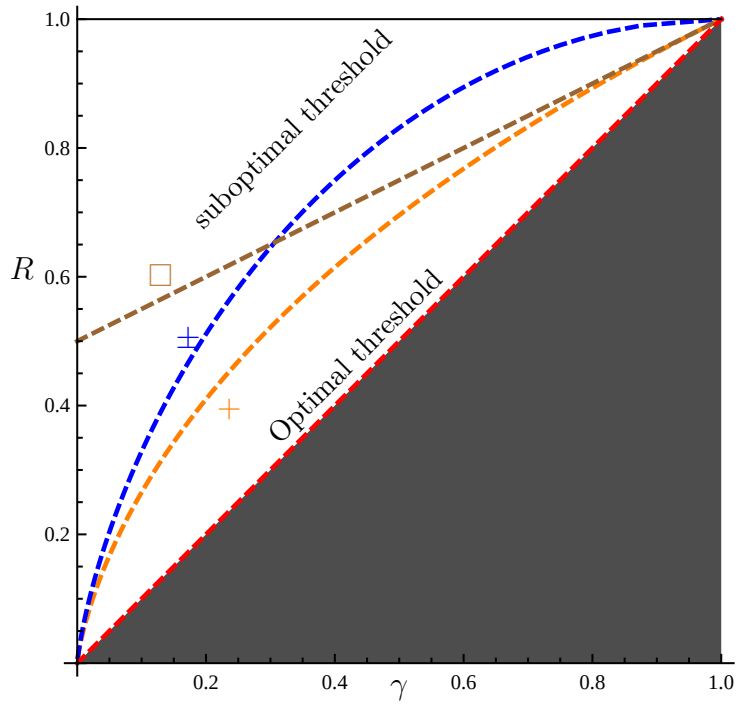

Figure 7.1: Comparison of suboptimal thresholds (7.4) – (7.6) with optimal threshold.

which implies that R_χ has infinite slope at $\gamma = 0$. Therefore when $\gamma \ll 1$, the ℓ_1 and AMP decoders require on the order of $2s \log_{\mathrm{e}} \frac{n}{s}$ measurements to successfully recover the unknown s-sparse vector. In contrast, s measurements suffice when using an optimal decoder (or ℓ_0-minimization decoder). The LP or AMP decoders are also highly suboptimal for simple signals, since $\mathsf{R}_\square(\gamma)$ converges to $\frac{1}{2}$ instead of zero as $\gamma \to 0$. This suboptimality is due to the fact that the LP feasibility decoder (7.3) simply finds any x^n in the hypercube $[0,1]^n$ that is compatible with the linear measurements. Such a decoding strategy does not enforce the typical discrete structure of the signal, since most of the entries saturate at 0 or 1 equiprobably. Alternatively, the following decoder achieves the optimal γ: define

$$\mathsf{T}(x^n) = \frac{1}{n} \sum_{i=1}^{n} \left(\mathbf{1}_{\{x_i=0\}}, \mathbf{1}_{\{x_i \notin \{0,1\}\}}, \mathbf{1}_{\{x_i=1\}} \right).$$

The decoder outputs the solution to $\boldsymbol{A}x^n = y^k$ such that $\mathsf{T}(x^n)$ is closest to $\left(\frac{1-\gamma}{2}, \gamma, \frac{1-\gamma}{2}\right)$ (in total variation distance for example).

7.2 Noisy compressed sensing

In the noisy case, we consider the AMP decoder [6] and the ℓ_1-penalized least-squares (i.e. LASSO) decoder [13]:

$$g^{\lambda}_{\mathrm{LASSO}}(y) = \operatorname*{argmin}_{x \in \mathbb{R}^n} \frac{1}{2} \|y - \boldsymbol{A}x\|_2 + \lambda \|x\|_1 , \tag{7.8}$$

where $\lambda > 0$ is a regularization parameter. For Gaussian sensing matrices and Gaussian observation noise, the asymptotic mean-square error achieved by LASSO for a fixed λ

$$\lim_{n\to\infty} \frac{1}{n}\mathbb{E}\left[\left\|X^n - g^{\lambda}_{\text{LASSO}}(\boldsymbol{A}X^n + \sigma N^k)\right\|^2\right] \tag{7.9}$$

can be determined as a function of P_X, λ and σ by applying [125, Theorem 1.4] with $\psi(x,y) = (x-y)^2$. For sparse X distributed according to (1.4), the following *formal* expressions for the worst-case (or asymptotic)[3] noise sensitivity of LASSO with optimized λ is proposed in [6, Proposition 3.1(1.a)] and verified by extensive empirical results: for any P_X of the mixture form (1.4),

$$\zeta_{\text{LASSO}}(X,R) = \xi_{\text{LASSO}}(X,R) = \begin{cases} \frac{R\,\mathsf{R}_{\pm}(\gamma)}{R-\mathsf{R}_{\pm}(\gamma)} & R > \mathsf{R}_{\pm}(\gamma) \\ \infty & R \le \mathsf{R}_{\pm}(\gamma). \end{cases} \tag{7.10}$$

where $\mathsf{R}_{\pm}(\gamma)$ is given in (7.5). Furthermore, numerical evidence suggests that (7.10) also applies to the AMP algorithm [6]. In view of (7.10), the phase transition thresholds of noise sensitivity for the LASSO and AMP decoder are both $\mathsf{R}_{\pm}(\gamma)$, whereas the optimal threshold is γ as a consequence of Theorem 49. Therefore the phase transition boundaries are identical to Fig. 7.1 and the same observation in Section 7.1 applies. For sparse signals of the form (1.2) with $\gamma = 0.1$, Fig. 7.2 compares those expressions for the asymptotic noise sensitivity of LASSO or AMP algorithms to the optimal noise sensitivity predicted by Theorem 48 based on replica heuristics. Note that the phase transition threshold of LASSO is approximately 3.3 times the optimal. Also, as the measurement rate R tends to infinity, the noise sensitivity of LASSO converges to $\mathsf{R}_{\pm}(\gamma) > 0$ rather than zero, while the optimal noise sensitivity vanishes at the speed of $\frac{1}{R}$.

[3]The worst-case and asymptotic noise sensitivity coincide since the worst-case noise variance tends to zero. This is because the least-favorable prior in [6, (2.5)] is supported on $\{0, \pm\infty\}$.

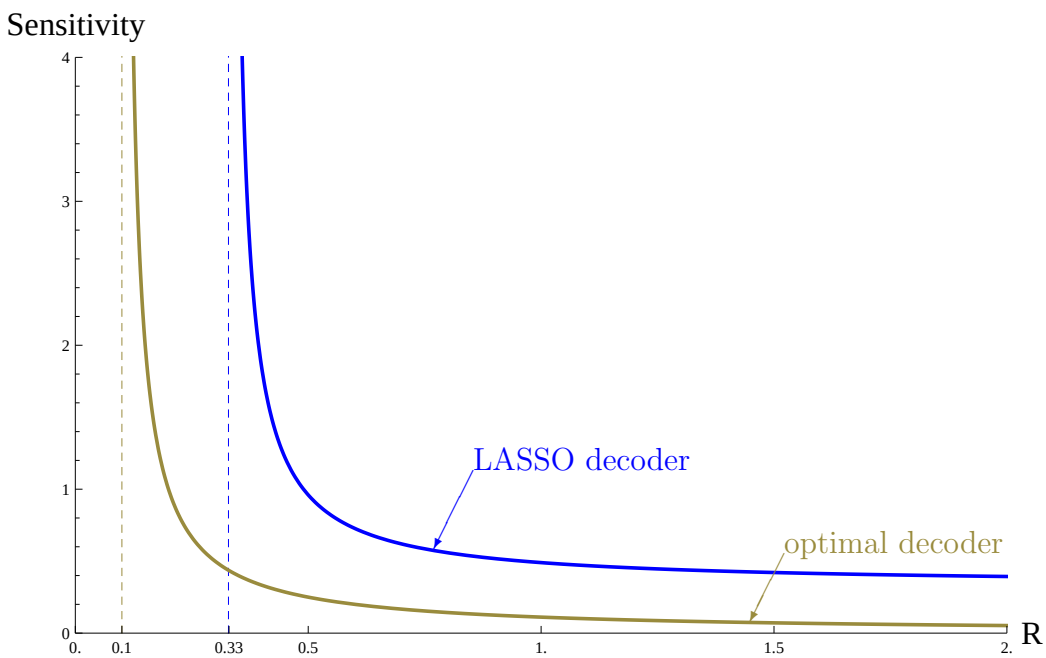

Figure 7.2: Phase transition of the asymptotic noise sensitivity: sparse signal model (1.2) with $\gamma = 0.1$.

Chapter 8

Conclusions and future work

In this thesis we adopt a Shannon-theoretic framework for compressed sensing by modeling input signals as random processes rather than individual sequences. We investigate the phase transition thresholds (minimum measurement rate) of reconstruction error probability (noiseless observations) and normalized MMSE (noisy observations) achievable by optimal nonlinear, optimal linear, and random linear encoders combined with the corresponding optimal decoders. For discrete-continuous input distributions, which are most relevant for compressed sensing applications, the optimal phase transition threshold is shown to be the information dimension of the input, i.e., the weight of the analog part, regardless of the specific discrete and absolutely continuous component. In the memoryless case, this corresponds to the fraction of analog symbols in the source realization. This observation might suggest that the fundamental importance of the mixed discrete-continuous nature of the signal, of which sparsity (or simplicity discussed in Section 1.3) is just one manifestation.

To conclude this chapter, we discuss several open problems in theoretical compressed sensing as future research directions.

8.1 Universality of sensing matrix ensemble

Random linear coding is not only an important achievability proof technique in Shannon theory, but also an inspiration to obtain efficient schemes in modern coding theory and compressed sensing.

In the noiseless case, Theorem 39 shows that for discrete-continuous mixtures, the information dimension γ can be achieved with Lebesgue-a.e. linear encoders. This implies that random linear encoders drawn from any absolutely continuous distribution achieve the desired error probability almost surely. However, it is not clear whether the optimal measurement rate γ can be obtained using discrete ensembles, e.g., Rademacher matrices.

In the noisy case, one of our main findings is that *Gaussian* sensing matrices achieve the same phase transition threshold as optimal nonlinear encoding. However, it should be noted that the argument used in the proof of Theorem 49 relies crucially on the Gaussianness of the sensing matrices because of two reasons:

- The upper bound on the distribution function of the least singular value in Lemma 23 is a direct consequence of the upper bound on its density due to Edelman [113], which only holds in the Gaussian case. In fact, we only need that the exponent in (5.161) diverges as $t \to 0$. It is possible to generalize this result to other sub-Gaussian ensembles with densities by adapting the arguments in [126, Theorem 1.1]. However, it should be noted that in general Lemma 23 does not hold for discrete ensembles (e.g. Rademacher), because the least singular value always has a mass at zero with a fixed exponent;
- Due to the rotational invariance of the Gaussian ensemble, the result in Lemma 24 does not depend on the basis of the subspace.

The universal optimality of random sensing matrices with non-Gaussian i.i.d. entries in terms of phase transition thresholds remains an open question question.

8.2 Near-optimal practical reconstruction algorithms

From Chapter 7 we see that the phase transition threshold of state-of-the-art decoding algorithms (e.g., LP-based and AMP decoders) still lie far from the optimal boundary. Designing low-complexity reconstruction schemes that attain the optimal thresholds remains an outstanding challenge. In particular, for mixture distributions (1.4), there are no known low-complexity algorithm whose phase transition threshold has finite slope at $\gamma = 0$.

In the case of sparse input distribution (1.2), it is interesting to notice that the optimal phase transition boundary in Fig. 7.1 can be attained by ℓ_0-minimization, while the threshold of ℓ_1-minimization is known to be strictly suboptimal. A natural question to ask is the performance of ℓ_q-minimization for $0 < q < 1$, which has been the focus of very active research, e.g., [127, 128, 129, 130, 131]. It is possible that as q moves from one to zero, the phase-transition boundary of ℓ_q-minimization evolves continuously from the suboptimal ℓ_1-boundary to the optimal ℓ_0-boundary. However, very little is known except the fact the the performance of ℓ_q-minimization does not deteriorate as q decreases [127]. Only upper bounds on the transition boundary are known [130, p. 11]. Note that as q decreases, the complexity of the ℓ_q-decoder increases as the optimization problem becomes more non-convex. Therefore the knowledge about the precise threshold will allow us to strike a good balance between decoding performance and computational complexity among this family of algorithms.

8.3 Performance of optimal sensing matrices

The majority of the compressed sensing literature is devoted to either randomly generated measurement matrices or sufficient conditions on individual measurement matrices that guarantee successful recovery of sparse vectors. Theorems 32 and 49 indicate that when *optimal* decoders are used, random sensing matrices achieve the

optimal phase transition threshold in both noiseless and noisy cases. However, with pragmatic *suboptimal* recovery algorithms, it might be possible to improve the performance by using optimized sensing matrices. Next we consider ℓ_p-minimization decoders with $p \in [0,1]$. Given a measurement matrix $\mathbf{A} \in \mathbb{R}^{k\times n}$, define the ℓ_p-minimization decoder as

$$g_p(y) = \arg\min_{z:\mathbf{A}z=y} \|z\|_p . \tag{8.1}$$

Now we define the following non-asymptotic fundamental limits of ℓ_p-minimization decoder by allowing to optimize over measurement matrices.

Definition 25. A $k \times n$ matrix $\mathbf{A}$ is called (s,n,p)-*successful* if using $\mathbf{A}$ as the measurement matrix, the ℓ_p-minimization decoder can recover all s-sparse vectors on $\mathbb{R}^n$, i.e., $g_p \circ \mathbf{A} = \mathsf{id}$ restricted on $\Sigma_{s,n}$, the collection of all s-sparse n-dimensional vectors. Define $k_p^*(s,n)$ as the minimum number of measurements that guarantee the success of ℓ_p-minimization decoder in recovering all s-sparse vectors, i.e., $k_p^*(s,n) = \min\{k \in \mathbb{N} : \exists \mathbf{A} \in \mathbb{R}^{k\times n}$ that is (s,n,p)-successful$\}$.

From Remark 17 we know that

$$k_0^*(s,n) = \min\{2s, n\}. \tag{8.2}$$

Also, it is trivial to show that for all p, $k_p^*(1,n) = \min\{2,n\}$. For the special case of $p=1$, [132, Proposition 2.3] characterizes a given matrix being $(s,n,1)$-successful in terms of its null space. However, determining $k_1^*(s,n)$ remains an open question. In the high-dimensional linearly sparse regime, we define

$$\mathsf{R}^*_{\ell_p}(\gamma) = \limsup_{n\to\infty} \frac{k_p^*(\lfloor \gamma n \rfloor, n)}{n}. \tag{8.3}$$

Then $\mathsf{R}^*_{\ell_p}(\gamma) \in [0,1]$ for all p and γ. From (8.2) we have

$$\mathsf{R}^*_{\ell_0}(\gamma) = \min\{2\gamma, 1\}, \tag{8.4}$$

while $\mathsf{R}^*_{\ell_p}(\gamma)$ is not known for $p \neq 0$.

Remark 31. If one relaxes Definition 25 and only requires that all s-sparse *positive* vectors are successfully recovered, it can be shown ([133, Corollary 1.1] and [134]) that ℓ_1-minimization attains the ℓ_0 performance in (8.2), achieved by a Fourier or Vandermonde matrix.

8.4 Compressive signal processing

The current information system consists of three stages: data acquisition, compression and processing. Compressed sensing permeates compression in the data collection process and achieves drastic reduction of data amount. Extending this rationale, it is more efficient to combine all three stages by directly processing the

information condensed from compressible signals. Fundamental limits of this more general paradigm, namely, *compressive signal processing*, are largely unknown.

As a concrete example, consider the problem of support recovery from noisy random linear measurements, where the goal is not to recover the unknown sparse vector with high fidelity, but to recover its support faithfully. It has been shown that if the SNR is fixed and sensing matrices are i.i.d., reliable support recovery (vanishing block error rate) requires infinite rate, because number of necessary measurements scales as $\Theta(n \log n)$ [135, 136]. On the other hand, if we tolerate any positive fraction of error (bit-error rate), finite rate can be achieved [9]. Note that the above conclusions apply to random sensing matrices. The performance of *optimal* sensing matrices in support recovery is a fundamental yet unexplored problem. It is yet to be determined whether asymptotic reliable support recovery is possible with finite SNR and finite measurement rate using optimized sensing matrices and decoders.

Appendix A

Proofs in Chapter 2

A.1 The equivalence of (2.7) to (2.4)

Proof. Let X be a random vector in $\mathbb{R}^n$ and denote its distribution by μ. Due to the equivalence of ℓ_p-norms, it is sufficient to show (2.7) for $p = \infty$. Recalling the notation for mesh cubes in (3.8), we have

$$H([X]_m) = \mathbb{E}\left[\log \frac{1}{\mu(C_m(2^m[X]_m))}\right]. \tag{A.1}$$

For any $0 < \epsilon < 1$, there exists $m \in \mathbb{N}$, such that $2^{-m} \leq \epsilon < 2^{m-1}$. Then

$$C_m([x]_m 2^m) \subset B_\infty(x, 2^{-m}) \subset B_\infty(x, \epsilon) \tag{A.2}$$

$$\subset B_\infty(x, 2^{-(m-1)}) \subset B_\infty([x]_m, 2^{-(m-2)}). \tag{A.3}$$

As a result of (A.2), we have

$$\mathbb{E}\left[\log \frac{1}{\mu(B_\infty(X, \epsilon))}\right] \leq H([X]_m) \tag{A.4}$$

On the other hand, note that $B_\infty([x]_m, 2^{-(m-2)})$ is a disjoint union of 8^n mesh cubes. By (A.3), we have

$$\mathbb{E}\left[\log \frac{1}{\mu(B_\infty(X, \epsilon))}\right] \geq \mathbb{E}\left[\log \frac{1}{\mu(B_\infty([X]_m, 2^{-(m-2)}))}\right] \tag{A.5}$$

$$\geq H([X]_m) - n \log 8. \tag{A.6}$$

Combining (A.4) and (A.6) yields

$$\frac{H([X]_m) - n \log 8}{m} \leq \frac{\mathbb{E}\left[\log \frac{1}{\mu(B_\infty(X,\epsilon))}\right]}{\log \frac{1}{\epsilon}} \leq \frac{H([X]_m)}{m-1} \tag{A.7}$$

By Lemma 1, sending $\epsilon \downarrow 0$ and $m \to \infty$ yields (2.7). □

A.2 Proofs of Theorem 1

Lemma 28. *For all $p, q \in \mathbb{N}$,*

$$H(\langle X^n \rangle_p) \leq H(\langle X^n \rangle_q) + n \log \left(\left\lceil \frac{p}{q} \right\rceil + 1 \right). \tag{A.8}$$

Proof. For notational conciseness, we only give the proof for the scalar case ($n = 1$).

$$H(\langle X \rangle_p) \leq H(\langle X \rangle_p, \langle X \rangle_q) \tag{A.9}$$
$$= H(\langle X \rangle_p | \langle X \rangle_q) + H(\langle X \rangle_q). \tag{A.10}$$

Note that for any $l \in \mathbb{Z}$,

$$H\left(\langle X \rangle_p \middle| \langle X \rangle_q = \frac{l}{q} \right) = H\left(\langle X \rangle_p \middle| \frac{l}{q} \leq \frac{X}{q} < \frac{l+1}{q} \right) \tag{A.11}$$
$$= H\left(\lfloor pX \rfloor \middle| \frac{pl}{q} \leq pX < \frac{p(l+1)}{q} \right). \tag{A.12}$$

Given $pX \in \left[\frac{pl}{q}, \frac{p(l+1)}{q} \right)$, the range of $\lfloor pX \rfloor$ is upper bounded by $\left\lceil \frac{p}{q} \right\rceil + 1$. Therefore for all $l \in \mathbb{Z}$,

$$H\left(\langle X \rangle_p \middle| \langle X \rangle_q = \frac{l}{q} \right) \leq \log \left(\left\lceil \frac{p}{q} \right\rceil + 1 \right). \tag{A.13}$$

Hence $H(\langle X \rangle_p | \langle X \rangle_q)$ admits the same upper bound and (A.8) holds. □

Lemma 29 ([137, p. 2102])**.** *Let W be an $\mathbb{N}$-valued random variable. Then $H(W) < \infty$ if $\mathbb{E}[\log W] < \infty$.*

Proof of Theorem 1. Using Lemma 28 with $p = m, q = 1$ and $p = 1, q = m$, we have

$$H(\lfloor X^n \rfloor) - \log 2 \leq H(\langle X^n \rangle_m) \leq H(\lfloor X^n \rfloor) + \log m. \tag{A.14}$$

2) ⇒ 1): When $H(\lfloor X^n \rfloor)$ is finite, dividing both sides of (A.14) by $\log m$ and letting $m \to \infty$ results in 1).

1) ⇒ 2): Suppose $H(\lfloor X^n \rfloor) = \infty$. By (A.14), $H(\langle X^n \rangle_m) = \infty$ for every , and 1) fails. This also proves (2.14).

3) ⇔ 2): By the equivalence of norms on $\mathbb{R}^n$, we shall consider the ℓ_∞ ball. Define

$$g(X^n, \epsilon) \triangleq \mathbb{E}\left[\log \frac{1}{P_{X^n}(B_\infty(X^n, \epsilon))} \right], \tag{A.15}$$

which is decreasing in ϵ. Let $2^{-m} \leq \epsilon < 2^{-m+1}$, where $m \in \mathbb{N}$. By (A.4) and (A.6) (or [23, Lemma 2.3]), we have

$$g(X^n, 2^{-m}) \leq H([X^n]_m) \leq g(X^n, 2^{-m}) + n \log 8. \tag{A.16}$$

The equivalence between 3) and 2) then follows from (A.14) and (A.16).

3) $\Leftrightarrow$ 4): By Lemma 32 in Appendix A.5.

