

Optimal Phase Transitions in Compressed Sensing

Yihong Wu, and Sergio Verdú, *Fellow, IEEE*

Abstract—Compressed sensing deals with efficient recovery of analog signals from linear encodings. This paper presents a statistical study of compressed sensing by modeling the input signal as an i.i.d. process with known distribution. Three classes of encoders are considered, namely optimal nonlinear, optimal linear and random linear encoders. 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. The optimal phase transition threshold is determined as a functional of the input distribution and compared to suboptimal thresholds achieved by popular reconstruction algorithms. In particular, we show that 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 in the weak-noise regime.

Keywords: Compressed sensing, Shannon theory, phase transition, Rényi information dimension, MMSE dimension, random matrix, joint source-channel coding.

I. INTRODUCTION

A. Setup

Compressed sensing [3], [4] is a signal processing technique that compresses analog vectors by means of a linear transformation. By leveraging prior knowledge of the signal structure (e.g., sparsity) and by designing efficient nonlinear 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: A real vector $x^n \in \mathbb{R}^n$ is mapped into $y^k \in \mathbb{R}^k$ by an encoder (or compressor) $f: \mathbb{R}^n \rightarrow \mathbb{R}^k$. The decoder (or decompressor) $g: \mathbb{R}^k \rightarrow \mathbb{R}^n$ 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 dimensionality compression ratio, is given by

$$R = \frac{k}{n}. \quad (1)$$

Most of the compressed sensing literature focuses on the setup where

Manuscript received ???; revised ???. This work was supported in part by the National Science Foundation (NSF) under Grant CCF-1016625 and by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under Grant CCF-0939370. The results of this paper were presented in part at the Third Annual School of Information Theory, University of Southern California, Los Angeles CA, August 5–8, 2010 [1] and the IEEE International Symposium on Information Theory, Cambridge, MA, July 1–6, 2012 [2].

Y. Wu is with the Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104 (e-mail: yihongwu@wharton.upenn.edu) and S. Verdú is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: verdu@princeton.edu).




Fig. 1. Compressed sensing: an abstract setup.

- performance is measured on a worst-case basis with respect to x^n .
- the encoder is constrained to be a linear mapping characterized by a $k \times n$ matrix \mathbf{A} , called the *sensing* or *measurement matrix*, which is usually assumed to be random, and known at the decoder.¹
- the decoder is a low-complexity algorithm which is robust with respect to observation noise, for example, decoders based on convex optimizations such as ℓ_1 -minimization [6] and ℓ_1 -penalized least-squares (i.e. LASSO) [7], greedy algorithms such as matching pursuit [8], graph-based iterative decoders such as approximate message passing (AMP) [9], fast iterative shrinkage-thresholding algorithm (FISTA) [10], etc.

In contrast, in this paper we formulate an information-theoretic fundamental limit in the following setup:

- the input vector x^n is random with a known distribution and performance is measured on an average basis.²
- in addition to the performance that can be achieved by the optimal sensing matrix, we also investigate the optimal performance that can be achieved by any *nonlinear* encoder.
- the decoder is *optimal*:³
 - In the noiseless case, it is required to be *Lipschitz continuous* for the sake of robustness;
 - In the noisy case, it is the minimum mean-square error (MMSE) estimator, i.e., the conditional expectation of the input vector given the noisy measurements.

Due to the constraints of actual measuring devices in certain applications of compressed sensing (e.g., MRI [23], high-resolution radar imaging [24]), one does not have the freedom to optimize over all possible sensing matrices. Therefore we consider both optimized as well as random sensing matrices and investigate their respective fundamental limits achieved by the corresponding optimal decoders.

¹Alternative notations have been used to denote the signal dimension and the number of measurements, e.g., (m, n) in [4] and (N, K) in [5].

²Similar Bayesian modeling is followed in some of the compressed sensing literature, for example, [11], [9], [12], [13], [14], [15], [16], [17], [18], [19].

³The performance of optimal decoders for support recovery in the noisy case has been studied in [20], [21], [22] on a worst-case basis.