4) $\Leftrightarrow$ 5): For all $\mathsf{snr} > \mathsf{snr}' > 0$

$$I(\mathsf{snr}) = I(\mathsf{snr}') + \frac{1}{2}\int_{\mathsf{snr}'}^{\mathsf{snr}} \mathsf{mmse}(X,\gamma)\mathrm{d}\gamma \tag{A.17}$$

$$\leq I(\mathsf{snr}') + \frac{1}{2}\log\frac{\mathsf{snr}}{\mathsf{snr}'}, \tag{A.18}$$

where (A.17) and (A.18) are due to (4.58) and (4.8) respectively.

(2.12) $\Rightarrow$ 2):

$$\mathbb{E}\left[\log(|X|+1)\right] < \infty \tag{A.19}$$

$$\Rightarrow \quad \mathbb{E}\left[\log(\lfloor|X|\rfloor+1)\right] < \infty \tag{A.20}$$

$$\Rightarrow \quad H(\lfloor|X|\rfloor+1) < \infty \tag{A.21}$$

$$\Rightarrow \quad H(\lfloor|X|\rfloor) < \infty \tag{A.22}$$

$$\Rightarrow \quad H(\lfloor X\rfloor) < \infty \tag{A.23}$$

where

- (A.21): by Lemma 29.
- (A.23): by

$$H(\lfloor X\rfloor) \leq H(\lfloor|X|\rfloor) + H\left(\lfloor X\rfloor|\lfloor|X|\rfloor\right) \tag{A.24}$$

$$\leq H(\lfloor|X|\rfloor) + \log 2. \tag{A.25}$$

□

A.3 Proofs of Lemmas 1 – 2

Proof of Lemma 1. Fix any $m \in \mathbb{N}$ and $l \in \mathbb{N}$, such that $2^{l-1} \leq m < 2^l$. By Lemma 28, we have

$$H([X]_{l-1}) \leq H(\langle X\rangle_m) + \log 3, \tag{A.26}$$

$$H(\langle X\rangle_m) \leq H([X]_l) + \log 3. \tag{A.27}$$

Therefore

$$\frac{H([X]_{l-1}) - \log 3}{l} \leq \frac{H(\langle X\rangle_m)}{\log m} \leq \frac{H([X]_l) + \log 3}{l-1} \tag{A.28}$$

and hence (2.16) and (2.17) follow. □

Proof of Lemma 2. Note that

$$H(\langle X\rangle_m) \leq H\left(\frac{\lceil mX\rceil}{m}\right) + H\left(\langle X\rangle_m \middle| \frac{\lceil mX\rceil}{m}\right) \tag{A.29}$$

$$\leq H\left(\frac{\lceil mX\rceil}{m}\right) + \log 2, \tag{A.30}$$

$$H\left(\frac{\lceil mX\rceil}{m}\right) \leq H\left(\langle X\rangle_m\right) + H\left(\frac{\lceil mX\rceil}{m} \middle| \langle X\rangle_m\right) \tag{A.31}$$

$$\leq H\left(\langle X\rangle_m\right) + \log 2. \tag{A.32}$$

The same bound holds for rounding. □

A.4 Proof of Lemma 3

First we present several auxiliary results regarding entropy of quantized random variables:

Lemma 30. *Let U^n and V^n be integer-valued random vectors. If $-B_i \leq U_i - V_i \leq A_i$ almost surely, then*

$$|H(U^n) - H(V^n)| \leq \sum_{i=1}^{n} \log(1 + A_i + B_i). \tag{A.33}$$

Proof. $H(U^n) \leq H(V^n) + H(U^n - V^n|V^n) \leq H(V^n) + \sum_{i=1}^{n} \log(1 + A_i + B_i)$. □

The next lemma shows that the entropy of quantized linear combinations is close to the entropy of linear combinations of quantized random variables, as long as the coefficients are integers. The proof is a simple application of Lemma 30.

Lemma 31. *Let $[\cdot]_m$ be defined in (2.15). For any $a^K \in \mathbb{Z}^K$ and any $X_1^n, \ldots, X_K^n$,*

$$\left| H\left(\sum_{j=1}^{K} a_j [X_i^n]_m\right) - H\left(\left[\sum_{j=1}^{K} a_j X_i^n\right]_m\right)\right| \leq n \log\left(2 + 2\sum_{j=1}^{K} |a_j|\right). \tag{A.34}$$

Proof of Lemma 3. The translation and scale-invariance of $d(\cdot)$ are immediate consequence of the alternative definition of information dimension in (2.8). To prove (2.20) and (2.21), we apply Lemma 31 to obtain

$$|H([X+Y]_m) - H([X]_m + [Y]_m)| \leq \log 6. \tag{A.35}$$

and note that

$$\max\{H([X]_m), H([Y]_m)\} \leq H([X]_m + [Y]_m) \leq H([X]_m) + H([Y]_m). \tag{A.36}$$

Since (2.22) is a simple consequence of independence, it remains to show (2.23). Let U, V, W be integer valued. By the data processing theorem for mutual information, we have $I(V; U+V+W) \le I(V; V+W)$, i.e.,

$$H(U+V+W) - H(V+W) \le H(U+W) - H(W). \tag{A.37}$$

Replacing U, V, W by $[X]_m, [Y]_m, [Z]_m$ respectively and applying Lemma 31, we have

$$H([X+Y+Z]_m) - H([Y+Z]_m) \le H([X+Z]_m) - H([Z]_m) + 9\log 2. \tag{A.38}$$

Dividing both sides by m and sending m to ∞ yield (2.23). □

A.5 A non-asymptotic refinement of Theorem 9

If $H(\lfloor X^n \rfloor) = \infty$, then, in view of Theorem 1, we have $I(X^n, \mathsf{snr}) = \infty$ and $d(X^n) = \infty$, in which case Theorem 9 holds automatically. Next we assume that $H(\lfloor X^n \rfloor) < \infty$, then $I(X^n, \mathsf{snr})$ and $g(X^n, \epsilon)$ are both finite for all snr and ϵ, where $g(X^n, \epsilon)$ is defined in (A.15). Next we prove the following non-asymptotic result, which, in view of the alternative definition of information dimension in (2.7), implies Theorem 9.

Lemma 32. *Let $H(\lfloor X^n \rfloor) < \infty$. For any $\mathsf{snr} > 0$,*

$$\frac{n}{2}\log\frac{6}{\mathrm{e}\pi} \le I(X^n, \mathsf{snr}) - g\left(X^n, 2\sqrt{3\mathsf{snr}^{-1}}\right) \le n\log 3, \tag{A.39}$$

where $g(X^n, \cdot)$ is defined in (A.15).

Proof. First consider the scalar case. Let U be uniformly distributed on $[-\sqrt{3}, \sqrt{3}]$, which has zero mean and unit variance. Then [138]

$$I(X, \mathsf{snr}) \le I(X; \sqrt{\mathsf{snr}}X + U) \le I(X, \mathsf{snr}) + \frac{1}{2}\log\frac{\pi\mathrm{e}}{6}, \tag{A.40}$$

where $\frac{1}{2}\log\frac{\pi\mathrm{e}}{6} = D(P_U \,||\, \mathcal{N}(0,1))$ is the non-Gaussianness of U.

Next we provide alternative lower and upper bounds to $I(X; \sqrt{\mathsf{snr}}X + U)$:

(Lower bound) Let $\epsilon = \frac{1}{\sqrt{\mathsf{snr}}}$. Abbreviate the ℓ_∞ ball $B_\infty(x, \epsilon)$ by $B(x, \epsilon)$. Then the density of $Y = X + \epsilon U$ is given by

$$p_Y(y) = \frac{P_X(B(y, \sqrt{3}\epsilon))}{2\sqrt{3}\epsilon}. \tag{A.41}$$

Then

$$I(X;\sqrt{\mathsf{snr}}X+U) = D(P_{Y|X}\,||\,P_Y\,|\,P_X) \tag{A.42}$$

$$= \mathbb{E}\left[\log\frac{1}{P_X(B(X+\epsilon U,\sqrt{3}\epsilon))}\right] \tag{A.43}$$

$$\geq \mathbb{E}\left[\log\frac{1}{P_X(B(X,2\sqrt{3}\epsilon))}\right], \tag{A.44}$$

where (A.44) is due to $B(X+\epsilon U,\sqrt{3}\epsilon)\subset B(X,2\sqrt{3}\epsilon)$, which follows from $|U|\leq\sqrt{3}$ and the triangle inequality.

(Upper bound) Let Q_Y denote the distribution of $X+3\epsilon U$, whose density is

$$q_Y(y) = \frac{P_X(B(y,3\sqrt{3}\epsilon))}{6\sqrt{3}\epsilon}. \tag{A.45}$$

Then

$$I(X;\sqrt{\mathsf{snr}}X+U) \leq D(P_{Y|X}\,||\,Q_Y\,|\,P_X) \tag{A.46}$$

$$= \mathbb{E}\left[\log\frac{3}{P_X(B(X+\epsilon U,3\sqrt{3}\epsilon))}\right] \tag{A.47}$$

$$\leq \mathbb{E}\left[\log\frac{1}{P_X(B(X,2\sqrt{3}\epsilon))}\right]+\log 3,, \tag{A.48}$$

where (A.48) is due to $B(X+\epsilon U,3\sqrt{3}\epsilon)\supset B(X,2\sqrt{3}\epsilon)$.

Assembling (A.40), (A.44) and (A.48) yields (A.39). The proof for the vector case follows from entirely analogous fashion. □

A.6 Proof of Lemma 5

Proof. Let U be an open interval (a,b). Then $F_j(U)=(w_j+ra,w_j+rb)$ and

$$\bigcup_j F_j(U)\subset U \Leftrightarrow \min_j w_j+ra\geq a \text{ and } \max_j w_j+rb\leq b. \tag{A.49}$$

and for $i\neq j$,

$$F_i(U)\cap F_j(U)=\emptyset \Leftrightarrow |w_i-w_j|\geq r(b-a). \tag{A.50}$$

Therefore, there exist $a<b$ that satisfy (A.49) and (A.50) if and only if (2.82) holds. □

A.7 Proof of Lemma 6

Proof. Let X be of the form in (2.80). Let $M = \max_{1\le j\le m} w_j$ and $m_k = \lfloor r^{-k} \rfloor$ where $k \in \mathbb{N}$. Then $|X| \le \frac{rM}{1-r}$. By Lemma 1,

$$d_\alpha(X) = \lim_{k\to\infty} \frac{H_\alpha(\lfloor m_k X \rfloor)}{k \log \frac{1}{r}}. \tag{A.51}$$

Since

$$\left| \lfloor m_k X \rfloor - \sum_{i=0}^{k-1} r^{-i} W_{k-i} \right| \le \frac{2rM}{1-r}, \tag{A.52}$$

applying Lemma 30 to (A.52) and substituting into (A.51) yield

$$d_\alpha(X) \le \lim_{k\to\infty} \frac{H_\alpha(\sum_{i=0}^{k-1} r^{-i} W_{k-i})}{k \log \frac{1}{r}} \le \frac{H_\alpha(P)}{\log \frac{1}{r}}, \tag{A.53}$$

where we have used $H_\alpha(\sum_{i=0}^{k-1} r^{-i} W_{k-i}) \le H_\alpha(W^k) = kH_\alpha(P)$.

□

A.8 Proof of Theorem 12

First we present several auxiliary results from sumset theory for entropy.

Lemma 33 (Submodularity [139, Lemma A.2])**.** *If X_0, X_1, X_2, X_{12} are random variables such that $X_0 = f(X_1) = g(X_2)$ and $X_{12} = h(X_1, X_2)$, where f, g, h are deterministic functions. Then*

$$H(X_0) + H(X_{12}) \le H(X_1) + H(X_2). \tag{A.54}$$

In the sequel, G denotes an arbitrary abelian group.

Lemma 34 (Ruzsa's triangle inequality [139, Theorem 1.7])**.** *Let X and Y be G-valued random variables. Define the* Ruzsa distance[1] *as*

$$\Delta(X, Y) \triangleq H(X' - Y') - \frac{1}{2}H(X') - \frac{1}{2}H(Y'), \tag{A.55}$$

where X' and Y' are independent copies of X and Y. Then

$$\Delta(X, Z) \le \Delta(X, Y) + \Delta(Y, Z), \tag{A.56}$$

that is,

$$H(X - Z) \le H(X - Y) + H(Y - Z) - H(Y). \tag{A.57}$$

[1] $\Delta(\cdot,\cdot)$ is not a metric since $\Delta(X, X) > 0$ unless X is deterministic.

The following entropic analogue of Plünnecke-Ruzsa's inequality [140] is implicitly implied in [141, Proposition 1.3], as pointed out by Tao [142, p. 203].

Lemma 35. *Let $X, Y_1, \ldots, Y_m$ be independent G-valued random variables. Then*

$$H(X + Y_1 + \ldots + Y_m) \leq H(X) + \sum_{i=1}^{m} [H(X + Y_i) - H(X)]. \tag{A.58}$$

The following entropy inequality for group-valued random variables is based on results from [143, 144, 145], whose proof can be found in [146, Section II-E].

Lemma 36.

$$\frac{1}{2} \leq \frac{\Delta(X, X)}{\Delta(X, -X)} \leq 2. \tag{A.59}$$

Lemma 37 (Sum-difference inequality [139, (2.2)])**.** *Let X and Y be independent G-valued random variables. Then $\Delta(X, Y) \leq 3\Delta(X, -Y)$, i.e.,*

$$H(X - Y) \leq 3H(X + Y) - H(X) - H(Y). \tag{A.60}$$

The foregoing entropy inequalities deal with linear combinations of independent random variables with coefficients being ± 1. In Theorem 50 we present a new upper bound on the entropy of the linear combination of independent random variables with general integer coefficients. The following two lemmas lie at the core of the proof, which allow us to reduce the case from p to either $p \pm 1$ (additively) or $\frac{p}{2}$ (multiplicatively).

Lemma 38 (Additive descent)**.** *Let X, Y and Z be independent G-valued random variables. Let $p, r \in \mathbb{Z} \backslash \{0\}$. Then*

$$H(pX + Y) \leq H((p - r)X + rZ + Y) + H(X - Z) - H(Z). \tag{A.61}$$

Proof. Applying Lemma 33 with $X_1 = ((p - r)X + rZ + Y, X - Z)$, $X_2 = (X, Y)$, $X_0 = pX + Y$ and $X_{12} = (X, Y, Z)$ yields (A.61). □

Lemma 39 (Multiplicative descent)**.** *Let X, X' and Y be independent G-valued random variables where X' has the same distribution as X. Let $p \in \mathbb{Z} \backslash \{0\}$. Then*

$$H(pX + Y) \leq H\left(\frac{p}{2}X + Y\right) + H(2X - X') - H(X). \tag{A.62}$$

Proof. By (A.57), we have

$$H(pX+Y) \le H\left(\frac{p}{2}X+Y\right) + H\left(pX - \frac{p}{2}X'\right) - H\left(\frac{p}{2}X\right). \tag{A.63}$$

$$= H\left(\frac{p}{2}X+Y\right) + H(2X'-X) - H(X). \tag{A.64}$$

□

Theorem 50. *Let X and Y be independent G-valued random variables. Let $p, q \in \mathbb{Z}\backslash\{0\}$. Then*

$$\begin{aligned} &H(pX+qY) - H(X+Y) \\ &\le 9(l_p+l_q)H(X+Y) - (5l_p+4l_q)H(Y) - (5l_q+4l_p)H(X), \end{aligned} \tag{A.65}$$

where $l_p \triangleq 1 + \lfloor \log|p| \rfloor$.

It is interesting to compare Theorem 50 and its proof techniques to its additive-combinatorial counterpart [147, Theorem 1.3], which gives an upper bound on the cardinality of the sum of dilated subsets.

Proof of Theorem 50. Let $\{X_k\}$ and $\{Y_k\}$ are sequences of independent copies of X and Y respectively. By (A.57), we have

$$H(pX+qY) \le H(pX - X_1) + H(X+qY) - H(X) \tag{A.66}$$

$$\le H(pX+Y) + H(qY+X) + H(X+Y) - H(X) - H(Y). \tag{A.67}$$

Next we upper bound $H(pX+Y)$. First we assume that p is positive. If p is even, applying (A.62) yields

$$H(pX+Y) \le H\left(\frac{p}{2}X+Y\right) + H(2X - X_1) - H(X). \tag{A.68}$$

If p is odd, applying (A.61) with $r=1$ then applying (A.62), we have

$$\begin{aligned} &H(pX+Y) \\ &\le H((p-1)X + Y_1 + Y) + H(X-Y) - H(Y) && \text{(A.69)} \\ &\le H\left(\frac{p-1}{2}X + Y_1 + Y\right) + H(X-Y) - H(Y) + H(2X - X_1) - H(X) && \text{(A.70)} \\ &\le H\left(\frac{p-1}{2}X + Y_1 + Y\right) + 2H(X-Y) + 2H(X+Y) - 3H(Y) - H(X), && \text{(A.71)} \end{aligned}$$

where the last equality holds because

$$H(2X - X_1) \le H(X+Y) + H(X - X_1 - Y) - H(Y) \tag{A.72}$$

$$\le 2H(X+Y) + H(X-Y) - 2H(Y), \tag{A.73}$$

where (A.72) and (A.73) are due to (A.61) and Lemma 35 respectively.

Combining (A.68) and (A.71) gives the following upper bound

$$\begin{aligned}
&H(pX+Y)\\
&\le H\left(\left\lfloor\frac{p}{2}\right\rfloor X+Y_1+Y\right)+2H(X-Y)+2H(X+Y)-3H(Y)-H(X) \qquad \text{(A.74)}\\
&\le H\left(\left\lfloor\frac{p}{2}\right\rfloor X+Y_1+Y\right)+8H(X+Y)-5H(Y)-3H(X), \qquad \text{(A.75)}
\end{aligned}$$

where (A.75) is due to Lemma 37. Repeating the above procedure with p replaced by $\left\lfloor\frac{p}{2}\right\rfloor$, we have

$$\begin{aligned}
H(pX+Y) &\le H\left(X+\sum_{i=1}^{l_p} Y_i\right)+l_p(8H(X+Y)-5H(Y)-3H(X)) \qquad \text{(A.76)}\\
&\le H(X)+l_p(9H(X+Y)-5H(Y)-4H(X)), \qquad \text{(A.77)}
\end{aligned}$$

where (A.77) is due to Lemma 35. For negative p, the above proof remains valid if we apply (A.61) with $r=-1$ instead of 1. Then (A.77) holds with p replaced by $-p$. Similarly, we have

$$H(qY+X)\le H(Y)+l_q(9H(X+Y)-5H(X)-4H(Y)). \qquad \text{(A.78)}$$

Assembling (A.77) – (A.78) with (A.67), we obtain (A.65). □

Finally, we prove Theorem 12 by applying Theorem 50.

Proof. First we assume that $p=p'=1$. Substituting $[X]_m$ and $[Y]_m$ into (A.65) and applying Lemma 30, we have

$$\begin{aligned}
&H(p[X]_m+q[Y]_m)-H([X]_m+[Y]_m)\\
&\le 9(l_p+l_q)H([X]_m+[Y]_m)-(5l_p+4l_q)H([Y]_m)-(5l_q+4l_p)H([X]_m)+O(1).
\end{aligned} \qquad \text{(A.79)}$$

Dividing both sides by m and sending $m\to\infty$ yield

$$\begin{aligned}
&d(pX+qY)-d(X+Y)\\
&\le 9(l_p+l_q)d(X+Y)-(5l_p+4l_q)d(Y)-(5l_q+4l_p)d(X). \qquad \text{(A.80)}
\end{aligned}$$

Then

$$\begin{aligned}
&18(l_p+l_q)[d(pX+qY)-d(X+Y)]\\
&=[(5l_p+4l_q)d(Y)-(5l_q+4l_p)d(X)-9(l_p+l_q)d(X+Y)]\\
&\quad+(5l_p+4l_q)[d(pX+qY)-d(Y)-d(X+Y)]\\
&\quad+(5l_q+4l_p)[d(pX+qY)-d(X)-d(X+Y)]+9(l_p+l_q)d(pX+qY) \qquad \text{(A.81)}\\
&\le -[d(pX+qY)-d(X+Y)]+9(l_p+l_q)d(pX+qY), \qquad \text{(A.82)}
\end{aligned}$$

where (A.82) is due to (A.80), (2.21) and (2.11). Therefore

$$d(pX+qY)-d(X+Y) \leq \frac{9(l_p+l_q)}{18(l_p+l_q)+1} = \frac{1}{2} - \frac{1}{36(l_p+l_q)+2}, \tag{A.83}$$

which is the desired (2.96).

Finally, by the scale-invariance of information dimension, setting $Z = \frac{Y}{pp'}$ gives

$$d(pX+qY)-d(p'X+q'Y) = d(X+p'qZ)-d(X+pq'Z) \tag{A.84}$$

$$\leq \frac{1}{2} - \epsilon(p'q, pq'). \tag{A.85}$$

□

Remark 32. For the special case of $p=2$ and $q=1$, we proceed as follows:

$$H(2X+Y) \leq \Delta(X,X) + H(X+X'+Y) \tag{A.86}$$

$$\leq \Delta(X,X) + 2H(X+Y) - H(Y). \tag{A.87}$$

$$\leq 4H(X+Y) - 2H(Y) - H(X). \tag{A.88}$$

where

- (A.86): by (A.61);
- (A.87): by Lemma 35;
- (A.88): by triangle inequality $\Delta(X,X) \leq 2\Delta(X,-Y)$.

By (A.88), we have $d(2X+Y)-d(X+Y) \leq 3d(X+Y)-d(X)-2d(Y)$. Consequently, the same manipulation in (A.81) – (A.83) gives

$$d(2X+Y)-d(X+Y) \leq \frac{3}{7}. \tag{A.89}$$

Similarly for $p=1$ and $q=-1$, from Lemma 37, we have:

$$d(X-Y) \leq 3d(X+Y)-d(X)-d(Y), \tag{A.90}$$

which implies

$$d(X-Y)-d(X+Y) \leq \frac{2}{5}. \tag{A.91}$$

A.9 Proof of Theorem 13

The following basic result from the theory of rational approximation and continued fractions ([148, 149]) is instrumental in proving Theorem 13. For every irrational x, there exists a unique $x_0 \in \mathbb{Z}$ and a unique sequence $\{x_k\} \subset \mathbb{N}$, such that x is given

by the infinite continued fraction

$$x = x_0 + \cfrac{1}{x_1 + \cfrac{1}{x_2 + \ddots}} \tag{A.92}$$

The rational number

$$\frac{p_k(x)}{q_k(x)} \triangleq x_0 + \cfrac{1}{x_1 + \cfrac{1}{x_2 + \cfrac{1}{\ddots + \cfrac{1}{x_k}}}}, \tag{A.93}$$

is called the kth-*convergent* of x. Convergents provide the best rational approximation to irrationals [149].

Lemma 40. *Let x be an irrational real number and N a positive integer. Define*

$$E_N(x) = \left\{ (p,q) \in \mathbb{Z}^2 : |p - qx| \le \frac{1}{2q}, |p|, |q| \le N \right\}. \tag{A.94}$$

Then

$$|E_N(x)| \le 1 + 2\log N. \tag{A.95}$$

Proof. By [149, Theorem 19, p. 30], $\frac{p}{q}$ is necessarily a convergent of x. By [149, Theorem 12, p. 13], for any irrational x, the sequence $\{q_k(x)\}$ formed by denominators of successive convergents grows exponentially: $q_k(x) \ge 2^{\frac{k-1}{2}}$ for all $k \ge 1$,[2] which implies the desired (A.95). □

The next lemma shows that convolutions of homogeneous self-similar measures with common similarity ratio are also self-similar (see also [48, p. 1341] and the references therein), but with possible overlaps which might violate the open set conditions. Therefore the formula in (2.85) does not necessarily apply to convolutions.

Lemma 41. *Let X and X' be independently distributed according to μ and μ' respectively, which are homogeneous self-similar with identical similarity ratio r. Then for any $a, a' \in \mathbb{R}$, the distribution of $aX + a'X'$ is also homogeneous self-similar with similarity ratio r.*

Proof. According to Remark 3, X and X' satisfy $X \overset{\mathrm{D}}{=} rX + W$ and $X' \overset{\mathrm{D}}{=} rX' + W'$, respectively, where W and W' are both finitely valued and $\{X, W, X', W'\}$ are independent. Therefore $aX + a'X' \overset{\mathrm{D}}{=} r(aX + a'X') + aW + a'W'$, which implies the desired conclusion. □

[2] In fact, $\lim_{n\to\infty} q_n(x)^{\frac{1}{n}} = \gamma$ for Lebesgue a.e. x, where $\gamma = \mathrm{e}^{\frac{\pi^2 \log \mathrm{e}}{12}}$ is the Khinchin-Lévy constant [149, p. 66].

Proof of Theorem 13. Let $\lambda \triangleq \frac{p'q}{pq'}$ be irrational. In view of the scale-invariance of information dimension, to establish (2.98), it is equivalent to show

$$\sup_{X \perp\!\!\!\perp Y} d(X + \lambda Y) - d(X + Y) = \frac{1}{2}. \tag{A.96}$$

Next we construct a sequence of distributions $\{\mu_N\}$ such that for X and Y are independent and both distributed according to μ_N, we have

$$d(X + \lambda Y) - d(X + Y) \geq \frac{1}{2} - O\left(\frac{\log \log N}{\log N}\right), \quad N \to \infty, \tag{A.97}$$

which, in view of (2.95), implies the desired (A.96).

Let $N \geq 2$ be an integer. Denote by μ_N the invariant measure generated by the IFS (2.73) with $r_j = \frac{1}{N^2}, w_j = N(j-1)$ and $p_j = \frac{1}{N}$, $j = 1, \ldots, N$. Then μ_N is homogeneous self-similar and satisfies the open set condition (in fact the strong separation condition), in view of Lemma 5. By Theorem 11, $d(\mu_N) = \frac{1}{2}$ for all N.

Alternatively, μ_N can be defined by the N-ary expansions as follows: Let X be distributed according to μ_N. Then the even digits in the N-ary expansion of X are independent and equiprobable, while the odd digits are frozen to be zero. In view of the discussion in Section 2.5, the information dimension of μ_N is the normalized entropy rate of the digits, which is equal to $\frac{1}{2}$.

Let X and Y be independently distributed according to μ_N. Then

$$X = \sum_{i \geq 1} X_i N^{-2i}, \tag{A.98}$$

$$Y = \sum_{i \geq 1} Y_i N^{-2i}, \tag{A.99}$$

where $\{X_i\}$ and $\{Y_i\}$ are both i.i.d. and equiprobable on $N\{0, \ldots, N-1\}$. Then $X + Y$ also has a homogeneous self-similar with similarity ratio N^{-2} and satisfies the open set condition. Applying Theorem 11 yields

$$d(X + Y) = \frac{H(X_1 + Y_1)}{2 \log N} \leq \frac{\log(2N - 1)}{2 \log N}. \tag{A.100}$$

Next we prove that $d(X + \lambda Y)$ is close to one. By Lemma 41, $X + \lambda Y$ is self-similar homogeneous with similarity ratio $\frac{1}{N^2}$. By Theorem 10, $d(X + \lambda Y)$ exists. Moreover,

$$d(X + \lambda Y) \geq \frac{H(\langle X + \lambda Y \rangle_{N^2}) - 10}{2 \log N}. \tag{A.101}$$

Denote $X' = \langle X \rangle_{N^2}$ and $Y' = \langle Y \rangle_{N^2}$. Let $N' \triangleq 4N^2$. To further lower bound the entropy in (A.101), note that

$$\begin{aligned} H(\langle X + \lambda Y \rangle_{N^2}) &\geq H(\langle X + \lambda Y \rangle_{N'}) - \log 5 && \text{(A.102)} \\ &\geq H(\langle X' + \lambda Y' \rangle_{N'}) - \log 5 - \log(3 + 2|\lambda|) && \text{(A.103)} \\ &= H(X') + H(Y') - H(X', Y' | \langle X' + \lambda Y' \rangle_{N'}) \\ &\quad - \log 5 - \log(3 + 2|\lambda|), && \text{(A.104)} \end{aligned}$$

where (A.102), (A.103) and (A.104) follow from Lemma 28, Lemma 30 and the independence of X' and Y', respectively.

Next we upper bound the conditional entropy in (A.104) using Lemma 40. By the definition of μ_N, X' is valued in $\frac{1}{N}\{0, \ldots, N-1\}$. Fix $k \in \mathbb{Z}$. Conditioned on the event $\left\{\langle X' + \lambda Y' \rangle_{N'} = \frac{k}{N'}\right\}$, define the set of all possible realizations of (X', Y'):

$$C_k = \left\{ \left(\frac{x}{N}, \frac{y}{N}\right) : \left| \frac{x}{N} + \lambda \frac{y}{N} - \frac{k}{N'} \right| < \frac{1}{N'}, x, y \in \mathbb{Z}_+, x < N, y < N \right\}. \tag{A.105}$$

To upper bound $|C_k|$, note that if $(\frac{x}{N}, \frac{y}{N})$ and $(\frac{u'}{N}, \frac{v'}{N})$ both belong to C_k, then

$$\left| \frac{x - x'}{N} + \lambda \frac{y - y'}{N} \right| \leq \frac{2}{N'}, \tag{A.106}$$

which, in view of $|y - y'| < N$ and $N' = 4N^2$, implies that

$$|x - x' + \lambda(y - y')| \leq \frac{2N}{N'} \leq \frac{1}{2|y - y'|}. \tag{A.107}$$

In view of (A.94), (A.107) implies that $(x - x', y - y') \in E_N(\lambda)$. In other words,

$$C_k - C_k \subset \frac{1}{N} E_N(\lambda). \tag{A.108}$$

By assumption, λ is irrational. Consequently, by Lemma 40,

$$|C_k| \leq |E_N(\lambda)| \leq 1 + 2 \log N. \tag{A.109}$$

Since (A.109) holds for all k, we have

$$H(X', Y' | \langle X' + \lambda Y' \rangle_{N'}) \leq \log(1 + 2 \log N). \tag{A.110}$$

Plugging (A.110) into (A.104) yields

$$H(\langle X + \lambda Y \rangle_{N^2}) \geq H(\langle X \rangle_{N^2}) + H(\langle Y \rangle_{N^2}) - O(\log \log N). \tag{A.111}$$

Dividing both sides of (A.111) by $2\log N$ and substituting into (A.101), we have

$$d(X+\lambda Y) \geq \frac{H(\langle X\rangle_{N^2})}{2\log N} + \frac{H(\langle Y\rangle_{N^2})}{2\log N} - O\left(\frac{\log\log N}{\log N}\right) \tag{A.112}$$

$$= 1 - O\left(\frac{\log\log N}{\log N}\right), \tag{A.113}$$

where we have used $H(\langle X\rangle_{N^2}) = H(\langle Y\rangle_{N^2}) = \log N$. Combining (A.100) and (A.113), we obtain the desired (A.96), hence completing the proof of Theorem 13. □

Appendix B

Proofs in Chapter 4

B.1 Proof of Theorem 19

Proof. Invariance of the MMSE functional under translation is obvious. Hence for any $\alpha, \beta \neq 0$,

$$
\begin{aligned}
&\mathsf{mmse}(\alpha X + \gamma, \beta N + \eta, \mathsf{snr}) \\
&= \mathsf{mmse}(\alpha X, \beta N, \mathsf{snr}) && \text{(B.1)}\\
&= \mathsf{mmse}(\alpha X | \alpha\sqrt{\mathsf{snr}} X + \beta N) && \text{(B.2)}\\
&= |\alpha|^2 \mathsf{mmse}\left(X \middle| \frac{|\alpha|}{|\beta|}\sqrt{\mathsf{snr}} X + \operatorname{sgn}(\alpha\beta) N\right) && \text{(B.3)}\\
&= |\alpha|^2 \mathsf{mmse}\left(X, \operatorname{sgn}(\alpha\beta) N, \frac{|\alpha|}{|\beta|}\sqrt{\mathsf{snr}}\right) && \text{(B.4)}\\
&= |\alpha|^2 \mathsf{mmse}\left(\operatorname{sgn}(\alpha\beta) X, N, \frac{|\alpha|}{|\beta|}\sqrt{\mathsf{snr}}\right). && \text{(B.5)}
\end{aligned}
$$

Therefore

$$
\begin{aligned}
&\underline{\mathscr{D}}(\alpha X + \gamma, \beta N + \eta) \\
&= \liminf_{\mathsf{snr}\to\infty} \frac{\mathsf{snr} \cdot \mathsf{mmse}(\alpha X + \gamma, \beta N + \eta, \mathsf{snr})}{\mathsf{var}(\beta N + \eta)} && \text{(B.6)}\\
&= \liminf_{\mathsf{snr}\to\infty} \frac{\mathsf{snr} \cdot |\alpha|^2 \mathsf{mmse}\left(X, \operatorname{sgn}(\alpha\beta) N, \frac{|\alpha|}{|\beta|}\sqrt{\mathsf{snr}}\right)}{|\beta|^2 \mathsf{var} N} && \text{(B.7)}\\
&= \underline{\mathscr{D}}(X, \operatorname{sgn}(\alpha\beta) N) && \text{(B.8)}\\
&= \underline{\mathscr{D}}(\operatorname{sgn}(\alpha\beta) X, N), && \text{(B.9)}
\end{aligned}
$$

where (B.8) and (B.9) follow from (B.4) and (B.5) respectively. The claims in Theorem 19 are special cases of (B.8) and (B.9). The proof for $\overline{\mathscr{D}}$ follows analogously. □

B.2 Calculation of (4.17), (4.18) and (4.93)

In this appendix we compute $\mathsf{mmse}(X, N, \mathsf{snr})$ for three different pairs of (X, N). First we show (4.18), where X is uniformly distributed in $[0, 1]$ and N has the density in (4.16) with $\alpha = 3$. Let $\epsilon = \frac{1}{\sqrt{\mathsf{snr}}}$. Then $\mathbb{E}[X|X + \epsilon N = y] = \frac{q_1}{q_0}(y)$, where

$$q_0(y) = \frac{1}{\epsilon}\mathbb{E}\left[f_N\left(\frac{y-x}{\epsilon}\right)\right] \tag{B.10}$$

$$= \begin{cases} 1 - \frac{\epsilon^2}{y^2} & \epsilon < y < 1 + \epsilon \\ \frac{\epsilon^2(-1+2y)}{(-1+y)^2 y^2} & y > 1 + \epsilon. \end{cases} \tag{B.11}$$

and

$$q_1(y) = \frac{1}{\epsilon}\mathbb{E}\left[X f_N\left(\frac{y-x}{\epsilon}\right)\right] \tag{B.12}$$

$$= \begin{cases} \frac{2\epsilon^{\frac{3}{2}}}{\sqrt{y}} + y - 3\epsilon & \epsilon < y < 1 + \epsilon \\ \frac{2\epsilon^{\frac{3}{2}}}{\sqrt{y}} + \frac{(3-2y)\epsilon^{\frac{3}{2}}}{(y-1)^{\frac{3}{2}}} & y > 1 + \epsilon. \end{cases} \tag{B.13}$$

Then

$$\begin{aligned} &\mathsf{mmse}(X, N, \mathsf{snr}) \\ &= \mathbb{E}[X^2] - \mathbb{E}\left[(\mathbb{E}[X|X + \epsilon N])^2\right] && \text{(B.14)} \\ &= \mathbb{E}[X^2] - \int_0^\infty \frac{q_1^2(y)}{q_0(y)} \mathrm{d}y && \text{(B.15)} \\ &= 2\epsilon^2 \left[\log\left(1 + \frac{1}{2\epsilon}\right) - 2 + 8\epsilon \coth^{-1}(1 + 4\epsilon)\right], && \text{(B.16)} \end{aligned}$$

where we have used $\mathbb{E}[X^2] = \frac{1}{3}$. Taking Taylor expansion on (B.16) at $\epsilon = 0$ yields (4.18). For $\alpha = 2$, (4.17) can be shown in similar fashion.

To show (4.93), where X and N are both uniformly distributed in $[0, 1]$, we note that

$$q_0(y) = \frac{1}{\epsilon}[\min\{y, 1\} - (y - \epsilon)^+] \tag{B.17}$$

$$q_1(y) = \frac{1}{2\epsilon}[\min\{y^2, 1\} - ((y - \epsilon)^+)^2], \tag{B.18}$$

where $(x)^+ \triangleq \max\{x, 0\}$. Then (4.93) can be obtained using (B.15).

B.3 Proof of Theorem 21

The outline of the proof is as follows. From

$$\mathsf{mmse}(X, \mathsf{snr}|U) \leq \mathsf{mmse}(X, \mathsf{snr}), \tag{B.19}$$

we immediately obtain the inequalities in (4.61) and (4.62). Next we prove that equalities hold if U is discrete. Let $\mathcal{U} = \{u_i : i = 1, \ldots, n\}$ denote the alphabet of U with $n \in \mathbb{N} \cup \{\infty\}$, $\alpha_i = \mathbb{P}\{U = u_i\}$. Denote by μ_i the distribution of X given $U = u_i$. Then the distribution of X is given by the following mixture:

$$\mu = \sum_{i=1}^{n} \alpha_i \mu_i. \tag{B.20}$$

Our goal is to establish

$$\overline{\mathscr{D}}(\mu) = \sum_{i=1}^{n} \alpha_i \overline{\mathscr{D}}(\mu_i), \tag{B.21}$$

$$\underline{\mathscr{D}}(\mu) = \sum_{i=1}^{n} \alpha_i \underline{\mathscr{D}}(\mu_i). \tag{B.22}$$

After recalling an important lemma due to Doob in Appendix B.3.1, we decompose the proof of (B.21) and (B.22) into four steps, which are presented in Appendices B.3.2 – B.3.5 respectively:

1. We prove the special case of $n = 2$ and $\mu_1 \perp \mu_2$;
2. We consider $\mu_1 \ll \mu_2$;
3. Via the Hahn-Lebesgue decomposition and induction on n, the conclusion is extended to any finite mixture;
4. We prove the most general case of countable mixture ($n = \infty$).

B.3.1 Doob's relative limit theorem

The following lemma is a combination of [35, Exercise 2.9, p.243] and [150, Theorem 1.6.2, p.40]:

Lemma 42. *Let μ and ν be two Radon measures on $\mathbb{R}^n$. Define the density of μ with respect to ν by*

$$\frac{\mathrm{D}\mu}{\mathrm{D}\nu}(x) = \lim_{\epsilon \downarrow 0} \frac{\mu(B(x, \epsilon))}{\nu(B(x, \epsilon))}, \tag{B.23}$$

where $B(x, \epsilon)$ denotes the open ball of radius ϵ centered at x. If $\mu \perp \nu$, then

$$\frac{\mathrm{D}\mu}{\mathrm{D}\nu} = 0, \quad \nu\text{-a.e.} \tag{B.24}$$

If $\mu \ll \nu$, then

$$\frac{\mathrm{D}\mu}{\mathrm{D}\nu} = \frac{\mathrm{d}\mu}{\mathrm{d}\nu}, \quad \mu\text{-a.e.} \tag{B.25}$$

The Lebesgue-Besicovitch differentiation theorem is a direct consequence of Lemma 42:

Lemma 43 ([150, Corollary 1.7.1])**.** *Let ν be a Radon measure on $\mathbb{R}^n$ and $g \in L^1_{\mathrm{loc}}(\mathbb{R}^n, \nu)$. Then*

$$\lim_{\epsilon \downarrow 0} \frac{1}{\nu(B(x,\epsilon))} \int_{B(x,\epsilon)} |g(y) - g(x)| \nu(\mathrm{d}y) = 0 \tag{B.26}$$

holds for ν-a.e. $x \in \mathbb{R}^n$.

It is instructive to reformulate Lemma 42 in a probabilistic context: Suppose μ and ν are probability measures and X and Z are random variables distributed according to μ and ν respectively. Let N be uniformly distributed in $[-1, 1]$ and independent of $\{X, Z\}$. Then $X + \epsilon N$ has the following density:

$$f_{X+\epsilon N}(x) = \frac{1}{2\epsilon} \mu(B(x,\epsilon)), \tag{B.27}$$

hence the density of μ with respect to ν can be written as

$$\frac{\mathrm{D}\mu}{\mathrm{D}\nu}(x) = \lim_{\epsilon \downarrow 0} \frac{f_{X+\epsilon N}(x)}{f_{Z+\epsilon N}(x)}. \tag{B.28}$$

A natural question is whether (B.28) still holds if N has a non-uniform distribution. In [151, Theorem 4.1], Doob gave a sufficient condition for this to be true, which is satisfied in particular by Gaussian-distributed N [151, Theorem 5.2]:

Lemma 44. *For any $z \in \mathbb{R}$, let $\varphi_z(\cdot) = \varphi(z + \cdot)$. Under the assumption of Lemma 42, if*

$$\int \varphi_z\left(\frac{y-x}{\epsilon}\right) \mu(\mathrm{d}y), \int \varphi_z\left(\frac{y-x}{\epsilon}\right) \nu(\mathrm{d}y) \tag{B.29}$$

are finite for all $\epsilon > 0$ and $x \in \mathbb{R}$, then

$$\frac{\mathrm{D}\mu}{\mathrm{D}\nu}(x) = \lim_{\epsilon \downarrow 0} \frac{\int \varphi_z\left(\frac{y-x}{\epsilon}\right) \mu(\mathrm{d}y)}{\int \varphi_z\left(\frac{y-x}{\epsilon}\right) \nu(\mathrm{d}y)} \tag{B.30}$$

holds for ν-a.e. x.

Consequently we have counterparts of Lemmas 42 and 43:

Lemma 45. *Under the condition of Lemma 42, if $\mu \perp \nu$, then*

$$\lim_{\epsilon \downarrow 0} \frac{\int \varphi_z\left(\frac{y-x}{\epsilon}\right) \mu(\mathrm{d}y)}{\int \varphi_z\left(\frac{y-x}{\epsilon}\right) \nu(\mathrm{d}y)} = 0, \quad \nu\text{-a.e.} \tag{B.31}$$

If $\mu \ll \nu$, then

$$\lim_{\epsilon \downarrow 0} \frac{\int \varphi_z\left(\frac{y-x}{\epsilon}\right)\mu(\mathrm{d}y)}{\int \varphi_z\left(\frac{y-x}{\epsilon}\right)\nu(\mathrm{d}y)} = \frac{\mathrm{d}\mu}{\mathrm{d}\nu}, \quad \mu\text{-a.e.} \tag{B.32}$$

Lemma 46. *Under the condition of Lemma 43,*

$$\lim_{\epsilon \downarrow 0} \frac{\int |g(y)-g(x)|\varphi_z\left(\frac{y-x}{\epsilon}\right)\nu(\mathrm{d}y)}{\int \varphi_z\left(\frac{y-x}{\epsilon}\right)\nu(\mathrm{d}y)} = 0 \tag{B.33}$$

holds for ν-a.e. x.

B.3.2 Mixture of two mutually singular measures

We first present a lemma which enables us to truncate the input or the noise. The point of this result is that the error term depends only on the truncation threshold but not on the observation.