B. Phase transition

The general goal is to investigate the fundamental tradeoff between reconstruction fidelity and measurement rate as $n \rightarrow \infty$, as a functional of the signal and noise statistics.

When the measurements are noiseless, the 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., independent and identically distributed (i.i.d.) ones, 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 approach one for any sequence of encoder-decoder pairs. Such a phenomenon is known as *phase transition* in statistical physics. In information-theoretic parlance, we say that the *strong converse* holds.

When the measurement is noisy, exact analog signal recovery is obviously impossible and we gauge the reconstruction fidelity by the *noise sensitivity*, defined as the ratio between the 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 the noise sensitivity is bounded for all noise variances, and below which the noise sensitivity blows up as the noise variance tends to zero.

C. Signal model

Sparse vectors, supported on a subspace with dimension smaller than n , play an important role in signal processing and statistical models. A stochastic model that captures sparsity is the following mixture distribution [12], [25], [9], [14], [17], [18], [19]:

$$P = (1 - \gamma)\delta_0 + \gamma P_c, \quad (2)$$

where δ_0 denotes the Dirac measure at 0, P_c is a probability measure absolutely continuous with respect to the Lebesgue measure, and $0 \leq \gamma \leq 1$. Consider a random vector X^n independently drawn from P . By the weak law of large numbers, $\frac{1}{n} \|X^n\|_0 \xrightarrow{\mathbb{P}} \gamma$, where the “ ℓ_0 norm” $\|\cdot\|_0$ denotes the number of non-zeros of a vector. This corresponds to the regime of *proportional (or linear) sparsity*. In (2), the weight on the continuous part γ parametrizes the signal sparsity and P_c serves as the prior distribution of non-zero entries.

Generalizing (2), we henceforth consider *discrete-continuous mixed* distributions (i.e., *elementary distributions* [26]):

$$P_X = (1 - \gamma)P_d + \gamma P_c, \quad (3)$$

where P_d is a discrete probability measure and P_c is an absolutely continuous probability measure. For simplicity we focus on i.i.d. input processes in this paper. Note that apart from sparsity, there are other signal structures that have been previously explored in the compressed sensing literature. For example, the so-called *simple* signal in infrared absorption spectroscopy [27, Example 3, p. 914] is such that each entry of the signal vector is constrained to lie in the unit interval, with most of the entries saturated at the boundaries (0 or 1). Similar to the rationale that leads to (2), an appropriate statistical

model for simple signals is a mixture of a Bernoulli distribution and an absolutely continuous distribution supported on the unit interval, which is a particular instance of (3). Although most of the results in the present paper hold for arbitrary input distributions, with no practical loss of generality, we will be focusing on discrete-continuous mixtures (i.e., without singular components) because of their relevance to compressed sensing applications.

D. Main contributions

We introduced the framework of almost lossless analog compression in [12] as a Shannon-theoretic formulation of noiseless compressed sensing. Under regularity conditions on the encoder or the decoder, [12] derives various coding theorems for the minimal measurement rate involving the information dimension of the input distribution, introduced by Alfréd Rényi in 1959 [28]. Along with the Minkowski and MMSE dimension, we summarize a few relevant properties of Rényi information dimension in Section II. The most interesting regularity constraints are the linearity of the compressor and Lipschitz continuity (robustness) of the decompressor, which are considered *separately* in [12]. Section III gives a brief summary of the non-asymptotic version of these results. In addition, in this paper we also consider the fundamental limit when linearity and Lipschitz continuity are imposed *simultaneously*. For i.i.d. discrete-continuous mixtures, we show that the minimal measurement rate is given by the input information dimension, i.e., the weight γ of the absolutely continuous part. Moreover, the Lipschitz constant of the decoder can be chosen independently of n , as a function of the gap between the measurement rate and γ . This results in the optimal phase transition threshold of error probability in noiseless compressed sensing.