Lemma 47. *For $K > 0$ define $X_K = X\mathbf{1}_{\{|X|\le K\}}$ and $\bar{X}_K = X - X_K$. Then for all Y,*

$$|\mathsf{mmse}(X_K|Y) - \mathsf{mmse}(X|Y)| \le 3\left\|\bar{X}_K\right\|_2 \|X\|_2 . \tag{B.34}$$

Proof.

$$\begin{aligned}
\mathsf{mmse}(X_K|Y) &= \|X_K - \mathbb{E}[X_K|Y]\|_2^2 && \text{(B.35)}\\
&\le \left(\|X - \mathbb{E}[X|Y]\|_2 + \left\|\bar{X}_K - \mathbb{E}[\bar{X}_K|Y]\right\|_2\right)^2 && \text{(B.36)}\\
&\le \left(\mathsf{mmse}(X|Y)^{\frac{1}{2}} + \sqrt{\mathsf{var}\bar{X}_K}\right)^2 && \text{(B.37)}\\
&= \mathsf{mmse}(X|Y) + \mathsf{var}\bar{X}_K + 2\sqrt{\mathsf{var}\bar{X}_K\ \mathsf{mmse}(X|Y)} && \text{(B.38)}\\
&\le \mathsf{mmse}(X|Y) + 3\left\|\bar{X}_K\right\|_2 \|X\|_2 . && \text{(B.39)}
\end{aligned}$$

The other direction of (B.34) follows in entirely analogous fashion. □

Let X_i be a random variable with distribution μ_i, for $i = 1, 2$. Let U be a random variable independent of $\{X_1, X_2\}$, taking values on $\{1, 2\}$ with probability $0 < \alpha_1 < 1$ and $\alpha_2 = 1 - \alpha_2$ respectively. Then the distribution of X_U is

$$\mu = \alpha_1\mu_1 + \alpha_2\mu_2. \tag{B.40}$$

Fixing $M > 0$, we define

$$g_{u,\epsilon}(y) = \mathbb{E}\left[\varphi\left(\frac{y - X_u}{\epsilon}\right)\right], \tag{B.41}$$

$$f_{u,\epsilon}(y) = \mathbb{E}\left[\frac{y - X_u}{\epsilon}\varphi\left(\frac{y - X_u}{\epsilon}\right)\mathbf{1}_{\{|\frac{y-X_u}{\epsilon}|\le M\}}\right] \tag{B.42}$$

and

$$g_\epsilon = \alpha_1 g_{1,\epsilon} + \alpha_2 g_{2,\epsilon}, \tag{B.43}$$
$$f_\epsilon = \alpha_1 f_{1,\epsilon} + \alpha_2 f_{2,\epsilon}. \tag{B.44}$$

Then the densities of $Y_u = X_u + \epsilon N_{\mathsf{G}}$ and $Y_U = X_U + \epsilon N_{\mathsf{G}}$ are respectively given by

$$q_{u,\epsilon}(y) = \frac{1}{\epsilon} g_{u,\epsilon}(y) \tag{B.45}$$
$$q_\epsilon(y) = \frac{1}{\epsilon} g_\epsilon(y). \tag{B.46}$$

We want to show

$$\mathsf{mmse}(X_U, \mathsf{snr}) - \mathsf{mmse}(X_U, \mathsf{snr}|U) = o\left(\frac{1}{\mathsf{snr}}\right), \tag{B.47}$$

i.e., the benefit of knowing the true distribution is merely $o\left(\frac{1}{\mathsf{snr}}\right)$ in the high-SNR regime. To this end, let $\epsilon = \frac{1}{\sqrt{\mathsf{snr}}}$ and define

$$W = N_{\mathsf{G}} \mathbf{1}_{\{|N_{\mathsf{G}}| \le M\}}, \tag{B.48}$$
$$A_M(\epsilon) = \mathsf{mmse}(W|Y_U) - \mathsf{mmse}(W|Y_U, U) \ge 0. \tag{B.49}$$

By the orthogonality principle,

$$A_M(\epsilon) = \mathbb{E}[(W - \hat{W}(Y_U))^2] - \mathbb{E}[(W - \hat{W}(Y_U, U))^2] \tag{B.50}$$
$$= \mathbb{E}[(\hat{W}(Y_U) - \hat{W}(Y_U, U))^2] \tag{B.51}$$

where

$$\hat{W}(y, u) = \mathbb{E}[W|Y_U = y, U = u] = \frac{f_{u,\epsilon}(y)}{g_{u,\epsilon}(y)} \tag{B.52}$$

and

$$\hat{W}(y) = \mathbb{E}[W|Y_U = y] = \frac{f_\epsilon(y)}{g_\epsilon(y)}. \tag{B.53}$$

Therefore

$$A_M(\epsilon) = \sum_{u=1}^{2} \alpha_u \, \mathbb{E}[(\hat{W}(Y_U, u) - \hat{W}(Y_U))^2 | U = u] \tag{B.54}$$

$$= \sum_{u=1}^{2} \alpha_u \, \mathbb{E}\left[\left(\frac{f_{u,\epsilon}}{g_{u,\epsilon}} - \frac{f_\epsilon}{g_\epsilon}\right)^2 (Y_U) \Big| U = u\right] \tag{B.55}$$

$$= \alpha_1 \alpha_2^2 \, \mathbb{E}\left[\left(\frac{f_{1,\epsilon} g_{2,\epsilon} - f_{2,\epsilon} g_{1,\epsilon}}{g_{1,\epsilon} g_\epsilon}\right)^2 (Y_1)\right] +$$

$$\alpha_2 \alpha_1^2 \, \mathbb{E}\left[\left(\frac{f_{1,\epsilon} g_{2,\epsilon} - f_{2,\epsilon} g_{1,\epsilon}}{g_{2,\epsilon} g_\epsilon}\right)^2 (Y_2)\right] \tag{B.56}$$

$$= \frac{\alpha_1 \alpha_2}{\epsilon} \int \frac{(f_{1,\epsilon} g_{2,\epsilon} - f_{2,\epsilon} g_{1,\epsilon})^2}{g_\epsilon g_{1,\epsilon} g_{2,\epsilon}} \mathrm{d}y \tag{B.57}$$

$$= \alpha_1 \alpha_2 \, \mathbb{E}\left[\frac{g_{1,\epsilon}(Y_U) g_{2,\epsilon}(Y_U)}{g_\epsilon^2(Y_U)} (\hat{W}(Y_U, 1) - \hat{W}(Y_U, 2))^2\right] \tag{B.58}$$

where

- (B.55): by (B.52) and (B.53).
- (B.56): by (B.43) and (B.44).
- (B.57): by (B.45) and (B.46).

Next we show that as $\epsilon \to 0$, the quantity defined in (B.49) vanishes:

$$A_M(\epsilon) = o(1). \tag{B.59}$$

Indeed,

$$A_M(\epsilon) \leq 4M^2 \alpha_1 \alpha_2 \mathbb{E}\left[\frac{g_{1,\epsilon} g_{2,\epsilon}}{g_\epsilon^2} \circ Y_U\right] \tag{B.60}$$

$$\leq 4M^2 \alpha_1 \alpha_2 \left(\mathbb{E}\left[\frac{g_{2,\epsilon}}{g_\epsilon} \circ Y_1\right] + \mathbb{E}\left[\frac{g_{1,\epsilon}}{g_\epsilon} \circ Y_2\right]\right), \tag{B.61}$$

where

- (B.60): by (B.58) and
$$|\hat{W}(y, u)| \leq M \tag{B.62}$$
- (B.61): by (B.43) and (B.44), we have for $u = 1, 2$,
$$\frac{g_{u,\epsilon}}{g_\epsilon} \leq \frac{1}{\alpha_u}. \tag{B.63}$$

Write

$$\mathbb{E}\left[\frac{g_{2,\epsilon}}{g_\epsilon} \circ Y_1\right] = \int \varphi(z) \mathrm{d}z \int \frac{g_{2,\epsilon}}{g_\epsilon}(x + \epsilon z) \mu_1(\mathrm{d}x). \tag{B.64}$$

Fix z. Since

$$\frac{g_{1,\epsilon}}{g_{2,\epsilon}}(x+\epsilon z) = \frac{\mathbb{E}\left[\varphi_z\left(\frac{y-X_1}{\epsilon}\right)\right]}{\mathbb{E}\left[\varphi_z\left(\frac{y-X_2}{\epsilon}\right)\right]}, \tag{B.65}$$

and $\mu_1 \perp \mu_2$, applying Lemma 45 yields

$$\lim_{\epsilon\downarrow 0} \frac{g_{2,\epsilon}}{g_{1,\epsilon}}(x+\epsilon z) = 0, \tag{B.66}$$

for μ_1-a.e. x. Therefore

$$\frac{g_{2,\epsilon}}{g_\epsilon}(x+\epsilon z) = o(1) \tag{B.67}$$

for μ_1-a.e. x. In view of (B.64) and (B.63), we have

$$\mathbb{E}\left[\frac{g_{2,\epsilon}}{g_\epsilon} \circ Y_1\right] = o(1) \tag{B.68}$$

by the dominated convergence theorem, and in entirely analogous fashion:

$$\mathbb{E}\left[\frac{g_{1,\epsilon}}{g_\epsilon} \circ Y_2\right] = o(1). \tag{B.69}$$

Substituting (B.68) and (B.69) into (B.61), we obtain (B.59).

By Lemma 47,

$$\begin{aligned} &\limsup_{\epsilon\downarrow 0}[\mathsf{mmse}(N_{\mathsf{G}}|Y_U) - \mathsf{mmse}(N_{\mathsf{G}}|Y_U, U)] \\ \leq\ &\lim_{\epsilon\downarrow 0} A_M(\epsilon) + 6\left\|N_{\mathsf{G}}\right\|_2 \left\|N_{\mathsf{G}}\mathbf{1}_{\{|N_{\mathsf{G}}|>M\}}\right\|_2 \qquad \text{(B.70)} \\ =\ &6\left\|N_{\mathsf{G}}\mathbf{1}_{\{|N_{\mathsf{G}}|>M\}}\right\|_2. \qquad \text{(B.71)} \end{aligned}$$

By the arbitrariness of M, letting $\epsilon = \frac{1}{\sqrt{\mathsf{snr}}}$ yields

$$\begin{aligned} 0 =\ &\limsup_{\epsilon\downarrow 0}[\mathsf{mmse}(N_{\mathsf{G}}|Y_U) - \mathsf{mmse}(N_{\mathsf{G}}|Y_U, U)] \qquad \text{(B.72)} \\ =\ &\limsup_{\mathsf{snr}\downarrow 0} \mathsf{snr}\left(\mathsf{mmse}(X_U, \mathsf{snr}) - \mathsf{mmse}(X_U, \mathsf{snr}|U)\right), \qquad \text{(B.73)} \end{aligned}$$

which gives the desired result (B.47).

B.3.3 Mixture of two absolutely continuous measures

Now we assume $\mu_1 \ll \mu_2$. In view of the proof in Appendix B.3.2, it is sufficient to show (B.59). Denote the Radon-Nikodym derivative of μ_1 with respect to μ_2 by

$$h = \frac{\mathrm{d}\mu_1}{\mathrm{d}\mu_2}, \tag{B.74}$$

where $h : \mathbb{R} \to \mathbb{R}_+$ satisfies $\int h \mathrm{d}\mu_2 = 1$.

From (B.57), we have

$$A_M(\epsilon) = \alpha_1\alpha_2\mathbb{E}\left[\frac{g_{1,\epsilon}}{g_\epsilon}\left(\frac{f_{1,\epsilon}}{g_{1,\epsilon}} - \frac{f_{2,\epsilon}}{g_{2,\epsilon}}\right)^2 \circ Y_2\right] \tag{B.75}$$

$$= \alpha_1\alpha_2 \int \varphi(z)\mathrm{d}z \int_{F^c} \frac{g_{1,\epsilon}}{g_\epsilon}\left(\frac{f_{1,\epsilon}}{g_{1,\epsilon}} - \frac{f_{2,\epsilon}}{g_{2,\epsilon}}\right)^2 (x+\epsilon z)\mu_2(\mathrm{d}x) \tag{B.76}$$

$$+ \alpha_1\alpha_2 \int \varphi(z)\mathrm{d}z \int_{F} \frac{(f_{1,\epsilon}g_{2,\epsilon} - f_{2,\epsilon}g_{1,\epsilon})^2}{g_\epsilon g_{1,\epsilon} g_{2,\epsilon}^2}(x+\epsilon z)\mu_2(\mathrm{d}x), \tag{B.77}$$

where

$$F = \{x : h(x) > 0\}. \tag{B.78}$$

If $h(x) = 0$, applying Lemma 45 we obtain

$$\frac{g_{1,\epsilon}}{g_\epsilon}(x+\epsilon z) = o\,(1) \tag{B.79}$$

for μ_2-a.e. x and every z. In view of (B.62), we conclude that the integrand in (B.76) is also $o\,(1)$.

If $h(x) > 0$, then

$$|(f_{1,\epsilon}g_{2,\epsilon} - f_{2,\epsilon}g_{1,\epsilon})(x+\epsilon z)| \tag{B.80}$$

$$= \Big| \int \frac{y-x_1}{\epsilon}\varphi_z\left(\frac{y-x_1}{\epsilon}\right)\mathbf{1}_{\{|\frac{y-x_1}{\epsilon}|\le M\}} h(x_1)\mu_2(\mathrm{d}x_1) \int \varphi_z\left(\frac{y-x_2}{\epsilon}\right)\mu_2(\mathrm{d}x_2) \tag{B.81}$$

$$- \int \frac{y-x_1}{\epsilon}\varphi_z\left(\frac{y-x_1}{\epsilon}\right)\mathbf{1}_{\{|\frac{y-x_1}{\epsilon}|\le M\}} \mu_2(\mathrm{d}x_1) \int \varphi_z\left(\frac{y-x_2}{\epsilon}\right) h(x_2)\mu_2(\mathrm{d}x_2)\Big| \tag{B.82}$$

$$\le M \iint |h(x_1) - h(x_2)|\varphi_z\left(\frac{y-x_1}{\epsilon}\right)\varphi_z\left(\frac{y-x_2}{\epsilon}\right)\mu_2(\mathrm{d}x_1)\mu_2(\mathrm{d}x_2) \tag{B.83}$$

$$\le 2M g_{2,\epsilon}(x+\epsilon z) \int |h(x_1) - h(x)|\varphi_z\left(\frac{y-x_1}{\epsilon}\right)\mu_2(\mathrm{d}x_1) \tag{B.84}$$

$$= o\left(g_{2,\epsilon}^2(x+\epsilon z)\right) \tag{B.85}$$

for μ_2-a.e. x and every z, which follows from applying Lemma 46 to $h \in L^1(\mathbb{R}, \mu_2)$. By Lemma 45, we have

$$\lim_{\epsilon\downarrow 0} \frac{g_{2,\epsilon}^2}{g_{1,\epsilon}g_\epsilon}(x) = \frac{1}{h(x)(\alpha_1 + \alpha_2 h(x))} \tag{B.86}$$

Combining (B.85) and (B.86) yields that the integrand in (B.77) is $o\,(1)$. Since the integrand is bounded by $\frac{4M^2}{\alpha_1}$, (B.47) follows from applying the dominated convergence theorem to (B.75).

B.3.4 Finite Mixture

Dealing with the more general case where μ_1 and μ_2 are arbitrary probability measures, we perform the Hahn-Lebesgue decomposition [150, Theorem 1.6.3] on μ_1 with respect to μ_2, which yields

$$\mu_1 = \beta_1 \nu_{\mathrm{c}} + \beta_2 \nu_{\mathrm{s}}, \tag{B.87}$$

where $0 \leq \beta_1 = 1 - \beta_2 \leq 1$ and ν_{c} and ν_{s} are two probability measures such that $\nu_{\mathrm{c}} \ll \mu_2$ and $\nu_{\mathrm{s}} \perp \mu_2$. Consequently, $\nu_{\mathrm{s}} \perp \nu_{\mathrm{c}}$. Therefore

$$\begin{aligned}
\mu &= \alpha_1 \mu_1 + \alpha_2 \mu_2 && \text{(B.88)}\\
&= \beta_1(\alpha_1 \nu_{\mathrm{c}} + \alpha_2 \mu_2) + \beta_2(\alpha_1 \nu_{\mathrm{s}} + \alpha_2 \mu_2) && \text{(B.89)}\\
&= \beta_1 \mu^* + \beta_2 \nu^* && \text{(B.90)}
\end{aligned}$$

where

$$\begin{aligned}
\mu^* &= \alpha_1 \nu_{\mathrm{c}} + \alpha_2 \mu_2, && \text{(B.91)}\\
\nu^* &= \alpha_1 \nu_{\mathrm{s}} + \alpha_2 \mu_2. && \text{(B.92)}
\end{aligned}$$

Then

$$\begin{aligned}
\overline{\mathscr{D}}(\mu) &= \beta_1 \overline{\mathscr{D}}(\mu^*) + \beta_2 \overline{\mathscr{D}}(\nu^*) && \text{(B.93)}\\
&= \beta_1[\alpha_1 \overline{\mathscr{D}}(\nu_{\mathrm{c}}) + \alpha_2 \overline{\mathscr{D}}(\mu_2)] + \beta_2[\alpha_1 \overline{\mathscr{D}}(\nu_{\mathrm{s}}) + \alpha_2 \overline{\mathscr{D}}(\mu_2)] && \text{(B.94)}\\
&= \alpha_1[\beta_1 \overline{\mathscr{D}}(\nu_{\mathrm{c}}) + \beta_2 \overline{\mathscr{D}}(\nu_{\mathrm{s}})] + \alpha_2 \overline{\mathscr{D}}(\mu_2) && \text{(B.95)}\\
&= \alpha_1 \overline{\mathscr{D}}(\beta_1 \nu_{\mathrm{c}} + \beta_2 \nu_{\mathrm{s}}) + \alpha_2 \overline{\mathscr{D}}(\mu_2) && \text{(B.96)}\\
&= \alpha_1 \overline{\mathscr{D}}(\mu_1) + \alpha_2 \overline{\mathscr{D}}(\mu_2), && \text{(B.97)}
\end{aligned}$$

where

- (B.93): applying the results in Appendix B.3.3 to (B.90), since by assumption $\alpha_2 > 0$, we have $\mu^* \ll \nu^*$.
- (B.94), (B.96): applying the results in Appendix B.3.2 to (B.91), (B.92) and (B.87), since $\nu_{\mathrm{c}} \ll \mu_2$, $\nu_{\mathrm{s}} \perp \mu_2$ and $\nu_{\mathrm{s}} \perp \nu_{\mathrm{c}}$.

Similarly, $\underline{\mathscr{D}}(\mu) = \alpha_1 \underline{\mathscr{D}}(\mu_1) + \alpha_2 \underline{\mathscr{D}}(\mu_2)$. This completes the proof of (B.21) and (B.22) for $n = 2$.

Next we proceed by induction on n: Suppose that (B.21) holds for $n = N$. For $n = N + 1$, assume that $\alpha_{N+1} < 1$, then

$$\overline{\mathscr{D}}\left(\sum_{i=1}^{N+1} \alpha_i \mu_i\right) = \overline{\mathscr{D}}\left((1-\alpha_{N+1})\sum_{i=1}^{N} \frac{\alpha_i}{1-\alpha_{N+1}}\mu_i + \alpha_{N+1}\mu_{N+1}\right) \tag{B.98}$$

$$= (1-\alpha_{N+1})\overline{\mathscr{D}}\left(\sum_{i=1}^{N} \frac{\alpha_i}{1-\alpha_{N+1}}\mu_i\right) + \alpha_{N+1}\overline{\mathscr{D}}(\mu_{N+1}) \tag{B.99}$$

$$= \sum_{i=1}^{N+1} \alpha_i \overline{\mathscr{D}}(\mu_i), \tag{B.100}$$

where (B.99) and (B.100) follow from the induction hypothesis. Therefore, (B.21) and (B.22) hold for any $n \in \mathbb{N}$.

B.3.5 Countable Mixture

Now we consider $n = \infty$: without loss of generality, assume that $\sum_{i=1}^{N} \alpha_i < 1$ for all $N \in \mathbb{N}$. Then

$$\overline{\mathscr{D}}(\mu) = \sum_{i=1}^{N} \alpha_i \overline{\mathscr{D}}(\mu_i) + (1 - \sum_{i=1}^{N} \alpha_i)\overline{\mathscr{D}}(\nu_N) \tag{B.101}$$

$$\leq \sum_{i=1}^{N} \alpha_i \overline{\mathscr{D}}(\mu_i) + (1 - \sum_{i=1}^{N} \alpha_i), \tag{B.102}$$

where

- (B.101): we have denoted $\nu_N = \frac{\sum_{i=N+1}^{\infty} \alpha_i \mu_i}{1-\sum_{i=1}^{N} \alpha_i}$.
- (B.102): by Theorem 18.

Sending $N \to \infty$ yields $\overline{\mathscr{D}}(\mu) \leq \sum_{i=1}^{\infty} \alpha_i \overline{\mathscr{D}}(\mu_i)$, and in entirely analogous fashion, $\overline{\mathscr{D}}(\mu) \geq \sum_{i=1}^{\infty} \alpha_i \overline{\mathscr{D}}(\mu_i)$. This completes the proof of Theorem 21.

Remark 33. Theorem 21 also generalizes to non-Gaussian noise. From the above proof, we see that (B.47) holds for all noise densities that satisfy Doob's relative limit theorems, in particular, those meeting the conditions in [151, Theorem 4.1], e.g., uniform (by (B.28)) and exponential and Cauchy density ([151, Theorems 5.1]).

More generally, notice that Lemma 44 deals with convolutional kernels which correspond to additive-noise channels. In [151, Theorem 3.1], Doob also gave a result for general kernels. Therefore, it is possible to extend the results in Theorem 21 to general channels .

B.4 Proof of Theorem 22

Proof. Let $p_i = \mathbb{P}\{X = x_i\}$, $\epsilon = \frac{1}{\sqrt{\mathsf{snr}}}$ and $Y_\epsilon = X + \epsilon N$. In view of (4.45), it is equivalent to show that

$$\mathsf{mmse}(N|Y_\epsilon) = o(1). \tag{B.103}$$

Fix $\delta > 0$. Since $N \in L^2$, there exists $K > 0$, such that

$$\mathbb{E}\left[N^2 \mathbf{1}_{\{|N|>K\}}\right] < \delta/2. \tag{B.104}$$

Since $f_N(N) < \infty$ a.s., we can choose $J > 0$ such that

$$\mathbb{P}\{f_N(N) > J\} < \frac{\delta}{2K^2}. \tag{B.105}$$

Define

$$E_\delta = \{z : f_N(z) \le J, |z| \le K\}. \tag{B.106}$$

and $N_\delta = N\mathbf{1}_{E_\delta}(N)$. Then we have

$$\begin{aligned}\mathbb{E}\left[(N - N_\delta)^2\right] &\le K^2 \mathbb{P}\{f_N(N) > J\} + \mathbb{E}\left[N^2 \mathbf{1}_{\{|N|>K\}}\right] &\text{(B.107)}\\ &\le \delta, &\text{(B.108)}\end{aligned}$$

where (B.108) follows from (B.104) and (B.105).

The optimal estimator for N_δ based on Y_ϵ is given by

$$\hat{N}_\delta(y) = \frac{\sum_j p_j \frac{y-x_j}{\epsilon} \mathbf{1}_{E_\delta}\left(\frac{y-x_j}{\epsilon}\right) f_N\left(\frac{y-x_j}{\epsilon}\right)}{\sum_j p_j f_N\left(\frac{y-x_j}{\epsilon}\right)} \tag{B.109}$$

Then the MMSE of estimating N_δ based on Y_ϵ is

$$\begin{aligned}\mathsf{mmse}(N_\delta|Y_\epsilon) &= \sum_i p_i \mathbb{E}\left[(N_\delta - \hat{N}_\delta(x_i + N))^2\right] &\text{(B.110)}\\ &= \sum_i p_i \int_{E_\delta} (z - \hat{N}_\delta(x_i + z))^2 f_N(z)\mathrm{d}z \\ &\quad + \sum_i p_i \int_{E_\delta^c} \hat{N}_\delta(x_i + z)^2 f_N(z)\mathrm{d}z &\text{(B.111)}\\ &\le \sum_i p_i \int_{E_\delta} g_{z,i}(\epsilon) f_N(z)\mathrm{d}z + K^2 \mathbb{P}\{N \notin E_\delta\} &\text{(B.112)}\\ &\le J \sum_i p_i \int_{-K}^{K} g_{z,i}(\epsilon)\mathrm{d}z + \delta &\text{(B.113)}\end{aligned}$$

where

- (B.112): by $|\hat{N}_\delta(y)| \le K$, since $N_\delta \le K$ a.s. We have also defined

$$g_{z,i}(\epsilon) = (z - \hat{N}_\delta(x_i + z))^2. \tag{B.114}$$

- (B.113): by (B.106) and

$$
\begin{aligned}
& K^2\mathbb{P}\left\{N \notin E_\delta\right\} \\
\leq\ & K^2\mathbb{P}\left\{f_N(N) > J\right\} + K^2\mathbb{P}\left\{|N| > K\right\} && \text{(B.115)}\\
\leq\ & K^2\mathbb{P}\left\{f_N(N) > J\right\} + \mathbb{E}\left[N^2\mathbf{1}_{\{|N|>K\}}\right] && \text{(B.116)}\\
\leq\ & \delta, && \text{(B.117)}
\end{aligned}
$$

where (B.117) follows from (B.104) and (B.105).

Next we show that for all i and $z \in E_\delta$, as $\epsilon \to 0$, we have

$$
g_{z,i}(\epsilon) = o\left(1\right). \tag{B.118}
$$

Indeed, using (B.109),

$$
\begin{aligned}
g_{z,i}(\epsilon) =\ & \left[\frac{\sum_j p_j \frac{x_i - x_j}{\epsilon}\mathbf{1}_{E_\delta}\left(\frac{x_i-x_j}{\epsilon}+z\right) f_N\left(\frac{x_i-x_j}{\epsilon}+z\right)}{\sum_j p_j f_N\left(\frac{x_i-x_j}{\epsilon}+z\right)}\right]^2 && \text{(B.119)}\\
\leq\ & \frac{\left[\sum_{j\neq i} p_j \frac{x_i - x_j}{\epsilon}\mathbf{1}_{E_\delta}\left(\frac{x_i-x_j}{\epsilon}+z\right) f_N\left(\frac{x_i-x_j}{\epsilon}+z\right)\right]^2}{p_i^2 f_N(z)^2} && \text{(B.120)}\\
\leq\ & \frac{J^2(K+|z|)^2}{p_i^2 f_N(z)^2}\left[\sum_{j\neq i} p_j \mathbf{1}_{E_\delta}\left(\frac{x_i - x_j}{\epsilon}+z\right)\right]^2 && \text{(B.121)}\\
\leq\ & \left[\frac{J(K+|z|)}{p_i f_N(z)}\mathbb{P}\left\{X \neq x_i, \left|\frac{x_i - X}{\epsilon}+z\right| \leq K\right\}\right]^2 && \text{(B.122)}\\
=\ & o\left(1\right), && \text{(B.123)}
\end{aligned}
$$

where

- (B.121): by (B.106).
- (B.123): by the boundedness of E_δ and the dominated convergence theorem.

By definition in (B.114), we have

$$
g_{z,i}(\epsilon) \leq (z+K)^2. \tag{B.124}
$$

In view of (B.113) and dominated convergence theorem, we have

$$
\limsup_{\epsilon\to 0} \mathsf{mmse}(N_\delta|Y_\epsilon) \leq \delta \tag{B.125}
$$

Then

$$
\begin{aligned}
\limsup_{\epsilon\to 0}\sqrt{\mathsf{mmse}(N|Y_\epsilon)} \leq\ & \limsup_{\epsilon\to 0}\left\|N - \hat{N}_\delta(Y_\epsilon)\right\|_2 && \text{(B.126)}\\
\leq\ & \limsup_{\epsilon\to 0}\sqrt{\mathsf{mmse}(N_\delta|Y_\epsilon)} + \left\|N - N_\delta\right\|_2 && \text{(B.127)}\\
\leq\ & 2\sqrt{\delta} && \text{(B.128)}
\end{aligned}
$$

where

- (B.126): by the suboptimality of $\hat{N}_\delta$.
- (B.126): by the triangle inequality.
- (B.128): by (B.108) and (B.125).

By the arbitrariness of δ, the proof of (B.103) is completed.

□

B.5 Proof of Theorems 23 – 25

We first compute the optimal MSE estimator under absolutely continuous noise N. Let $\epsilon = \frac{1}{\sqrt{\mathsf{snr}}}$. The density of $Y_\epsilon = X + \epsilon N$ is

$$q_0(y) = \frac{1}{\epsilon}\mathbb{E}\left[f_N\left(\frac{y-X}{\epsilon}\right)\right]. \tag{B.129}$$

Denote

$$q_1(y) = \frac{1}{\epsilon}\mathbb{E}\left[X f_N\left(\frac{y-X}{\epsilon}\right)\right]. \tag{B.130}$$

Then the MMSE estimator of X given Y_ϵ is given by

$$\hat{X}(y) = \mathbb{E}[X|Y_\epsilon = y] = \frac{q_1(y)}{q_0(y)} = \frac{\mathbb{E}\left[X f_N\left(\frac{y-X}{\epsilon}\right)\right]}{\mathbb{E}\left[f_N\left(\frac{y-X}{\epsilon}\right)\right]}. \tag{B.131}$$

B.5.1 Proof of Theorem 24

Proof. By (4.44), we have

$$\mathsf{snr}\cdot\mathsf{mmse}(X,N,\mathsf{snr}) = \mathsf{mmse}(N|Y_\epsilon), \tag{B.132}$$

Due to Theorem 18, we only need to show

$$\underline{\mathscr{D}}(X,N) \geq 1, \tag{B.133}$$

which, in view of (B.132), is equivalent to

$$\liminf_{\epsilon\downarrow 0}\mathsf{mmse}(N|Y_\epsilon) \geq \mathsf{var}N. \tag{B.134}$$

The optimal estimator for N given Y_ϵ is given by

$$\mathbb{E}[N|Y_\epsilon = y] = \frac{\mathbb{E}[N f_X(y-\epsilon N)]}{\mathbb{E}[f_X(y-\epsilon N)]} \tag{B.135}$$

Fix an arbitrary positive δ. Since $N \in L^2(\Omega)$, there exists $M > 0$, such that

$$\mathbb{E}\left[N^2\mathbf{1}_{\{|N|>M\}}\right] < \delta^2. \tag{B.136}$$

Then

$$\begin{aligned}
&\sqrt{\mathsf{mmse}(N|Y_\epsilon)} \\
&= \sqrt{\mathbb{E}[(N-\mathbb{E}[N|Y_\epsilon])^2]} && \text{(B.137)}\\
&\geq \sqrt{\mathbb{E}[(N-\mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}|Y_\epsilon])^2]} - \sqrt{\mathbb{E}[(\mathbb{E}[N\mathbf{1}_{\{|N|>M\}}|Y_\epsilon])^2]} && \text{(B.138)}\\
&\geq \sqrt{\mathbb{E}[(N-\mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}|Y_\epsilon])^2]} - \sqrt{\mathbb{E}[N^2\mathbf{1}_{\{|N|>M\}}]} && \text{(B.139)}\\
&\geq \sqrt{E_{M,\epsilon}} - \delta && \text{(B.140)}
\end{aligned}$$

where

- (B.138): by writing $\mathbb{E}[N|Y_\epsilon] = \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}|Y_\epsilon] + \mathbb{E}[N\mathbf{1}_{\{|N|>M\}}|Y_\epsilon]$ and the triangle inequality.
- (B.139): by $\mathbb{E}[(\mathbb{E}[U|V])^2] \leq \mathbb{E}[U^2]$ for all $U, V \in L^2(\Omega)$.
- (B.140): by (B.136) and

$$E_{M,\epsilon} \triangleq \mathbb{E}[(N-\mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}|Y_\epsilon])^2]. \tag{B.141}$$

Define

$$\begin{aligned}
p_M(y;\epsilon) &= \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}} f_X(y-\epsilon N)] && \text{(B.142)}\\
q(y;\epsilon) &= \mathbb{E}[f_X(y-\epsilon N)] && \text{(B.143)}\\
\hat{N}_M(y;\epsilon) &= \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}|Y_\epsilon = y] = \frac{p_M(y;\epsilon)}{q(y;\epsilon)}. && \text{(B.144)}
\end{aligned}$$

Suppose $f \in C^b$. Then by the bounded convergence theorem,

$$\begin{aligned}
\lim_{\epsilon\downarrow 0} p_M(x+\epsilon z;\epsilon) &= \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}] f_X(x), && \text{(B.145)}\\
\lim_{\epsilon\downarrow 0} q(x+\epsilon z;\epsilon) &= f_X(x) && \text{(B.146)}
\end{aligned}$$

hold for all $x, z \in \mathbb{R}$. Since $f_X(X) > 0$ a.s.,

$$\begin{aligned}
\lim_{\epsilon\downarrow 0} \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}|Y_\epsilon] &= \lim_{\epsilon\downarrow 0} \hat{N}_M(X+\epsilon N;\epsilon) && \text{(B.147)}\\
&= \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}] && \text{(B.148)}
\end{aligned}$$

holds a.s. Then by Fatou's lemma:

$$\liminf_{\epsilon\downarrow 0} E_{M,\epsilon} \geq \mathbb{E}\left[N - \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}]\right]^2 \geq \mathsf{var}N. \tag{B.149}$$

By (B.140),

$$\liminf_{\epsilon\downarrow 0} \sqrt{\mathsf{mmse}(N|Y_\epsilon)} \geq \liminf_{\epsilon\downarrow 0} \sqrt{E_{M,\epsilon}} - \delta \tag{B.150}$$

$$\geq \sqrt{\mathsf{var}N} - \delta. \tag{B.151}$$

By the arbitrariness of δ, we conclude that

$$\liminf_{\epsilon\downarrow 0} \mathsf{mmse}(N|Y_\epsilon) \geq \mathsf{var}N, \tag{B.152}$$

hence (B.133) holds. □

B.5.2 Proof of Theorem 23

Now we are dealing with X whose density is not necessarily continuous or bounded. In order to show that (B.145) and (B.146) continue to hold under the assumptions of the noise density in Theorem 23, we need the following lemma from [152, Section 3.2]:

Lemma 48 ([152, Theorem 3.2.1])**.** *Suppose the family of functions* $\{K_\epsilon : \mathbb{R}^d \to \mathbb{R}\}_{\epsilon>0}$ *satisfies the following conditions: for some constant* $\eta > 0$ *and* $C \in \mathbb{R}$*,*

$$\int_{\mathbb{R}^d} K_\epsilon(x)\mathrm{d}x = C \tag{B.153}$$

$$\sup_{x\in\mathbb{R}^d,\epsilon>0} \epsilon|K_\epsilon(x)| < \infty \tag{B.154}$$

$$\sup_{x\in\mathbb{R}^d,\epsilon>0} \frac{|x|^{1+\eta}}{\epsilon^\eta}|K_\epsilon(x)| < \infty \tag{B.155}$$

hold for all $\epsilon > 0$ *and* $x \in \mathbb{R}^d$*. Then for all* $f \in L^1_{\mathrm{loc}}(\mathbb{R}^d)$*,*

$$\lim_{\epsilon\downarrow 0} f * K_\epsilon(x) = Cf(x) \tag{B.156}$$

holds for Lebesgue-a.e. x*.*

Note that in the original version of Lemma 48 in [152, Section 3.2] $C = 1$, and K_ϵ is dubbed *approximation of the identity*. For $C \neq 0$ or 1, the same conclusion follows from scaling. The case of $C = 0$ can be shown as follows: take some kernel G_ϵ which is an approximation to the identity. Then $G + K$ is also an approximation to the identity. Then the conclusion for K follows by applying Lemma 48 to both G and $G + K$ and then subtracting the corresponding (B.156) limits.

Proof of Theorem 23. Based on the proof of Theorem 24, it is sufficient to show that (B.145) and (B.146) hold for Lebesgue-a.e. x and z. Fix z. First look at (B.146): introduce the following kernel which corresponds to the density of $\epsilon(N - z)$

$$K_\epsilon(x) = \frac{1}{\epsilon} f_N\left(\frac{x}{\epsilon} + z\right). \tag{B.157}$$

We check that K_ϵ is an approximation to the identity by verifying:

- (B.153): $\int_{\mathbb{R}} K_\epsilon(x)\mathrm{d}x = \int_{\mathbb{R}} f_N(u)\mathrm{d}u = 1$.
- (B.154): $\sup_{x\in\mathbb{R},\epsilon>0} \epsilon|K_\epsilon(x)| = \sup_{u\in\mathbb{R}} f_N(u) < \infty$, since f_N is bounded.
- (B.155): by boundedness of f_N and (4.69), we have: for some $\eta > 0$,

$$\sup_{u\in\mathbb{R}} |u|^{1+\eta} f_N(u) < \infty \tag{B.158}$$

then

$$\left(\sup_{x\in\mathbb{R},\epsilon>0} \frac{|x|^{1+\eta}}{\epsilon^\eta}|K_\epsilon(x)|\right)^{\frac{1}{1+\eta}} = \sup_{u\in\mathbb{R}} |u| f_N^{\frac{1}{1+\eta}}(u+z) \tag{B.159}$$

$$\leq \sup_{u\in\mathbb{R}} |u| f_N^{\frac{1}{1+\eta}}(u) + |z| \sup_{u\in\mathbb{R}} f_N^{\frac{1}{1+\eta}}(u) \tag{B.160}$$

$$< \infty. \tag{B.161}$$

Note that

$$q(x+\epsilon z;\epsilon) = f_X * K_\epsilon(x). \tag{B.162}$$

Since $f_X \in L^1(\mathbb{R})$, by Lemma 48, (B.146) holds for Lebesgue-a.e. x. Similarly, we define

$$G_\epsilon(x) = \left(\frac{x}{\epsilon} + z\right) \mathbf{1}_{\left\{\left|\frac{x}{\epsilon}+z\right|\leq M\right\}} \frac{1}{\epsilon} f_N\left(\frac{x}{\epsilon} + z\right). \tag{B.163}$$

Then it is verifiable that G_ϵ satisfies (B.153) – (B.155), with

$$\int G_\epsilon(x)\mathrm{d}x = \mathbb{E}[N\mathbf{1}_{\{|N|\leq M\}}]. \tag{B.164}$$

Since

$$p_M(x+\epsilon z;\epsilon) = f_X * G_\epsilon(x), \tag{B.165}$$

Lemma 48 implies that (B.145) holds for Lebesgue-a.e. x. □

B.5.3 Proof of Theorem 25

Proof. Without loss of generality, we shall assume that $\mathbb{E}N = 0$. Following the notation in the proof of Theorem 24, similar to $q(y;\epsilon)$ in (B.143), we define

$$p(y;\epsilon) = \mathbb{E}[N f_X(y - \epsilon N)]. \tag{B.166}$$

Then

$$\mathbb{E}[N|Y_\epsilon = y] = \frac{p(y;\epsilon)}{q(y;\epsilon)}. \tag{B.167}$$

Fix $x, z \in \mathbb{R}$. Since f_X, f_X' and f_X'' are all bounded, we can interchange derivatives with integrals [101, Theorem 2.27] and write

$$
\begin{aligned}
& q(x+\epsilon z;\epsilon) \\
&= \mathbb{E}[f_X(x+\epsilon(z-N))] && \text{(B.168)}\\
&= f_X(x) + \epsilon f_X'(x)\mathbb{E}(z-N) + \frac{\epsilon^2}{2} f_X''(x)\mathbb{E}(z-N)^2 + o\left(\epsilon^2\right) && \text{(B.169)}\\
&= f_X(x) + \epsilon f_X'(x) z + \frac{\epsilon^2}{2} f_X''(x)(z^2 + \mathsf{var}N) + o(\epsilon^2). && \text{(B.170)}
\end{aligned}
$$

Similarly, since $\mathbb{E}[|N|^3] < \infty$, we have

$$
\begin{aligned}
& p(x+\epsilon z;\epsilon) \\
&= \mathbb{E}[N f_X(x+\epsilon(z-N))] && \text{(B.171)}\\
&= -\epsilon f_X'(x)\,\mathsf{var}N + \frac{\epsilon^2}{2} f_X''(x)(\mathbb{E}[N^3] - 2z\,\mathsf{var}N) + o\left(\epsilon^2\right). && \text{(B.172)}
\end{aligned}
$$

Define the score function $\rho(x) = \frac{f_X'}{f_X}(x)$. Since $f_X(X) > 0$ a.s., by (B.167),

$$
\begin{aligned}
\mathbb{E}[N|Y_\epsilon] = \frac{\epsilon^2}{2}&\left[\frac{f_X''(X)}{f_X(X)}(\mathbb{E}[N^3] - 2N\mathsf{var}N) + 2N\rho^2(X)\mathsf{var}N\right] \\
&- \epsilon\rho(X)\mathsf{var}N + o\left(\epsilon^2\right). && \text{(B.173)}
\end{aligned}
$$

holds a.s. Then

$$
\begin{aligned}
\mathsf{mmse}(N|Y_\epsilon) &= \mathbb{E}[N - \mathbb{E}[N|Y_\epsilon]]^2 && \text{(B.174)}\\
&= \mathsf{var}N - \epsilon\mathbb{E}\left[N\right]\mathbb{E}[\rho(X)]\mathsf{var}N + \epsilon^2(\mathsf{var}N)^2\mathbb{E}[\rho^2(X)] - 2\epsilon^2(\mathsf{var}N)^2 J(X) \\
&\quad + \epsilon^2(\mathbb{E}[N]\mathbb{E}[N^3] - 2\mathsf{var}^2 N)\mathbb{E}\left[\frac{f_X''(X)}{f_X(X)}\right] + o\left(\epsilon^2\right) && \text{(B.175)}\\
&= \mathsf{var}N - \epsilon^2(\mathsf{var}N)^2 J(X) + o\left(\epsilon^2\right), && \text{(B.176)}
\end{aligned}
$$

where

- (B.175): by the bounded convergence theorem, because the ratio between the $o\left(\epsilon^2\right)$ term in (B.173) and ϵ^2 is upper bounded by $\rho^2(X)$ and $\frac{f_X''(X)}{f_X(X)}$, which are integrable by assumption.
- (B.176): by $\mathbb{E}\left[\rho(X)\right] = 0$, $\mathbb{E}\left[\rho^2(X)\right] = J(X)$ and $\mathbb{E}\left[\frac{f_X''(X)}{f_X(X)}\right] = 0$, in view of (4.84).

□

B.6 Proof of Theorem 28

By Remark 7, Theorem 28 holds trivially if $\mathbb{P}\{X=0\}=1$ or $\mathbb{P}\{X=1\}=1$. Otherwise, since $X=0$ if and only if $(X)_j=0$ for all $j\in\mathbb{N}$, but $\{(X)_j\}$ are i.i.d., therefore $\mathbb{P}\{X=0\}=0$. Similarly, $\mathbb{P}\{X=1\}=0$. Hence $0<X<1$ a.s.

To prove (4.105), we define the function $G:\mathbb{R}\to\mathbb{R}_+$

$$G(b)=M^{2b}\cdot\mathsf{mmse}(X,N,M^{2b}) \tag{B.177}$$
$$=\mathsf{mmse}(N,X,M^{-2b}), \tag{B.178}$$

where (B.178) follows from (4.44). The oscillatory behavior of G is given in the following lemma:

Lemma 49.

1. *For any $b\in\mathbb{R}$, $\{G(k+b):k\in\mathbb{N}\}$ is an nondecreasing nonnegative sequence bounded from above by $\mathsf{var}N$.*

2. *Define a function $\Psi:\mathbb{R}\to\mathbb{R}_+$ by*

$$\Psi(b)=\lim_{k\to\infty}G(k+b). \tag{B.179}$$

Then Ψ is a 1-periodic function, and the convergence in (B.179) is uniform *in $b\in[0,1]$. Therefore as $b\to\infty$,*

$$G(b)=\Psi(b)+o(1) \tag{B.180}$$

In view of Lemma 49, we define a $2\log M$-periodic function $\Phi_{X,N}:\mathbb{R}\to[0,1]$ as follows:

$$\Phi_{X,N}(t)=\frac{1}{\mathsf{var}N}\Psi\left(\frac{t}{2\log M}\right). \tag{B.181}$$

Having defined $\Phi_{X,N}$, (4.105), (4.106) and (4.107) readily follow from (B.180).

Next we prove (4.108) in the case of Gaussian N: in view of (B.181), it is equivalent to show

$$\int_0^1\Psi(b)\mathrm{d}b=d(X). \tag{B.182}$$

Denote

$$\mathsf{snr}=M^{2b}. \tag{B.183}$$

Recalling $G(b)$ defined in (B.178), we have

$$\int_0^b G(\tau)\mathrm{d}\tau=\int_1^{\mathsf{snr}}\gamma\,\mathsf{mmse}(X,\gamma)\frac{1}{2\gamma\log M}\mathrm{d}\gamma \tag{B.184}$$
$$=\frac{I(\mathsf{snr})-I(1)}{2\log M}, \tag{B.185}$$

where (B.185) follows from (4.49) and $I(\mathsf{snr})$ is defined in (4.50).