Our main results are presented in Section IV, which deals with the case where the measurements are corrupted by additive Gaussian noise. We consider three formulations of noise sensitivity: optimal nonlinear, optimal linear and random linear (with i.i.d. entries) encoder and the associated optimal decoder. In the case of i.i.d. 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 measurement matrices. Invoking the results in [29], we show that the calculation of the reconstruction error with random measurement matrices based on *heuristic* replica methods in [14] predicts the correct phase transition threshold. These results also serve as a rigorous verification of the replica calculations in [14] in the high-SNR regime (up to $o(\sigma^2)$ as the noise variance σ^2 vanishes).

The fact that randomly chosen sensing matrices turn out to incur no penalty in phase transition threshold with respect to optimal nonlinear encoders lends further importance to the conventional compressed sensing setup described in Section I-A.

In Section V, we compare the optimal phase transition threshold to the suboptimal threshold of several practical reconstruction algorithms under various input distributions.

In particular, we demonstrate that the thresholds achieved by the ℓ_1 -minimization decoder and the AMP decoder [25], [13] lie far from the optimal boundary, especially in the highly sparse regime which is most relevant to compressed sensing applications.

II. THREE DIMENSIONS

In this section we introduce three dimension concepts for sets and probability measures involved in various coding theorems in Sections III and IV.

A. Information dimension

A key concept in fractal geometry, in [28] Rényi defined the *information dimension* (also known as the *entropy dimension* [30]) of a probability distribution. It measures the rate of growth of the entropy of successively finer discretizations.

Definition 1. Let X be a real-valued random variable. Let $m \in \mathbb{N}$. The *information dimension* of X is defined as

$$d(X) = \lim_{m \rightarrow \infty} \frac{H(\lfloor mX \rfloor)}{\log m}. \quad (4)$$

If the limit in (4) does not exist, the \liminf and \limsup are called lower and upper information dimensions of X respectively, denoted by $\underline{d}(X)$ and $\bar{d}(X)$.

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

The information dimension of X is finite if and only if the mild condition

$$H(\lfloor X \rfloor) < \infty \quad (5)$$

is satisfied [12]. A sufficient condition for $d(X) < \infty$ is $\mathbb{E}[\log(1 + |X|)] < \infty$, much milder than finite mean or finite variance.

Equivalent definitions of information dimension include:⁴

- For an integer $M \geq 2$, write the M -ary expansion of X as

$$X = \lfloor X \rfloor + \sum_{i \geq 1} (X)_i M^{-i}, \quad (6)$$

where the i^{th} M -ary digit $(X)_i \triangleq \lfloor M^i X \rfloor - M \lfloor M^{i-1} X \rfloor$ is a discrete random variable taking values in $\{0, \dots, M-1\}$. Then $d(X)$ is the normalized entropy rate of the digits $\{(X)_i\}$:

$$d(X) = \lim_{m \rightarrow \infty} \frac{H((X)_1, \dots, (X)_m)}{m \log M}. \quad (7)$$

- Denote by $B(x, \delta)$ the open ball of radius δ centered at x . Then (see [31, Definition 4.2] and [12, Appendix A])

$$d(X) = \lim_{\delta \downarrow 0} \frac{\mathbb{E}[\log P_X(B(X, \delta))]}{\log \delta}. \quad (8)$$

⁴The lower and upper information dimension are given by the \liminf and \limsup respectively.

- The rate-distortion function of X with mean-square error distortion is given by

$$R_X(D) = \inf_{\mathbb{E}[\hat{X} - X]^2 \leq D} I(X; \hat{X}). \quad (9)$$

Then [32, Proposition 3.3]

$$d(X) = \lim_{D \downarrow 0} \frac{R_X(D)}{\frac{1}{2} \log \frac{1}{D}}. \quad (10)$$

- Let $N \sim \mathcal{N}(0, 1)$ be independent of X . The mutual information $I(X; \sqrt{\text{snr}}X + N)$ is finite if and only if (5) holds [33]. Then [34], [33]

$$d(X) = \lim_{\text{snr} \rightarrow \infty} \frac{I(X; \sqrt{\text{snr}}X + N)}{\frac{1}{2} \log \text{snr}}. \quad (11)$$

The alternative definition in (7) implies that $d(X^n) \leq n$ (as long as it is finite). For discrete-continuous mixtures, the information dimension is given by the weight of the absolutely continuous part.