Since $d(X)$ exists, in view of (2.66) and (2.67), we have

$$\lim_{b\to\infty}\frac{1}{b}\int_0^b G(\tau)\mathrm{d}\tau = \lim_{\mathsf{snr}\to\infty}\frac{2I(\mathsf{snr})}{\log\mathsf{snr}} = d(X). \tag{B.186}$$

Since the convergence in (B.179) is uniform in $b\in[0,1]$, for all $\epsilon>0$, there exists k_0 such that for all $k\geq k_0$ and $b\in[0,1]$,

$$|G(k+b)-\Phi_{X,N_{\mathsf{G}}}(b)|\leq\epsilon. \tag{B.187}$$

Then for all integers $k\geq k_0$, we have

$$\left|\frac{1}{k}\int_0^k G(\tau)\mathrm{d}\tau-\int_0^1\Psi(\tau)\mathrm{d}\tau\right| \leq \frac{2k_0}{k}+\frac{1}{k}\sum_{j=k_0}^{k-1}\int_0^1|G(j+\tau)-\Psi(\tau)|\mathrm{d}\tau \tag{B.188}$$

$$\leq \frac{2k_0}{k}+\frac{k-k_0}{k}\epsilon \tag{B.189}$$

where

- (B.188): G and Ψ map into $[0,1]$.
- (B.189): by (B.187).

By the arbitrariness of ϵ and (B.186), we have

$$\int_0^1\Psi(\tau)\mathrm{d}\tau=\lim_{k\to\infty}\frac{1}{k}\int_0^k G(\tau)\mathrm{d}\tau=d(X). \tag{B.190}$$

To finish the proof, we prove Lemma 49. Note that

$$G(k+1+b)=\mathsf{mmse}(N,X,M^{-2(k+1+b)}) \tag{B.191}$$

$$=\mathsf{mmse}(N|N+\sqrt{\mathsf{snr}}M^{k+1}X) \tag{B.192}$$

$$=\mathsf{mmse}(N|N+\sqrt{\mathsf{snr}}M^{k}(U+V)) \tag{B.193}$$

$$\geq\mathsf{mmse}(N|N+\sqrt{\mathsf{snr}}M^{k}V) \tag{B.194}$$

$$=\mathsf{mmse}(N|N+\sqrt{\mathsf{snr}}M^{k}X) \tag{B.195}$$

$$=G(k+b) \tag{B.196}$$

where

- (B.191): by (B.178);
- (B.192): by (4.43) and (B.183);
- (B.193): by the M-ary expansion of X in (2.47) and we have defined

$$U=(X)_1, \tag{B.197}$$

$$V=\sum_{j=1}^{\infty}(X)_{j+1}M^{-j}; \tag{B.198}$$

- (B.194): by the data-processing inequality of MMSE [87], since U is independent of V;
- (B.195): by $V \overset{\mathrm{D}}{=} X$ since $\{(X)_j\}$ is an i.i.d. sequence.

Therefore for fixed b, $G(k+b)$ is an nondecreasing sequence in k. By (B.191),

$$G(k+b) \le \mathsf{var} N, \tag{B.199}$$

hence $\lim_{k\to\infty} G(k+b)$ exists, denoted by $\Psi(b)$. The 1-periodicity of Ψ readily follows.

To prove (B.180), we show that there exist two functions $c_1, c_2 : \mathbb{R} \to \mathbb{R}_+$ depending on the distribution of X and N only, such that

$$\lim_{b\to\infty} c_i(b) = 0, \quad i = 1, 2 \tag{B.200}$$

and

$$0 \le \Psi(b) - G(b) \tag{B.201}$$

$$\le c_1^2(b) + c_2(b) + 2c_1(b)\sqrt{\mathsf{var} N + c_2(b)}. \tag{B.202}$$

Then (B.180) follows by combining (B.200) – (B.202) and sending $b \to \infty$. Inequalities (B.201) and (B.202) also show that the convergence in (B.179) is uniform in $b \in [0, 1]$.

To conclude this proof, we proceed to construct the desired functions c_1 and c_2 and prove (B.201) and (B.202): by monotonicity, for all $b \in \mathbb{R}$,

$$G(b) \le \Psi(b). \tag{B.203}$$

Hence (B.201) follows. To prove (B.202), fix $k \in \mathbb{N}$ and $K > 0$ to be specified later. Define

$$N_K = N\mathbf{1}_{\{|N|\le K\}} \tag{B.204}$$

$$\bar{N}_K = N - N_K \tag{B.205}$$

We use a suboptimal estimator to bound $G(k+b) - G(b)$. To streamline the proof, introduce the following notation:

$$W = M^k[X]_k \tag{B.206}$$

$$Z = M^k(X - [X]_k) \tag{B.207}$$

$$Y = \sqrt{\mathsf{snr}}X + N \tag{B.208}$$

$$Y' = \sqrt{\mathsf{snr}}M^k X + N \tag{B.209}$$

$$= \sqrt{\mathsf{snr}}W + \sqrt{\mathsf{snr}}Z + N, \tag{B.210}$$

where W is integer-valued. Note that since X has i.i.d. M-ary expansion, $\{N, W, Z\}$ are independent, and

$$Z \overset{\mathrm{D}}{=} X, \tag{B.211}$$

hence

$$Y' \overset{\mathrm{D}}{=} Y + \sqrt{\mathsf{snr}}W. \tag{B.212}$$

Based on Y', we use the following two-stage suboptimal estimator $\tilde{N}_K$ for N_K: first estimate W (the first k bits of X) based on Y' according to

$$\tilde{w}(y) = \left[\frac{y}{\sqrt{\mathsf{snr}}}\right]. \tag{B.213}$$

Then peel off $\tilde{W} = \tilde{w}(Y')$ from Y' and plug it into the optimal estimator for X based on Y

$$\hat{N}_K(y) = \mathbb{E}[N_K|Y = y] \tag{B.214}$$

to estimate X, i.e.,

$$\tilde{N}_K(y) = \hat{N}_K(y - \sqrt{\mathsf{snr}}\tilde{w}(y)). \tag{B.215}$$

Next we bound the probability of choosing the wrong W:

$$\begin{aligned}
\mathbb{P}\left\{\tilde{W} \neq W\right\} &= 1 - \mathbb{P}\left\{\left[\frac{N}{\sqrt{\mathsf{snr}}} + W + Z\right] = W\right\} && \text{(B.216)}\\
&= 1 - \mathbb{P}\left\{W \leq \frac{N}{\sqrt{\mathsf{snr}}} + W + Z < W + 1\right\} && \text{(B.217)}\\
&= 1 - \mathbb{P}\left\{-\sqrt{\mathsf{snr}}X \leq N < \sqrt{\mathsf{snr}}(1 - X)\right\} && \text{(B.218)}\\
&= \int_{(0,1)} P_X(\mathrm{d}x)[1 - \mathbb{P}\left\{-\sqrt{\mathsf{snr}}x \leq N < \sqrt{\mathsf{snr}}(1 - x)\right\}] && \text{(B.219)}\\
&\triangleq \kappa(\mathsf{snr}) \to 0, && \text{(B.220)}
\end{aligned}$$

where

- (B.216): by (B.210) and (B.213);
- (B.217): by the fact that $[x] = n$ if and only if $n \leq x < n + 1$;
- (B.218): by (B.211);
- (B.219): by the assumption that $X \in (0, 1)$ a.s.;
- (B.220): by the bounded convergence theorem.

Note that $\kappa(\mathsf{snr})$ is a nonincreasing nonnegative function depending only on the distribution of X and N.

Finally we choose K and analyze the performance of $\tilde{N}_K$. Observe from (B.216) that the probability of choosing the wrong W does *not* depend k. This allows us to choose K independent of k:

$$K = [\kappa(\mathsf{snr})]^{-\frac{1}{3}}. \tag{B.221}$$

Therefore,

$$\begin{aligned}
\sqrt{G(k + b)} &= \sqrt{\mathsf{mmse}(N|Y')} && \text{(B.222)}\\
&\leq \left\|N - N_K + N_K - \tilde{N}_K(Y')\right\|_2 && \text{(B.223)}\\
&\leq \left\|\bar{N}_K\right\|_2 + \left\|N_K - \hat{N}_K(Y' - \sqrt{\mathsf{snr}}\tilde{W})\right\|_2 && \text{(B.224)}
\end{aligned}$$

where

- (B.222): by (B.178) and (B.210);
- (B.223): by the suboptimality of $\tilde{N}_K$;
- (B.224): by (B.215).

Now

$$\mathbb{E}\left[(N_K - \hat{N}_K(Y' - \sqrt{\mathsf{snr}}\tilde{W}))^2\right] \le \mathbb{E}[N_K - \hat{N}_K(Y)]^2 + \mathbb{E}\left[(N_K - \hat{N}_K(Y' - \sqrt{\mathsf{snr}}\tilde{W}))^2 \mathbf{1}_{\{\tilde{W} \neq W\}}\right] \quad \text{(B.225)}$$

$$\le \mathsf{mmse}(N_K|Y) + 4K^2 \mathbb{P}\left\{\tilde{W} \neq W\right\} \quad \text{(B.226)}$$

$$\le G(b) + 3\left\|\bar{N}_K\right\|_2 \|N\|_2 + 4\kappa(\mathsf{snr})^{\frac{1}{3}} \quad \text{(B.227)}$$

where

- (B.225): by (B.212);
- (B.226): by $|\hat{N}_K(y)| = |\mathbb{E}[\hat{N}_K|Y = y]| \le K$ for all y, since $|N_K| \le K$ a.s.;
- (B.227): by Lemma 47, (B.220), (B.221) and $\mathsf{mmse}(N|Y) = G(b)$.

Define

$$c_1(b) = \left\|\bar{N}_K\right\|_2 \quad \text{(B.228)}$$

$$c_2(b) = 3\left\|\bar{N}_K\right\|_2 \|N\|_2 + 4\kappa(\mathsf{snr})^{\frac{1}{3}} \quad \text{(B.229)}$$

where b and K are related to snr through (B.183) and (B.221), respectively. Then substituting (B.227) into (B.224) yields

$$\sqrt{G(k+b)} \le c_1(b) + \sqrt{G(b) + c_2(b)}. \quad \text{(B.230)}$$

Note that the right hand side of (B.230) does not depend on k. By (B.179), sending $k \to \infty$ we obtain

$$\sqrt{\Psi(b)} \le c_1(b) + \sqrt{G(b) + c_2(b)}. \quad \text{(B.231)}$$

In view of (B.203), squaring[1] (B.231) on both sides and noticing that $G(b) \le \mathsf{var}N$ for all b, we have

$$\Psi(b) - G(b) \le c_1^2(b) + c_2(b) + 2c_1(b)\sqrt{\mathsf{var}N + c_2(b)}. \quad \text{(B.232)}$$

By (B.220) and (B.221), $K \to \infty$ as $b \to \infty$. Since $\mathbb{E}[N^2] < \infty$, $c_1(b)$ and $c_2(b)$ both tend to zero as $b \to \infty$, and (B.202) follows.

[1] In this step it is essential that $G(b)$ be bounded, because in general $\sqrt{a_n} = \sqrt{b_n} + o(1)$ does *not* imply that $a_n = b_n + o(1)$. For instance, $a_n = n + 1, b_n = n$.

B.7 Proof for Remark 8

We show that for any singular X and any binary-valued N,

$$\underline{\mathscr{D}}(X, N) = 0. \tag{B.233}$$

To this end, we need the following auxiliary results:

Lemma 50 (Mutually singular hypothesis)**.** *The optimal test for the binary hypothesis testing problem*

$$\begin{cases} H_0 : & P \\ H_1 : & Q \end{cases} \tag{B.234}$$

with prior $\mathbb{P}\{H_0\} \in (0, 1)$ *has zero error probability if and only if* $P \perp Q$*.*

Proof. By definition, $P \perp Q$ if and only if there exists an event A such that $P(A) = 1$ and $Q(A) = 0$. Then the test that decides H_0 if and only if A occurs yields zero error probability. □

Lemma 51 ([153, Theorem 10])**.** *Let* μ *be a probability measure on* $(\mathbb{R}, \mathcal{B})$ *that is mutually singular with respect to Lebesgue measure. Then there exists a Borel set* E *with* $\mu(E) = 1$ *and a non-empty perfect set* C *such that* $\{c + E : c \in C\}$ *is a family of disjoint sets.*

Proof of (B.233). In view of Theorem 19, we may assume that N is $\{0, 1\}$-valued without loss of generality. For any input X whose distribution μ is mutually singular to the Lebesgue measure, we show that there exists a vanishing sequence $\{\epsilon_n\}$, such that for all n,

$$\mathsf{mmse}(X|X + \epsilon_n N) = 0. \tag{B.235}$$

which implies that $\underline{\mathscr{D}}(X, N) = 0$.

By Lemma 51, there exists a Borel set E and a perfect set C, such that $\mu(E) = 1$ and $\{c + E : c \in C\}$ is a family of disjoint sets. Pick any $a \in C$. Since C is perfect, there exists $\{a_n\} \subset C$, such that $a_n \to a$ and $\epsilon_n \triangleq a_n - a > 0$. Since X and $X + \epsilon_n$ are supported on disjoint subsets E and $E + \epsilon_n$ respectively, their distributions are mutually singular. By Lemma 50, the optimal test for N based on $X + \epsilon_n N$ succeeds with probability one, which implies (B.235). □

Appendix C

Proofs in Chapter 5

C.1 Injectivity of the cosine matrix

We show that the cosine matrix defined in Remark 17 is injective on Σ_k. We consider a more general case. Let $l = 2k \leq n$ and $\{\omega_1, \ldots, \omega_n\} \subset \mathbb{R}$. Let $\mathbf{H}$ be an $l \times n$ matrix where $H_{ij} = \cos((i-1)\omega_j)$. We show that each $l \times l$ submatrix formed by columns of $\mathbf{H}$ is non-singular if and only if $\{\cos(\omega_1), \ldots, \cos(\omega_n)\}$ are distinct.

Let $\mathbf{G}$ be the submatrix consisting of the first l columns of $\mathbf{H}$. Then $G_{ij} = T_{i-1}(x_j)$ where $x_j = \cos(\omega_j)$ and T_m denotes the mth order Chebyshev polynomial of the first kind [154]. Note that $\det(\mathbf{G})$ is a polynomial in $(x_1, \ldots, x_l)$ of degree $l(l-1)/2$. Also $\det(\mathbf{G}) = 0$ if $x_i = x_j$ for some $i \neq j$. Therefore $\det(\mathbf{G}) = C \prod_{1 \leq i < j \leq l}(x_i - x_j)$. The constant C is given by the coefficient of the highest order term in the contribution from the main diagonal $\prod_{i=1}^{l} T_{i-1}(x_i)$. Since the leading coefficient of T_j is 2^{j-1}, we have $C = 2^{(l-1)(l-2)/2}$. Therefore,

$$\det(\mathbf{G}) = 2^{(l-1)(l-2)/2} \prod_{1 \leq i < j \leq l} [\cos(\omega_i) - \cos(\omega_j)] \neq 0. \tag{C.1}$$

C.2 Proof of Lemma 9

Proof. Since $\|\cdot\|_p$-norms are equivalent, it is sufficient to only consider $p = \infty$. Observe that $\epsilon \mapsto N_A(\epsilon)$ is non-increasing. Hence for any $2^{-m} \leq \epsilon < 2^{-(m-1)}$, we have

$$\frac{\log N_A(2^{-(m-1)})}{m} \leq \frac{\log N_A(\epsilon)}{\log \frac{1}{\epsilon}} \leq \frac{\log N_A(2^{-m})}{m-1}. \tag{C.2}$$

Therefore it is sufficient to restrict to $\epsilon = 2^{-m}$ and $m \to \infty$ in (3.2) and (3.3). To see the equivalence of covering by mesh cubes, first note that $\tilde{N}_A(2^{-m}) \geq N_A(2^{-m})$. On the other hand, any ℓ_∞-ball of radius 2^{-m} is contained in the union of 3^n mesh cubes of size 2^{-m} (by choosing a cube containing some point in the set together with its neighboring cubes). Thus $\tilde{N}_A(2^{-m}) \leq 3^n N_A(2^{-m})$. Hence the limits in (3.2) and (3.3) coincide with those in (3.12) and (3.13). $\square$

C.3 Proof of Lemma 11

Proof. By Pinsker's inequality,

$$D(P||Q) \geq \frac{1}{2} d^2(P,Q) \log \mathrm{e}, \tag{C.3}$$

where $d(P,Q)$ is the total variation distance between P and Q and $0 \leq d(P,Q) \leq 2$. In this case where $\mathcal{X}$ is countable,

$$d(P,Q) = \sum_{x \in \mathcal{X}} |P(x) - Q(x)|. \tag{C.4}$$

By [55, Lemma 2.7, p. 33], when $d(P,Q) \leq \frac{1}{2}$,

$$|H(P) - H(Q)| \leq d(P,Q) \log \frac{|\mathcal{X}|}{d(P,Q)} \tag{C.5}$$

$$\leq d(P,Q) \log |\mathcal{X}| + \mathrm{e}^{-1} \log \mathrm{e}. \tag{C.6}$$

When $d(P,Q) \geq \frac{1}{2}$, by (C.3), $D(P||Q) \geq \frac{1}{8} \log \mathrm{e}$; when $0 \leq d(P,Q) < 1/2$, by (C.6),

$$d(P,Q) \geq \frac{(|H(P) - H(Q)| - \mathrm{e}^{-1} \log \mathrm{e})^+}{\log |\mathcal{X}|}. \tag{C.7}$$

Using (C.3) again,

$$D(P||Q) \geq \frac{1}{2} \left[\frac{(|H(P) - H(Q)| - \mathrm{e}^{-1} \log \mathrm{e})^+}{\log |\mathcal{X}|} \right]^2 \log \mathrm{e}. \tag{C.8}$$

Since $|H(P) - H(Q)| \geq \delta$ holds in the minimization of (3.26) and (3.28), (3.30) is proved. □

C.4 Proof of Lemma 19

Lemma 52. *Let $A \subset \mathbb{R}^k$ be compact, and let $g : A \to \mathbb{R}^n$ be continuous. Then, there exists a Borel measurable function $f : g(A) \to A$ such that $g(f(x)) = x$ for all $x \in g(A)$.*

Proof. For all $x \in g(A)$, $g^{-1}(\{x\})$ is nonempty and compact since g is continuous. For each $i \in \{1, \ldots, k\}$, let the ith component of f be

$$f_i(x) = \min\{t_i : t^k \in g^{-1}(\{x\})\}, \tag{C.9}$$

where t_i is the ith coordinate of t^k. This defines $f : g(A) \to A$ which satisfies $g(f(x)) = x$ for all $x \in g(A)$. Now we claim that each f_i is *lower semicontinuous*, which implies that f is Borel measurable. To this end, we show that for any $a \in A$, $f^{-1}((a, \infty))$ is open. Assume the opposite, then there exists a sequence $\{y_m\}$ in

$g(A)$ that converges to $y \in g(A)$, such that $f(y_m) \leq a$ and $f(y) > a$. Due to the compactness of A, there exists a subsequence $f(y_{m_l})$ that converges to some point x in A. Therefore $x \leq a$. But by the continuity of g, we have

$$g(x) = \lim_{l\to\infty} g(f(y_{m_l})) = \lim_{l\to\infty} y_{m_l} = y.$$

Hence by definition of f and $f(y) > a$, we have $x > a$, which is a contradiction. Therefore f_i is lower semicontinuous. □

Proof of Lemma 19. Let $S^n \subset \mathbb{R}^n$ be $\lfloor Rn \rfloor$-rectifiable and assume that (5.99) holds for all $n \geq N$. Then by definition there exists a bounded subset $T^n \subset \mathbb{R}^{\lfloor Rn \rfloor}$ and a Lipschitz function $g_n : T^n \to \mathbb{R}^n$, such that $S^n = g_n(T^n)$. By continuity, g_n can be extended to the closure $\overline{T^n}$, and $\overline{S^n} = g_n(\overline{T^n})$. Since $\overline{T^n}$ is compact, by Lemma 52, there exists a Borel function $f_n : \mathbb{R}^n \to \mathbb{R}^{\lfloor Rn \rfloor}$, such that $g_n(f_n(x^n)) = x^n$ for all $x^n \in \overline{S^n}$. By Kirszbraun's theorem [104, 2.10.43], g_n can be extended to a Lipschitz function on the whole $\mathbb{R}^{\lfloor Rn \rfloor}$. Then

$$\mathbb{P}\left\{g_n(f_n(X^n)) = X^n\right\} \geq \mathbb{P}\left\{X^n \in \overline{S}^n\right\} \geq 1 - \epsilon \tag{C.10}$$

for all $n \geq N$, which proves the ϵ-achievability of R. □

C.5 Proof of Lemma 21

Proof. Suppose we can construct a mapping $\tau : [0,1]^k \to [0,1]^{k+1}$ such that

$$\left\|\tau(x^k) - \tau(y^k)\right\|_\infty \geq M^{-k}\hat{d}(x^k, y^k) \tag{C.11}$$

holds for all $x^k, y^k \in [0,1]^k$. By (C.11), τ is injective. Let $W' = \tau(W)$ and $g' = g \circ \tau^{-1}$. Then by (C.11) and the L-Lipschitz continuity of g,

$$\begin{aligned}
\left\|g'(x^k) - g'(y^k)\right\|_\infty &= \left\|g(\tau^{-1}(x^k)) - g(\tau^{-1}(y^k))\right\|_\infty && \text{(C.12)}\\
&\leq L\,\hat{d}(\tau^{-1}(x^k), \tau^{-1}(y^k)) && \text{(C.13)}\\
&\leq LM^k \left\|\tau^{-1}(x^k) - \tau^{-1}(y^k)\right\|_\infty && \text{(C.14)}
\end{aligned}$$

holds for all $x^k, y^k \in W'$. Hence $g' : W' \to \mathbb{R}^n$ is Lipschitz with respect to the ℓ_∞ distance, and it satisfies $g'(W') = g(W)$.