Theorem 1 ([28]). Assume that X has a discrete-continuous mixed distribution as in (3). If $H(\lfloor X \rfloor) < \infty$, then

$$d(X) = \gamma. \quad (12)$$

In the presence of a singular component, the information dimension does not admit a simple formula in general. One example where the information dimension can be explicitly determined is the *Cantor distribution*, which can be defined via the following ternary expansion

$$X = \sum_{i \geq 1} (X)_i 3^{-i}, \quad (13)$$

where $(X)_i$'s are i.i.d. and equiprobable on $\{0, 2\}$. Then P_X is absolutely singular with respect to the Lebesgue measure and $d(X) = \log_3 2 \approx 0.63$, in view of (7).

B. MMSE dimension

Introduced in [29], the MMSE dimension is an information measure that governs the high-SNR asymptotics of the MMSE in Gaussian noise. Denote the MMSE of estimating X based on Y by

$$\text{mmse}(X|Y) = \inf_f \mathbb{E}[(X - f(Y))^2] \quad (14)$$

$$= \mathbb{E}[(X - \mathbb{E}[X|Y])^2] = \mathbb{E}[\text{var}(X|Y)], \quad (15)$$

where the infimum in (14) is over all Borel measurable f . When Y is related to X through an additive Gaussian noise channel with gain $\sqrt{\text{snr}}$, i.e., $Y = \sqrt{\text{snr}}X + N$ with $N \sim \mathcal{N}(0, 1)$ independent of X , we denote

$$\text{mmse}(X, \text{snr}) = \text{mmse}(X|\sqrt{\text{snr}}X + N). \quad (16)$$

Definition 2. The *MMSE dimension* of X is defined as

$$\mathcal{D}(X) = \lim_{\text{snr} \rightarrow \infty} \text{snr} \cdot \text{mmse}(X, \text{snr}). \quad (17)$$

Useful if the limit in (17) does not exist, the \liminf and \limsup are called lower and upper MMSE dimensions of X respectively, denoted by $\underline{\mathcal{D}}(X)$ and $\bar{\mathcal{D}}(X)$.

It is shown in [29, Theorem 8] that the information dimensions are sandwiched between the MMSE dimensions: if (5) is satisfied, then

$$0 \leq \underline{\mathcal{D}}(X) \leq \underline{d}(X) \leq \bar{d}(X) \leq \bar{\mathcal{D}}(X) \leq 1. \quad (18)$$

For discrete-continuous mixtures, the MMSE dimension coincides with the information dimension:

Theorem 2 ([29, Theorem 15]). *If X has a discrete-continuous mixed distribution as in (3), then $\mathcal{D}(X) = \gamma$.*

It is possible that the MMSE dimension does not exist and the inequalities in (18) are strict. For example, consider the Cantor distribution in (13). Then the product $\text{snr} \cdot \text{mmse}(X, \text{snr})$ oscillates periodically in $\log \text{snr}$ between $\underline{\mathcal{D}}(X) \approx 0.62$ and $\bar{\mathcal{D}}(X) \approx 0.64$ [29, Theorem 16].

C. Minkowski dimension

In fractal geometry, the Minkowski dimension (also known as the box-counting dimension) [35] gauges the fractality of a subset in metric spaces, defined as the exponent with which the covering number grows. The (ϵ -)Minkowski dimension of a probability measure is defined as the lowest Minkowski dimension among all sets with measure at least $1 - \epsilon$ [36].

Definition 3 (Minkowski dimension). Let A be a nonempty bounded subset of \mathbb{R}^n . For $\delta > 0$, denote by $N_A(\delta)$ the δ -covering number of A , i.e., the smallest number of ℓ_2 -balls of radius δ needed to cover A . Define the (upper) Minkowski dimension of A as

$$\overline{\dim}_B A = \limsup_{\delta \rightarrow 0} \frac{\log N_A(\delta)}{\log \frac{1}{\delta}}. \quad (19)$$

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

$$\overline{\dim}_B^\epsilon(\mu) = \inf\{\overline{\dim}_B(A) : \mu(A) \geq 1 - \epsilon\}. \quad (20)$$

Minkowski dimension is always nonnegative and less than the ambient dimension n , with $\dim_B A = 0$ for any finite set A and $\dim_B A = n$ for any bounded set A with nonempty interior. An intermediate example is the middle-third Cantor set C in the unit interval: $\dim_B C = \log_3 2$ [35, Example 3.3].

III. NOISELESS COMPRESSED SENSING

A. Definitions

Definition 4 (Lipschitz continuity). Let $U \subset \mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}^k$. Define⁵

$$\text{Lip}(f) \triangleq \sup_{x \neq y} \frac{\|f(x) - f(y)\|}{\|x - y\|}. \quad (21)$$

If $\text{Lip}(f) \leq L$ for some $L \in \mathbb{R}_+$, we say that f is L -Lipschitz continuous, and $\text{Lip}(f)$ is called the *Lipschitz constant* of f .

⁵Throughout the paper, $\|\cdot\|$ denotes the ℓ_2 norm on the Euclidean space. It should be noted that the proof in the present paper relies crucially on the inner product structure endowed by the ℓ_2 norm. See Remark 5.

Remark 1. $\text{Lip}(\cdot)$ defines a pseudo-norm on the space of all functions.

The Shannon-theoretic fundamental limits of noiseless compressed sensing are defined as follows.

Definition 5. Let X^n be a random vector consisting of independent copies of X . Define the minimum ϵ -achievable rate to be the minimum of $R > 0$ such that there exists a sequence of encoders $f_n : \mathbb{R}^n \rightarrow \mathbb{R}^{\lfloor Rn \rfloor}$ and decoders $g_n : \mathbb{R}^{\lfloor Rn \rfloor} \rightarrow \mathbb{R}^n$, such that

$$\mathbb{P}\{g_n(f_n(X^n)) \neq X^n\} \leq \epsilon \quad (22)$$

for all sufficiently large n . The minimum ϵ -achievable rate is denoted by $R^*(X, \epsilon)$, $R(X, \epsilon)$ and $\hat{R}(X, \epsilon)$ depending on the class of allowable encoders and decoders as specified in Table I.⁶

TABLE I
REGULARITY CONDITIONS OF ENCODER/DECODERS AND
CORRESPONDING MINIMUM ϵ -ACHIEVABLE RATES.

Encoder	Decoder	Minimum ϵ -achievable rate
Linear	Continuous	$R^*(X, \epsilon)$
Continuous	Lipschitz	$R(X, \epsilon)$
Linear	Lipschitz	$\hat{R}(X, \epsilon)$

Remark 2. In Definition 5, R and \hat{R} are defined under the Lipschitz continuity assumption of the decoder, which does not preclude the case where the Lipschitz constants blow up as the dimension grows. For practical applications, decoders with bounded Lipschitz constants are desirable, which amounts to constructing a sequence of decoders with Lipschitz constant only depending on the rate and the input statistics. As we will show later, this is indeed possible for discrete-continuous mixtures.

B. Results

The following general result holds for any input process [12]:

Theorem 3. *For any X and any $0 \leq \epsilon \leq 1$,*

$$R^*(X, \epsilon) \leq R(X, \epsilon) \leq \hat{R}(X, \epsilon). \quad (23)$$

Moreover, (23) holds for arbitrary input processes that are not necessarily i.i.d..

The second inequality in (23) follows from the definitions, since

$$\hat{R}(X, \epsilon) \geq \max\{R^*(X, \epsilon), R(X, \epsilon)\}.$$

Far less intuitive is the first inequality, proved in [12, Section VI], which states that robust reconstruction is always harder to achieve than linear compression.

The following result is a finite-dimensional version of the general achievability result of linear encoding in [12, Theorem 18], which states that sets of low Minkowski dimension can

⁶It was shown in [12] that in the definition of R^* and R , the continuity constraint can be replaced by Borel measurability without changing the minimum rate.

be linearly embedded into low-dimensional Euclidean space probabilistically. This is a probabilistic generalization of the embeddability result in [37].