To complete the proof of the lemma, we proceed to construct the required τ. The essential idea is to *puncture* the M-ary expansion of x^k such that any component has *at most* k consecutive nonzero digits. For notational convenience, define $r : \mathbb{N} \to \{0, \ldots, k\}$ and $\eta_j : \mathbb{Z}_M^k \to \mathbb{Z}_M^{k+1}$ for $j = 0, \ldots, k$ as follows:

$$r(i) = i - \left\lfloor \frac{i}{k+1} \right\rfloor (k+1), \tag{C.15}$$

$$\eta_j(b^k) = (b_1, \ldots, b_j, 0, b_{j+1}, \ldots, b_k)^{\mathrm{T}}. \tag{C.16}$$

Define

$$\tau(x^k) = \sum_{i\in\mathbb{N}} \eta_{r(i)}((x^k)_i)M^{-i}. \tag{C.17}$$

A schematic illustration of τ in terms of the M-ary expansion is given in Fig. C.1.

$$x^k = \begin{pmatrix} (x_1)_1 & (x_2)_1 & \dots \\ \vdots & \vdots & \dots \\ (x_k)_1 & (x_k)_1 & \dots \end{pmatrix} \xrightarrow{\tau}$$

$$\tau(x^k) = \begin{pmatrix} 0 & (x_1)_2 & \dots & (x_1)_k & 0 & \dots \\ (x_1)_1 & 0 & \dots & (x_2)_k & (x_1)_{k+1} & \dots \\ (x_2)_1 & (x_2)_2 & \dots & (x_3)_k & (x_2)_{k+1} & \dots \\ \vdots & \vdots & \dots & \vdots & \vdots & \dots \\ (x_{k-1})_1 & (x_{k-1})_2 & \dots & (x_k)_k & (x_{k-1})_{k+1} & \dots \\ (x_k)_1 & (x_k)_2 & \dots & 0 & (x_k)_{k+1} & \dots \end{pmatrix}$$

Figure C.1: A schematic illustration of τ in terms of M-ary expansions.

Next we show that τ satisfies the expansiveness condition in (C.11). For any $x^k \neq y^k \in [0,1]^k$, let $\hat{d}(x^k, y^k) = M^{-l}$. Then by definition, for some $m \in \{1, \dots, k\}$, $\hat{d}(x_m, y_m) = M^{-l}$ and $l = \min\{i \in \mathbb{N} : (x_m)_i \neq (y_m)_i\}$. Without loss of generality, assume that $(x_m)_i = 1$ and $(y_m)_i = 0$. Then by construction of τ, there are no more than k consecutive nonzero digits in $\tau(x)$ or $\tau(y)$. Since the worst case is that $(x_m)_i$ and $(y_m)_i$ are followed by k 0's and k $(M-1)$'s respectively, we have

$$\left\|\tau(x^k) - \tau(x^k)\right\|_\infty \geq M^{-(k+l+1)} \tag{C.18}$$

$$= M^{-(1+k)}\hat{d}(x^k, y^k) \tag{C.19}$$

which completes the proof of (C.11). □

C.6 Proof of Lemma 26

Proof. Recall that $l = k + j$ where

$$j = \lceil \delta n \rceil. \tag{C.20}$$

Let us write $q : [S]_m \to [0,1]^l$ as

$$q = \begin{bmatrix} p \\ \tau \end{bmatrix}, \tag{C.21}$$

where $p : [S]_m \to [0,1]^k$ and $\tau : [S]_m \to [0,1]^j$. The idea to construct q that satisfies (5.197) is the following: for those *sufficiently separated* x and y in $[S]_m$, their images under p will be kept separated accordingly. For those x and y that are close together,

they can be resolved by the mapping τ. Suppose p and τ are constructed so that

$$\|p(x) - p(y)\|_\infty \geq 2^{-m} \tag{C.22}$$

holds for all x, y such that $\|x - y\|_\infty > 2^{-m}$ and

$$\|\tau(x) - \tau(y)\|_\infty \geq 2^{-1/\delta} \tag{C.23}$$

holds for all x, y such that $\|x - y\|_\infty = 2^{-m}$. Then by (C.21),

$$\|q(x)\|_\infty = \max\left\{\|p(x)\|_\infty, \|\tau(x)\|_\infty\right\}, \tag{C.24}$$

we arrive at a q that satisfies the required (5.190) for all distinct x and y in $[S]_m$.

First construct p. For $x \in [S]_m$, define

$$\mathcal{D}(x) = \{y \in S : [y]_m = x\} \neq \emptyset. \tag{C.25}$$

Then we could choose[1] one vector in $\mathcal{D}(x)$ to correspond with x and denote it by $\tilde{x}$. Since $g : T \to S$ is bijective, let the inverse of g be $f : S \to T$. By the 2^{-m}-stability of g on T,

$$\|a - b\|_\infty \leq 2^{-m} \Rightarrow \|g(a) - g(b)\|_\infty \leq 2^{-m} \tag{C.26}$$

holds for all $a, b \in T$. By contraposition,

$$\|u - v\|_\infty > 2^{-m} \Rightarrow \|f(u) - f(v)\|_\infty > 2^{-m} \tag{C.27}$$

holds for all $u, v \in S$.

Define $p(x) = f(\tilde{x})$. Consider $x \neq y \in [S]_m$. Note that $\|x - y\|_\infty$ is an integer multiple of 2^{-m} and $\|x - y\|_\infty \geq 2^{-m}$. If $\|x - y\|_\infty > 2^{-m}$, i.e., $\|x - y\|_\infty \geq 2^{-(m-1)}$, we have $|x_i - y_i| \geq 2^{-(m-1)}$ for some $i \in \{1, \ldots, n\}$. Without loss of generality, assume that

$$y_i \geq x_i + 2^{-(m-1)} \tag{C.28}$$

Also notice that by the choice of $\tilde{x}$ and $\tilde{y}$,

$$x_i \leq \tilde{x}_i < x_i + 2^{-m}, \tag{C.29}$$

$$y_i \leq \tilde{y}_i < y_i + 2^{-m}, \tag{C.30}$$

Then

$$\|\tilde{x} - \tilde{y}\|_\infty \geq |\tilde{x}_i - \tilde{y}_i| \tag{C.31}$$

$$\geq \tilde{y}_i - y_i + y_i - x_i + x_i - \tilde{x}_i \tag{C.32}$$

$$> 2^{-m}. \tag{C.33}$$

[1] The set $\mathcal{D}(x)$ is, in general, uncountable. Hence this construction of p is valid by adopting the Axiom of Choice.

where the last inequality follows from (C.28) – (C.30). Hence by (C.27), (C.22) holds for all x, y such that $\|x - y\|_\infty > 2^{-m}$.

We proceed to construct $\tau : [S]_m \to [0,1]^j$ by a *graph-theoretical* argument: consider a undirected graph (G, E) where the vertex set is $\mathsf{G} = [S]_m$, and the edge set $\mathsf{E} \subset [S]_m \times [S]_m$ is formed according to

$$(x, y) \in \mathsf{E} \Leftrightarrow \|x - y\|_\infty = 2^{-m} \tag{C.34}$$

Note that for any x in G, the number of possible y's such that $(x, y) \in \mathsf{E}$ is at most $2^n - 1$, since each component of y and x differ by at most the least significant bit. Therefore the *maximum degree* of this graph is upper bounded by $2^n - 1$, i.e.

$$\Delta(\mathsf{G}) \leq 2^n - 1. \tag{C.35}$$

The mapping τ essentially assigns a different j-dimensional real vector to each of the vertices in G. We choose these vectors from the following "grid":

$$\mathcal{L}_M = \left\{ \frac{1}{M} z^j : z^j \in \mathbb{Z}_M^j \right\} \subset [0,1]^j, \tag{C.36}$$

that is, the grid obtained by dividing each dimension of $[0,1]^j$ uniformly into M segments, where $M \in \mathbb{N}$. Hence $|\mathcal{L}_M| = M^l$. Choose $M = \lceil 2^{1/\delta} \rceil$, then by (C.20) and (C.35),

$$|\mathcal{L}_M| = \Delta(\mathsf{G}) + 1. \tag{C.37}$$

We define τ by coloring the graph G using colors from $\mathcal{L}_M$, such that neighboring vertices have different colors. This is possible in view of Lemma 27 and (C.37). Denote by $\tau(x)$ the color of $x \in \mathsf{G}$. Since the vectors in $\mathcal{L}_M$ are separated by at least M^{-1} in ℓ_∞ distance, by the construction of τ and G, (C.23) holds for all x, y such that $\|x - y\|_\infty = 2^{-m}$. This finishes the construction of τ and the proof of the lemma. □

Appendix D

Distortion-rate tradeoff of Gaussian inputs

In this appendix we show the expressions (6.30) – (6.32) for the minimal distortion, thereby completing the proof of Theorem 46.

D.1 Optimal encoder

Plugging the rate-distortion function of standard Gaussian

$$R_{X_{\mathsf{G}}}(D) = \frac{1}{2}\log^{+}\frac{1}{D} \tag{D.1}$$

into (6.6) yields the equality in (6.30).

D.2 Optimal linear encoder

We compute the minimal distortion $D^{*}_{\mathrm{L}}(X, R, \sigma^2)$ achievable with the optimal linear encoder. Let the sensing matrix by $\mathbf{H}$. Since X^n and $Y^k = \mathbf{H}X^n + \sigma N^k$ are jointly Gaussian, the conditional distribution of X^n given Y^k is $\mathcal{N}(\hat{X}^n, \Sigma_{X^n|Y^k})$, where

$$\hat{X}^n = \mathbf{H}^{\mathrm{T}}(\mathbf{H}\mathbf{H}^{\mathrm{T}} + \sigma^2\mathbf{I}_k)^{-1}Y^k \tag{D.2}$$

$$\Sigma_{X^n|Y^k} = \mathbf{I}_n - \mathbf{H}^{\mathrm{T}}(\mathbf{H}\mathbf{H}^{\mathrm{T}} + \sigma^2\mathbf{I}_k)^{-1}\mathbf{H} \tag{D.3}$$

$$= (\mathbf{I}_n + \sigma^{-2}\mathbf{H}^{\mathrm{T}}\mathbf{H})^{-1}. \tag{D.4}$$

where we used the matrix inversion lemma. Therefore, the optimal estimator is linear, given by (D.2). Moreover,

$$\mathsf{mmse}(X^n|Y^k) = \mathsf{Tr}(\Sigma_{X^n|Y^k}) \tag{D.5}$$

$$= \mathsf{Tr}((\mathbf{I}_n + \sigma^{-2}\mathbf{H}^{\mathrm{T}}\mathbf{H})^{-1}). \tag{D.6}$$

Choosing the best encoding matrix $\mathbf{H}$ boils down to the following optimization problem:

$$\begin{aligned} \min \ & \mathsf{Tr}((\mathbf{I}_n + \sigma^{-2}\mathbf{H}^{\mathrm{T}}\mathbf{H})^{-1}) \\ \text{s.t.} \ & \mathsf{Tr}(\mathbf{H}^{\mathrm{T}}\mathbf{H}) \leq k \\ & \mathbf{H} \in \mathbb{R}^{k\times n} \end{aligned} \tag{D.7}$$

Let $\mathbf{H}^{\mathrm{T}}\mathbf{H} = \mathbf{U}^{\mathrm{T}}\Lambda\mathbf{U}$, where $\mathbf{U}$ is an $n \times n$ orthogonal matrix and Λ is a diagonal matrix consisting of the eigenvalues of $\mathbf{H}^{\mathrm{T}}\mathbf{H}$, denoted by $\{\lambda_1, \ldots, \lambda_n\} \subset \mathbb{R}_+$. Then

$$\mathsf{Tr}((\mathbf{I}_n + \sigma^{-2}\mathbf{H}^{\mathrm{T}}\mathbf{H})^{-1}) = \sum_{i=1}^{n} \frac{1}{1+\sigma^{-2}\lambda_i} \tag{D.8}$$

$$\geq \frac{n}{1+\sigma^{-2}\frac{\mathsf{Tr}(\mathbf{H}^{\mathrm{T}}\mathbf{H})}{n}} \tag{D.9}$$

$$\geq \frac{n}{1+\frac{k}{n}\sigma^{-2}} \tag{D.10}$$

$$= \frac{n}{1+R\sigma^{-2}}. \tag{D.11}$$

where

- (D.9): by the strict convexity of $x \mapsto \frac{1}{1+\sigma^{-2}x}$ on $\mathbb{R}_+$ and $\mathsf{Tr}(\mathbf{H}^{\mathrm{T}}\mathbf{H}) = \sum_{i=1}^{n} \lambda_i$;
- (D.10): due to the energy constraint.

Hence

$$D^*_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2) \geq \frac{1}{1+R\sigma^{-2}}. \tag{D.12}$$

Next we consider two cases separately:

1. $R \geq 1 (k \geq n)$: the lower bound in (D.12) can be achieved by

$$\mathbf{H} = \begin{bmatrix} \sqrt{R}\mathbf{I}_n \\ 0 \end{bmatrix}. \tag{D.13}$$

2. $R < 1 (k < n)$: the lower bound in (D.12) is *not* achievable. This is because to achieve equality in (D.9), all λ_i be equal to R; however, $\mathsf{rank}(\mathbf{H}^{\mathrm{T}}\mathbf{H}) \leq \mathsf{rank}(\mathbf{H}) \leq k < n$ implies that at least $n-k$ of them are zero. Therefore the lower bound (D.11) can be further improved to:

$$\mathsf{Tr}((\mathbf{I}_n + \sigma^{-2}\mathbf{H}^{\mathrm{T}}\mathbf{H})^{-1}) = n - k + \sum_{\lambda_i > 0} \frac{1}{1+\sigma^{-2}\lambda_i} \tag{D.14}$$

$$\geq n - k + \frac{k}{1+\sigma^{-2}\frac{\mathsf{Tr}(\mathbf{H}^{\mathrm{T}}\mathbf{H})}{k}} \tag{D.15}$$

$$\geq n - \frac{k}{1+\sigma^2}. \tag{D.16}$$

Hence when $R < 1$,

$$D_{\mathrm{L}}^{*}(X_{\mathsf{G}}, R, \sigma^2) \geq 1 - \frac{R}{1+\sigma^2}, \tag{D.17}$$

which can be achieved by

$$\mathbf{H} = \begin{bmatrix} \mathbf{I}_k & 0 \end{bmatrix}, \tag{D.18}$$

that is, simply keeping the first k coordinates of X^n and discarding the rest.

Therefore the equality in (6.31) is proved.

D.3 Random linear encoder

We compute the distortion $D_{\mathrm{L}}(X, R, \sigma^2)$ achievable with random linear encoder $\boldsymbol{A}$. Recall that $\boldsymbol{A}$ has i.i.d. entries with zero mean and variance $\frac{1}{n}$. By (D.6),

$$\frac{1}{n}\mathsf{mmse}(X^n|\boldsymbol{A}X^n + \sigma N^k, \boldsymbol{A}) = \frac{1}{n}\mathbb{E}\left[\mathsf{Tr}((\mathbf{I}_n + \sigma^{-2}\boldsymbol{A}^{\mathrm{T}}\boldsymbol{A})^{-1})\right] \tag{D.19}$$

$$= \frac{1}{n}\mathbb{E}\left[\sum_{i=1}^{n} \frac{1}{1+\sigma^{-2}\lambda_i}\right], \tag{D.20}$$

where $\{\lambda_1, \ldots, \lambda_n\}$ are the eigenvalues of $\boldsymbol{A}^{\mathrm{T}}\boldsymbol{A}$.

As $n \to \infty$, the empirical distribution of the eigenvalues of $\frac{1}{R}\boldsymbol{A}^{\mathrm{T}}\boldsymbol{A}$ converges weakly to the Marčenko-Pastur law almost surely [34, Theorem 2.35]:

$$\nu_R(\mathrm{d}x) = (1-R)^{+}\delta_0(\mathrm{d}x) + \frac{\sqrt{(x-a)(x-b)}}{2\pi c x}\mathbf{1}_{[a,b]}(x)\mathrm{d}x \tag{D.21}$$

where

$$c = \frac{1}{R}, a = (1-\sqrt{c})^2, b = (1+\sqrt{c})^2. \tag{D.22}$$

Since $\lambda \mapsto \frac{1}{1+\sigma^{-2}\lambda}$ is continuous and bounded, applying the dominated convergence theorem to (D.20) and integrating with respect to ν_R gives

$$D_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2) = \lim_{n\to\infty}\frac{1}{n}\mathsf{mmse}(X^n|\boldsymbol{A}X^n + \sigma N^k, \boldsymbol{A})$$

$$= \int \frac{1}{1+\sigma^{-2}Rx}\nu_R(\mathrm{d}x) \tag{D.23}$$

$$= \frac{1}{2}\left(1 - R - \sigma^2 + \sqrt{(1-R)^2 + 2(1+R)\sigma^2 + \sigma^4}\right), \tag{D.24}$$

where (D.24) follows from [34, (1.16)].

Next we verify that the formula in Claim 1 which was based on replica calculations coincides with (D.24) in the Gaussian case. Since in this case $\mathsf{mmse}(X_{\mathsf{G}}, \mathsf{snr}) = \frac{1}{1+\mathsf{snr}}$,

(6.13) becomes

$$\frac{1}{\eta} = 1 + \frac{1}{\sigma^2}\mathsf{mmse}(X, \eta R\sigma^{-2}) \tag{D.25}$$

$$= 1 + \frac{1}{\sigma^2 + \eta R} \tag{D.26}$$

whose unique positive solution is given by

$$\eta_\sigma = \frac{R - 1 - \sigma^2 + \sqrt{(1-R)^2 + 2(1+R)\sigma^2 + \sigma^4}}{2R} \tag{D.27}$$

which lies in $(0, 1)$. According to (6.12),

$$D_{\mathrm{L}}(X_{\mathsf{G}}, R, \sigma^2) = \mathsf{mmse}(X_{\mathsf{G}}, \sigma^{-2}\eta_\sigma) \tag{D.28}$$

$$= \frac{1}{1 + \sigma^{-2}\eta_\sigma} \tag{D.29}$$

$$= \frac{2\sigma^2}{R - 1 + \sigma^2 + \sqrt{(1-R)^2 + 2(1+R)\sigma^2 + \sigma^4}}, \tag{D.30}$$

which can be verified, after straightforward algebra, to coincide with (D.24).

Bibliography

[1] E. Candés, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489 – 509, Feb. 2006.

[2] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289 – 1306, Apr. 2006.

[3] D. L. Donoho and J. Tanner, "Counting faces of randomly-projected polytopes when the projection radically lowers dimension," *Journal of the American Mathematical Society*, vol. 22, no. 1, pp. 1–53, 2009.

[4] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18 914 – 18 919, Nov. 2009.

[5] Y. Wu and S. Verdú, "Rényi information dimension: Fundamental limits of almost lossless analog compression," *IEEE Transactions on Information Theory*, vol. 56, no. 8, pp. 3721 – 3748, Aug. 2010.

[6] D. L. Donoho, A. Maleki, and A. Montanari, "The noise-sensitivity phase transition in compressed sensing," Department of Statistics, Stanford University, Tech. Rep. 2010-04, Apr. 2010. [Online]. Available: http://arxiv.org/abs/1004.1218

[7] D. Guo, D. Baron, and S. Shamai (Shitz), "A single-letter characterization of optimal noisy compressed sensing," in *Proceedings of the Forty-Seventh Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2009.

[8] G. Reeves and M. Gastpar, "Sampling bounds for sparse support recovery in the presence of noise," in *Proceedings of the 2008 IEEE International Symposium on Information Theory*, Toronto, Canada, Jul. 2008.

[9] G. Reeves and M. Gastpar, "Fundamental tradeoffs for sparsity pattern recovery," 2010, preprint. [Online]. Available: http://arxiv.org/abs/1006.3128

[10] Y. Kabashima, T. Wadayama, and T. Tanaka, "Statistical mechanical analysis of a typical reconstruction limit of compressed sensing," in *Proceedings of 2010 IEEE International Symposium on Information Theory*, Austin, TX, Jun. 2010.

[11] A. M. Tulino, G. Caire, S. Shamai (Shitz), and S. Verdú, "Support recovery with sparsely sampled free random matrices," in *Proceedings of 2011 IEEE International Symposium on Information Theory*, Saint Petersburg, Russia, Aug. 2011.

[12] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1999.

[13] R. Tibshirani, "Regression shrinkage and selection via the LASSO," vol. 58, no. 1, pp. 267–288, 1996.

[14] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.

[15] D. L. Donoho and J. Tanner, "Counting the faces of randomly-projected hypercubes and orthants, with applications," *Discrete and Computational Geometry*, vol. 43, no. 3, pp. 522–541, 2010.

[16] D. L. Donoho and J. Tanner, "Precise undersampling theorems," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 913–924, 2010.

[17] A. Rényi, *Probability Theory.* Amsterdam: North-Holland Publishing Company, 1970.

[18] K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, 2nd ed. New York: Wiley, 2003.

[19] R. Mañé, "On the dimension of the compact invariant sets of certain nonlinear maps," in *Dynamical Systems and Turbulence, Warwick 1980*, ser. Lecture Notes in Mathematics. Springer, 1981, vol. 898, pp. 230 – 242.

[20] A. Ben-Artzi, A. Eden, C. Foias, and B. Nicolaenko, "Hölder continuity for the inverse of Mañé's projection," *Journal of Mathematical Analysis and Applications*, vol. 178, pp. 22 – 29, 1993.

[21] B. R. Hunt and V. Y. Kaloshin, "Regularity of of embeddings of infinite-dimensional fractal sets into finite-dimensional spaces," *Nonlinearity*, vol. 12, no. 5, pp. 1263–1275, 1999.

[22] T. Blumensath and M. E. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Transactions on Information Theory*, vol. 55, no. 4, pp. 1872–1882, 2009.