Theorem 4. Let X^n be a random vector with $\overline{\dim}_B^\epsilon(P_{X^n}) \leq k$. Let $m > k$. Then for Lebesgue almost every $\mathbf{A} \in \mathbb{R}^{m \times n}$, there exists a $(1 - \frac{k}{m})$ -Hölder continuous function $g : \mathbb{R}^m \rightarrow \mathbb{R}^n$, i.e., $\|g(x) - g(y)\| \leq L\|x - y\|^{1 - \frac{k}{m}}$ for some $L > 0$ and all x, y , such that $\mathbb{P}\{g(\mathbf{A}X^n) \neq X^n\} \leq \epsilon$.

Remark 3. In Theorem 4, the decoder can be chosen as follows: by definition of $\overline{\dim}_B^\epsilon$, there exists $U \subset \mathbb{R}^n$, such that $\overline{\dim}_B(U) \leq k$. 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 $g(y^k) = 0$.

Generalizing [12, Theorem 9], a non-asymptotic converse for Lipschitz decoding is the following:

Theorem 5. For any random vector X^n , if there exists a Borel measurable $f : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and a Lipschitz continuous $g : \mathbb{R}^k \rightarrow \mathbb{R}^n$ such that $\mathbb{P}\{g(f(X^n)) \neq X^n\} \leq \epsilon$, then

$$k \geq \overline{\dim}_B^\epsilon(P_{X^n}) \geq \bar{d}(X^n) - \epsilon n. \quad (24)$$

Proof: Section VI-B. ■

Remark 4. An immediate consequence of (24) is that for general input processes, we have

$$R(X, \epsilon) \geq \limsup_{n \rightarrow \infty} \frac{\bar{d}(X^n)}{n} - \epsilon. \quad (25)$$

which, for i.i.d. inputs, becomes

$$R(X, \epsilon) \geq \bar{d}(X) - \epsilon. \quad (26)$$

In fact, combining the left inequality in (24) and the following concentration-of-measure result [12, Theorem 14]: for any $0 < \epsilon < 1$,

$$\liminf_{n \rightarrow \infty} \frac{\overline{\dim}_B^\epsilon(P_{X^n})}{n} \geq \bar{d}(X), \quad (27)$$

(26) can be superseded by the following *strong converse*:

$$R(X, \epsilon) \geq \bar{d}(X). \quad (28)$$

General achievability results for $R(X, \epsilon)$ rely on rectifiability results from geometric measure theory [38]. See [12, Section VII].

For discrete-continuous mixtures, we show that linear encoders and Lipschitz decoders can be realized *simultaneously* with *bounded* Lipschitz constants.

Theorem 6 (Linear encoding: discrete-continuous mixture). Let P_X be a discrete-continuous mixed distribution of the form (3), with the weight of the continuous part equal to γ . Then

$$R^*(X, \epsilon) = d(X) = \gamma \quad (29)$$

for all $0 < \epsilon < 1$. Moreover, if the discrete part P_d 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 a Lipschitz-constant independent of n :

$$L = \frac{\sqrt{R}}{R - \gamma} \left(\frac{R}{\gamma} \right)^{\frac{\gamma}{R - \gamma}} \exp \left(\frac{H(P_d)(1 - \gamma) + h(\gamma)}{R - \gamma} + \frac{1}{2} \right). \quad (30)$$

Consequently,

$$\hat{R}(X, \epsilon) = R^*(X, \epsilon) = d(X) = \gamma. \quad (31)$$

Proof: Section VI-B. ■

Combining Theorem 6 and [12, Theorem 10] yields the following tight result: for any i.i.d. input with a common distribution of the discrete-continuous mixture form in (3), whose discrete component has finite entropy, we have

$$R^*(X, \epsilon) = R(X, \epsilon) = \hat{R}(X, \epsilon) = \gamma \quad (32)$$

for all $0 < \epsilon < 1$. In the special case of sparse signals ($P_d = \delta_0$) with $s = \gamma n$ non-zeros, this implies that roughly s linear measurements are sufficient to recover the unknown vector with high probability. This agrees with the well-known result that $s + 1$ measurements are both necessary and sufficient to reconstruct an s -sparse vector probabilistically (see, e.g., [39]).

Remark 5. In the achievability proof of Theorem 6, our construction of a sequence of Lipschitz decoders with bounded Lipschitz constants independent of the dimension n only works for recovery performance measured in the ℓ_2 norm. The reason is two-fold: 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, Kirschbraun's theorem states that any Lipschitz mapping from a subset of a Hilbert space to a Hilbert space can be extended to the whole space with the same Lipschitz constant [40, Theorem 1.31, p. 21]. This result fails for general Banach spaces, in particular, for \mathbb{R}^n equipped with any ℓ_p -norm ($p \neq 2$) [40, p. 20]. Of course, by the equivalence of norms on finite-dimensional spaces, it is always possible to extend to a Lipschitz function with a larger Lipschitz constant; however, such soft analysis does not control the size of the Lipschitz constant, which may possibly blow up as the dimension increases. Nevertheless, (28) shows that even if we allow a sequence of decompressors with Lipschitz constants that diverges as $n \rightarrow \infty$, the compression rate is still lower bounded by $\bar{d}(X)$.

Remark 6 (Behavior of the Lipschitz constant). The Lipschitz constant of the decoder is a proxy to gauge the decoding robustness. It is interesting to investigate what is the smallest attainable Lipschitz constant as a for a given rate $R > \gamma$. Note that the constant in (30) depends exponentially on $\frac{1}{R - \gamma}$, which implies that the decoding becomes increasingly less robust as the rate approaches the fundamental limit. For sparse signals ($P_d = \delta_0$ hence $H(P_d) = 0$), (30) reduces to

$$L = \frac{\sqrt{eR}}{R - \gamma} \left(\frac{R}{\gamma^2} \right)^{\frac{\gamma}{R - \gamma}} \left(\frac{1}{1 - \gamma} \right)^{\frac{1 - \gamma}{R - \gamma}}. \quad (33)$$

It is unclear whether it is possible to achieve a Lipschitz constant that diverges polynomially as $R \rightarrow \gamma$.

Remark 7. Although too computationally intensive and numerically unstable (in fact discontinuous in general), in the conventional compressed sensing setup, the optimal decoder is an ℓ_0 -minimizer that seeks that sparsest solution compatible with the linear measurements. In our Bayesian setting, such a decoder does not necessarily minimize the probability of selecting the wrong signal. However, the ℓ_0 -minimization decoder does achieve the asymptotic fundamental limit $R^*(X, \epsilon)$ for any sparse $P_X = (1 - \gamma)\delta_0 + \gamma P_c$, since it is, in fact, even better than the asymptotically optimum decoder described in Remark 3. The optimality of the ℓ_0 -minimization decoder for sparse signals has also been observed in [17, Section IV-A1] based on replica heuristics.

Remark 8. Converse results for any linear encoder and decoder pair have been proposed before in other compressed sensing setups. For example, the result in [41, Theorem 3.1] assumes noiseless measurement with arbitrary sensing matrices and recovery algorithms, dealing with best sparse approximation under ℓ_1/ℓ_1 -stability guarantee. The following non-asymptotic lower bound on the number of measurements is shown: if there exist a sensing matrix $\mathbf{A} \in \mathbb{R}^{k \times n}$, a decoder $g: \mathbb{R}^k \rightarrow B_0^n(s)$ ⁷ and a constant $C > 0$, such that

$$\|x - g(\mathbf{A}x)\|_1 \leq \min_{z \in B_0^n(s)} C\|x - z\|_1 \quad (34)$$

for any $x \in \mathbb{R}^n$, then $k \geq \frac{s \log \frac{n}{2s}}{\log(4+2C)}$. However, this result does not directly apply to our setup because we are dealing with ℓ_2/ℓ_2 -stability guarantee with respect to the measurement noise, instead of the sparse approximation error of the input vector.