[23] Y. Peres and B. Solomyak, "Existence of L_q dimensions and entropy dimension for self-conformal measures," *Indiana University Mathematics Journal*, vol. 49, no. 4, pp. 1603–1621, 2000.

[24] A. Rényi, "On the dimension and entropy of probability distributions," *Acta Mathematica Hungarica*, vol. 10, no. 1 – 2, Mar. 1959.

[25] T. Kawabata and A. Dembo, "The rate-distortion dimension of sets and measures," *IEEE Transactions on Information Theory*, vol. 40, no. 5, pp. 1564 – 1572, Sep. 1994.

[26] A. Guionnet and D. Shlyakhtenko, "On classical analogues of free entropy dimension," *Journal of Functional Analysis*, vol. 251, no. 2, pp. 738 – 771, Oct. 2007.

[27] Y. Wu and S. Verdú, "Fundamental Limits of Almost Lossless Analog Compression," in *Proceedings of 2009 IEEE International Symposium on Information Theory*, Seoul, Korea, June 2009.

[28] Y. Wu and S. Verdú, "MMSE dimension," *IEEE Transactions on Information Theory*, no. 8, pp. 4857 – 4879, Aug. 2011.

[29] Y. Wu, S. Shamai (Shitz), and S. Verdú, "Degrees of Freedom of Interference Channel: a General Formula," in *Proceedings of 2011 IEEE International Symposium on Information Theory*, Saint Petersburg, Russia, August 2011.

[30] Y. B. Pesin, *Dimension Theory in Dynamical Systems: Contemporary Views and Applications.* Chicago, IL: University of Chicago Press, 1997.

[31] A. N. Kolmogorov, "On the shannon theory of information transmission in the case of continuous signals," *IEEE Transactions on Information Theory*, vol. 2, no. 4, pp. 102–108, Dec. 1956.

[32] A. N. Kolmogorov, "Theory of transmission of information," *American Mathematical Society Translations: Series 2*, no. 33, pp. 291–321, 1963.

[33] B. R. Hunt and V. Y. Kaloshin, "How projections affect the dimension spectrum of fractal measures," *Nonlinearity*, vol. 10, pp. 1031–1046, 1997.

[34] A. M. Tulino and S. Verdú, "Random matrix theory and wireless communications," *Foundations and Trends in Communications and Information Theory*, vol. 1, no. 1, Jun. 2004.

[35] P. Mattila, *Geometry of Sets and Measures in Euclidean Spaces: Fractals and Rectifiability.* Cambridge, United Kingdom: Cambridge University Press, 1999.

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

[37] V. Strassen, "The existence of probability measures with given marginals," *Annals of Mathematical Statistics*, vol. 36, no. 2, pp. 423–439, 1965.

[38] E. Çinlar, *Probability and Stochastics*. New York: Springer, 2011.

[39] I. Csiszár, "On the dimension and entropy of order α of the mixture of probability distributions," *Acta Mathematica Hungarica*, vol. 13, no. 3 – 4, pp. 245 – 255, Sep. 1962.

[40] R. S. Strichartz, "Self-similar measures and their fourier transforms III," *Indiana University Mathematics Journal*, vol. 42, 1993.

[41] C. Beck, "Upper and lower bounds on the renyi dimensions and the uniformity of multifractals," *Physica D: Nonlinear Phenomena*, vol. 41, no. 1, pp. 67–78, 1990.

[42] A. György, T. Linder, and K. Zeger, "On the rate-distortion function of random vectors and stationary sources with mixed distributions," *IEEE Transactions on Information Theory*, vol. 45, pp. 2110 – 2115, Sep. 1999.

[43] T. Linder and R. Zamir, "On the asymptotic tightness of the Shannon lower bound," *IEEE Transactions on Information Theory*, vol. 40, no. 6, pp. 2026 – 2031, Nov. 1994.

[44] P. Zador, "Asymptotic quantization error of continuous signals and the quantization dimension," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 139–149, 1982.

[45] S. Graf and H. Luschgy, "Foundations of quantization for probability distributions," 2000.

[46] K. Falconer, *Techniques in Fractal Geometry*. Chichester, UK: John Wiley, 1997.

[47] Y. Peres and B. Solomyak, "Self-similar measures and intersections of Cantor sets," *Transactions of the American Mathematical Society*, vol. 350, no. 10, pp. 4065–4087, 1998.

[48] D. Feng, N. Nguyen, and T. Wang, "Convolutions of equicontractive self-similar measures on the line," *Illinois Journal of Mathematics*, vol. 46, no. 4, pp. 1339–1351, 2002.

[49] J. E. Hutchinson, "Fractals and self-similarity," *Indiana University Mathematics Journal*, vol. 30, no. 5, pp. 713–747, 1981.

[50] Y. Peres, W. Schlag, and B. Solomyak, "Sixty years of Bernoulli convolutions," in *Fractal geometry and stochastics II*, ser. Progress in Probability, C. Bandt, S. Graf, and M. Zähle, Eds. Springer, 2000, vol. 46, pp. 39–65.

[51] J. S. Geronimo and D. P. Hardin, "An exact formula for the measure dimensions associated with a class of piecewise linear maps," *Constructive Approximation*, vol. 5, no. 1, pp. 89–98, 1989.

[52] N. Patzschke, "Self-conformal multifractal measures," *Advances in Applied Mathematics*, vol. 19, no. 4, pp. 486–513, 1997.

[53] K. S. Lau and S. M. Ngai, "Multifractal measures and a weak separation condition," *Advances in Mathematics*, vol. 141, no. 1, pp. 45–96, 1999.

[54] Y. Wu, S. Shamai (Shitz), and S. Verdú, "A general formula for degrees of freedom of the interference channel," submitted to *IEEE Transactions on Information Theory*, 2011.

[55] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, Inc., 1982.

[56] I. Csiszár, "Generalized cutoff rates and Rényi's information measures," *IEEE Transactions on Information Theory*, vol. 41, no. 1, pp. 26 – 34, 1995.

[57] P. N. Chen and F. Alajaji, "Csiszár's cutoff rates for arbitrary discrete sources," *IEEE Transactions on Information Theory*, vol. 47, no. 1, pp. 330 – 338, Janurary 2001.

[58] H. Shimokawa, "Rényi's entropy and error exponent of source coding with countably infinite alphabet," in *Proceedings of 2006 IEEE International Symposium on Information Theory*, Seattle, WA, Jul. 2006.

[59] C. Chang and A. Sahai, "Universal quadratic lower bounds on source coding error exponents," in *Conference on Information Sciences and Systems*, 2007.

[60] R. G. Gallager, *Information Theory and Reliable Communication*. New York, NY: John Wiley & Sons, Inc., 1968.

[61] M. S. Pinsker, *Information and Information Stability of Random Variables and Processes*. San Francisco, CA: Holden-Day, 1964.

[62] K. T. Alligood, T. Sauer, and J. A. Yorke, *Chaos: an Introduction to Dynamical Systems*. Springer Verlag, 1996.

[63] Y. Wu and S. Verdú, "MMSE dimension," in *Proceedings of 2010 IEEE International Symposium on Information Theory*, Austin, TX, June 2010.

[64] Y. Wu and S. Verdú, "Functional properties of MMSE and mutual information," submitted to *IEEE Transactions on Information Theory*, 2010.

[65] D. Guo, S. Shamai (Shitz), and S. Verdú, "Mutual Information and Minimum Mean-Square Error in Gaussian Channels," *IEEE Transactions on Information Theory*, vol. 51, no. 4, pp. 1261 – 1283, Apr. 2005.

[66] D. Guo, Y. Wu, S. Shamai (Shitz), and S. Verdú, "Estimation in Gaussian Noise: Properties of the Minimum Mean-square Error," *IEEE Transactions on Information Theory*, vol. 57, no. 4, pp. 2371 – 2385, Apr. 2011.

[67] M. S. Pinsker, "Optimal filtering of square-integrable signals in Gaussian noise," *Problemy Peredachi Informatsii*, vol. 16, no. 2, pp. 52–68, 1980.

[68] M. Nussbaum, "Minimax risk: Pinsker bound," in *Encyclopedia of Statistical Sciences*. New York: Wiley, 1999, pp. 451–460.

[69] B. Z. Bobrovsky and M. Zakai, "Asymptotic a priori estimates for the error in the nonlinear filtering problem," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 371–376, 1982.

[70] J. K. Ghosh, *Higher Order Asymptotics*. Hayward, CA: Institute of Mathematical Statistics, 1994.

[71] J. K. Ghosh, B. K. Sinha, and S. N. Joshi, "Expansions for posterior probability and integrated Bayes risk," in *Statistical Decision Theory and Related Topics III*. New York, NY: Academic Press, 1982, pp. 403–456.

[72] M. V. Burnašev, "Investigation of second order properties of statistical estimators in a scheme of independent observations," *Math. USSR Izvestiya*, vol. 18, no. 3, pp. 439–467, 1982.

[73] P. J. Huber, *Robust Statistics*. New York, NY: Wiley-Interscience, 1981.

[74] L. D. Brown, "Admissible estimators, recurrent diffusions, and insoluble boundary value problems," *Annals of Mathematical Statistics*, vol. 42, no. 3, pp. 855–903, 1971.

[75] V. V. Prelov and E. C. van der Meulen, "Asymptotics of Fisher Information under Weak Perturbation," *Problems of Information Transmission*, vol. 31, no. 1, pp. 14 – 22, 1995.

[76] M. S. Pinsker, V. V. Prelov, and S. Verdú, "Sensitivity of channel capacity," *IEEE Transactions on Information Theory*, vol. 41, no. 6, pp. 1877–1888, Nov. 1995.

[77] V. V. Prelov and S. Verdú, "Second-order asymptotics of mutual information," *IEEE Transactions on Information Theory*, vol. 50, no. 8, pp. 1567–1580, Aug. 2004.

[78] A. J. Stam, "Some inequalities satisfied by the quantities of information of Fisher and Shannon," *Information and Control*, vol. 2, no. 2, pp. 101–112, 1959.

[79] H. L. Van Trees, *Detection, Estimation, and Modulation theory, Part I*. New York, NY: Wiley, 1968.

[80] L. D. Brown and L. Gajek, "Information inequalities for the Bayes risk," *Annals of Statistics*, vol. 18, no. 4, pp. 1578–1594, 1990.

[81] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd ed. New York, NY: Springer, 1998.

[82] C. Weidmann and M. Vetterli, "Rate distortion behavior of sparse sources," Aug. 2010, in revision for *IEEE Transactions on Information Theory*.

[83] A. Lozano, A. M. Tulino, and S. Verdú, "Optimum power allocation for parallel Gaussian channels with arbitrary input distributions," *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 3033–3051, Jul. 2006.

[84] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed. New York, NY: McGraw-Hill, 1976.

[85] K. Knopp and F. Bagemihl, *Theory of Functions, Parts I and II*. Mineola, NY: Dover Publications, 1996.

[86] Y. Wu and S. Verdú, "Functional properties of MMSE," in *Proceedings of 2010 IEEE International Symposium on Information Theory*, Austin, TX, June 2010.

[87] R. Zamir, "A proof of the Fisher information inequality via a data processing argument," *IEEE Transactions on Information Theory*, vol. 44, no. 3, pp. 1246–1250, Aug. 1998.

[88] S. Verdú, "Mismatched estimation and relative entropy," *IEEE Transactions on Information Theory*, vol. 56, no. 8, pp. 3712–3720, Jul. 2010.

[89] T. Weissman, "The relationship between causal and noncausal mismatched estimation in continuous-time AWGN channels," *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4256–4273, Aug. 2010.

[90] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379 – 423, 623 – 656, 1948.

[91] J. R. Munkres, *Topology*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 2000.

[92] A. Montanari and E. Mossel, "Smooth compression, Gallager bound and nonlinear sparse graph codes," in *Proceedings of 2008 IEEE International Symposium on Information Theory*, Toronto, Canada, Jul. 2008.

[93] E. Candés, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, Aug. 2006.

[94] R. Calderbank, S. Howard, and S. Jafarpour, "Construction of a large class of deterministic matrices that satisfy a statistical isometry property," *IEEE Journal of Selected Topics in Signal Processing, special issue on Compressed Sensing*, vol. 29, no. 4, pp. 358–374, 2009.

[95] W. Xu and B. Hassibi, "Compressed sensing over the Grassmann manifold: A unified analytical framework," in *Proceedings of the Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing*, 2008, pp. 562–567.

[96] T. T. Cai, L. Wang, and G. Xu, "New bounds for restricted isometry constants," *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4388–4394, 2010.

[97] W. Li, "The sharp Lipschitz constants for feasible and optimal solutions of a perturbed linear program," *Linear algebra and its applications*, vol. 187, pp. 15–40, 1993.

[98] P. Halmos, *Naive Set Theory.* Princeton, NJ: D. Van Nostrand Company, 1960.

[99] K. Kuratowski, *Topology, vol. I.* New York, NY: Academic Press, 1966.

[100] J. T. Schwartz, *Nonlinear Functional Analysis.* New York, NY: Gordon and Breach Science Publishers, 1969.

[101] G. B. Folland, *Real Analysis: Modern Techniques and Their Applications*, 2nd ed. New York, NY: Wiley-Interscience, 1999.

[102] H. Steinhaus, "Sur les distances des points des ensembles de mesure positive," *Fundamenta Mathematicae*, vol. 1, pp. 93 – 104, 1920.

[103] G. J. Minty, "On the extension of Lipschitz, Lipschitz-Hölder continuous, and monotone functions," *Bulletin of the American Mathematical Society*, vol. 76, no. 2, pp. 334 – 339, 1970.

[104] H. Federer, *Geometric Measure Theory.* New York, NY: Springer-Verlag, 1969.

[105] D. Preiss, "Geometry of measures in $\mathbb{R}^n$: Distribution, rectifiability, and densities," *Annals of Mathematics*, vol. 125, no. 3, pp. 537 – 643, 1987.

[106] G. A. Edgar, *Integral, Probability, and Fractal Measures.* New York, NY: Springer, 1997.

[107] M. Keane and M. Smorodinsky, "Bernoulli schemes of the same entropy are finitarily isomorphic," *Annals of Mathematics*, vol. 109, no. 2, pp. 397–406, 1979.

[108] A. Del Junco, "Finitary codes between one-sided Bernoulli shifts," *Ergodic Theory and Dynamical Systems*, vol. 1, pp. 285 – 301, 1981.

[109] J. G. Sinaĭ, "A weak isomorphism of transformations with an invariant measure," *Doklady Akademii Nauk SSSR*, vol. 147, pp. 797 – 800, 1962, English translation, American Mathematical Society Translations (2) 57 (1966), 123 – 143.

[110] K. E. Petersen, *Ergodic Theory.* Cambridge, United Kingdom: Cambridge University Press, 1990.

[111] W. H. Schikhof, *Ultrametric Calculus : An Introduction to p-adic Analysis.* New York, NY: Cambridge University Press, 2006.

[112] M. A. Martin and P. Mattila, "k-dimensional regularity classifications for s-fractals," *Transactions of the American Mathematical Society*, vol. 305, no. 1, pp. 293 – 315, 1988.

[113] A. Edelman, "Eigenvalues and condition numbers of random matrices," *SIAM Journal on Matrix Analysis and Applications*, vol. 9, no. 4, pp. 543–560, Oct. 1988.

[114] R. A. Brualdi, *Introductory Combinatorics*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 1999.

[115] Y. Wu, "MMSE dimension and noisy compressed sensing," poster in *Third Annual School of Information Theory*, University of Southern California, Los Angeles CA, Aug. 2010.

[116] Y. Wu and S. Verdú, "Optimal phase transitions in compressed sensing," submitted to *IEEE Transactions on Information Theory*, Jul. 2011.

[117] T. Tanaka, "A statistical-mechanics approach to large-system analysis of CDMA multiuser detectors," *IEEE Transactions on Information Theory*, vol. 48, pp. 2888–2910, Nov. 2002.

[118] D. Guo and S. Verdú, "Randomly spread CDMA: Asymptotics via statistical physics," *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 1983–2010, Jun. 2005.

[119] S. P. Boyd and L. Vandenberghe, *Convex Optimization.* Cambridge, United Kingdom: Cambridge University Press, 2004.

[120] T. Berger, *Rate Distortion Theory: A Mathematical Basis for Data Compression.* Englewood Cliffs, NJ: Prentice-Hall, 1971.

[121] T. Goblick Jr., "Theoretical limitations on the transmission of data from analog sources," *IEEE Transactions on Information Theory*, vol. 11, no. 4, pp. 558–567, 1965.

[122] J. Ziv, "The behavior of analog communication systems," *IEEE Transactions on Information Theory*, vol. 16, no. 5, pp. 587–594, 1970.

[123] A. Maleki, "Approximate message-passing algorithms for compressed sensing," Ph.D. dissertation, Department of Electrical Engineering, Stanford University, Mar. 2011.

[124] J. D. Blanchard, C. Cartis, J. Tanner, and A. Thompson, "Phase transitions for greedy sparse approximation algorithms," *Applied and Computational Harmonic Analysis*, pp. 188 – 203, 2010.

[125] M. Bayati and A. Montanari, "The LASSO risk for Gaussian matrices," *preprint*, 2010. [Online]. Available: http://arxiv.org/abs/1008.2581

[126] M. Rudelson and R. Vershynin, "Smallest singular value of a random rectangular matrix," *Communications on Pure and Applied Mathematics*, vol. 62, pp. 1707–1739, 2009.

[127] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Transactions on Information Theory*, vol. 49, no. 12, pp. 3320–3325, 2003.

[128] S. Foucart and M. J. Lai, "Sparsest solutions of underdetermined linear systems via ℓ_q-minimization for $0 < q < 1$," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 395–407, 2009.

[129] M. E. Davies and R. Gribonval, "Restricted isometry constants where ℓ^p sparse recovery can fail for $0 < p \leq 1$," *IEEE Transactions on Information Theory*, vol. 55, no. 5, pp. 2203–2214, 2009.

[130] J. D. Blanchard, C. Cartis, and J. Tanner, "Compressed sensing: How sharp is the restricted isometry property?" *SIAM Review*, vol. 53, no. 1, pp. 105–125, 2011.

[131] M. Wang, W. Xu, and A. Tang, "On the performance of sparse recovery via ℓ_p-minimization $(0 \leq p \leq 1)$," to appear in *IEEE Transactions on Information Theory*, 2011.

[132] Y. Zhang, "Theory of compressive sensing via ℓ_1-minimization: A non-RIP analysis and extensions," Department of Computational and Applied Mathematics, Rice University, Tech. Rep. TR08–11, 2008.

[133] D. L. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," *Proceedings of the National Academy of Sciences*, vol. 102, no. 27, pp. 9446 – 9451, 2005.

[134] Y. Zhang, "A simple proof for recoverability of ℓ_1-minimization (I): the nonnegativity case," Department of Computational and Applied Mathematics, Rice University, Tech. Rep. TR05-10, 2005.

[135] M. J. Wainwright, "Information-theoretic limitations on sparsity recovery in the high-dimensional and noisy setting," *IEEE Transactions on Information Theory*, vol. 55, no. 12, pp. 5728–5741, Dec. 2009.

[136] A. K. Fletcher, S. Rangan, and V. K. Goyal, "Necessary and sufficient conditions for sparsity pattern recovery," *IEEE Transactions on Information Theory*, vol. 55, no. 12, pp. 5758–5772, Dec. 2009.

[137] E. C. Posner and E. R. Rodemich, "Epsilon entropy and data compression," *Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 2079–2125, Dec. 1971.

[138] S. Ihara, "On the capacity of channels with additive non-gaussian noise*," *Information and Control*, vol. 37, no. 1, pp. 34–39, 1978.

[139] T. Tao, "Sumset and inverse sumset theory for Shannon entropy," *Combinatorics, Probability & Computing*, vol. 19, no. 4, pp. 603–639, 2010.

[140] I. Z. Ruzsa, "An application of graph theory to additive number theory," *Scientia Ser. A*, vol. 3, pp. 97 – 109, 1989.

[141] V. A. Kaimanovich and A. M. Vershik, "Random walks on discrete groups: boundary and entropy," *Annals of Probability*, vol. 11, no. 3, pp. 457–490, 1983.

[142] T. Tao, *An Epsilon of Room, II: pages from year three of a mathematical blog.* Providence, RI: American Mathematical Society, 2011.

[143] M. Madiman, "On the entropy of sums," in *Proceedings of 2008 IEEE Information Theory Workshop*, Porto, Portugal, 2008, pp. 303–307.

[144] T. Tao and V. Vu, "Entropy methods," unpublished notes. [Online]. Available: http://www.math.ucla.edu/~tao/preprints/Expository/chapter_entropy.dvi

[145] I. Z. Ruzsa, "Entropy and sumsets," *Random Structures and Algorithms*, vol. 34, pp. 1–10, Jan. 2009.

[146] M. Madiman and I. Kontoyiannis, "The entropies of the sum and the difference of two IID random variables are not too different," in *Proceedings of 2010 IEEE International Symposium on Information Theory*, Austin, TX, June 2010, pp. 1369–1372.

[147] B. Bukh, "Sums of dilates," *Combinatorics, Probability and Computing*, vol. 17, no. 05, pp. 627–639, 2008.

[148] J. W. S. Cassels, *An Introduction to Diophantine Approximation.* Cambridge, United Kingdom: Cambridge University Press, 1972.

[149] A. Y. Khinchin, *Continued Fractions.* Chicago, IL: The University of Chicago Press, 1964.

[150] L. C. Evans and R. F. Gariepy, *Measure Theory and Fine Properties of Functions.* Boca Raton, FL: CRC Press, 1992.

[151] J. Doob, "Relative limit theorems in analysis," *Journal d'Analyse Mathématique*, vol. 8, no. 1, pp. 289–306, 1960.

[152] E. M. Stein and R. Shakarchi, *Real analysis: measure theory, integration, and Hilbert spaces.* Princeton, NJ: Princeton University Press, 2005.

[153] V. Prokaj, "A Characterization of Singular Measures," *Real Analysis Exchange*, vol. 29, pp. 805–812, 2004.

[154] G. Szegö, *Orthogonal polynomials*, 4th ed. Providence, RI: American Mathematical Society, 1975.