IV. NOISY COMPRESSED SENSING

A. Setup

The basic setup of noisy compressed sensing is a joint source-channel coding problem as shown in Fig. 2, where we




Fig. 2. Noisy compressed sensing setup.

assume that

- The source X^n consists of i.i.d. copies of a real-valued random variable X with unit variance.
- The channel is stationary memoryless with i.i.d. additive Gaussian noise σN^k where $N^k \sim \mathcal{N}(0, \mathbf{I}_k)$.
- Unit average power constraint on the encoded signal:

$$\frac{1}{k} \mathbb{E}[\|f_n(X^n)\|_2^2] \leq 1. \quad (35)$$

- The reconstruction error is gauged by the per-symbol MSE distortion:

$$d(x^n, \hat{x}^n) = \frac{1}{n} \|\hat{x}^n - x^n\|_2^2. \quad (36)$$

⁷ $B_0^n(s) = \{x \in \mathbb{R}^n : \|x\|_0 \leq s\}$ denotes the collection of all s -sparse n -dimensional vectors.

In this setup, the fundamental question is: For a given noise variance and measurement rate, what is the lowest 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

$$\inf_g \mathbb{E} [\|X^n - g(\hat{Y}^k)\|^2] = \text{mmse}(X^n | f(X^n) + \sigma N^k). \quad (37)$$

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 rates 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.

B. Distortion-rate tradeoff

For a fixed noise variance σ^2 , we define three distortion-rate functions that correspond to *optimal* encoding, *optimal linear* encoding and *random linear* encoding, respectively. In the remainder of this section, we fix $k = \lfloor Rn \rfloor$.

1) Optimal encoder:

Definition 6. The minimal distortion achieved by the optimal encoding scheme is given by:

$$D^*(X, R, \sigma^2) \triangleq \limsup_{n \rightarrow \infty} \frac{1}{n} \inf_f \{ \text{mmse}(X^n | f(X^n) + \sigma N^k) : \mathbb{E}[\|f(X^n)\|_2^2] \leq Rn \}. \quad (38)$$

For stationary ergodic sources, the asymptotic optimization problem in (38) can be solved by applying Shannon's joint source-channel coding separation theorem [42, Section XI], which states that the lowest rate, R , that achieves distortion D is given by

$$R = \frac{R_X(D)}{C(\sigma^2)}, \quad (39)$$

where $R_X(\cdot)$ is the rate-distortion function of X in (9) 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

$$D^*(X, R, \sigma^2) = R_X^{-1} \left(\frac{R}{2} \log(1 + \sigma^{-2}) \right). \quad (40)$$

In general, optimal joint source-channel encoders are nonlinear [43]. 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 $n \rightarrow \infty$. Such a construction results in an encoder that is finite-valued, hence nonlinear. For fixed n and k , linear encoders are in general suboptimal.

2) *Optimal linear encoder*: To analyze the fundamental limit of conventional 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 (35) simplifies to

$$\mathbb{E}[\|\mathbf{H}X^n\|_2^2] = \mathbb{E}[X^{nT}\mathbf{H}^T\mathbf{H}X^n] = \text{Tr}(\mathbf{H}^T\mathbf{H}) = \|\mathbf{H}\|_F^2 \leq k, \quad (41)$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

Definition 7. Define the optimal distortion achievable by linear encoders as:

$$D_L^*(X, R, \sigma^2) \triangleq \limsup_{n \rightarrow \infty} \frac{1}{n} \inf_{\mathbf{H}} \left\{ \text{mmse}(X^n | \mathbf{H}X^n + \sigma N^k) : \|\mathbf{H}\|_F^2 \leq k \right\}. \quad (42)$$

3) *Random linear encoder*: We consider the ensemble performance of random linear encoders and relax the power constraint in (41) to hold on average:

$$\mathbb{E}[\|\mathbf{A}\|_F^2] \leq k. \quad (43)$$

In particular, we focus on the following ensemble of random sensing matrices, for which (43) holds with equality:

Definition 8. Let \mathbf{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:

$$\begin{aligned} D_L(X, R, \sigma^2) &\triangleq \limsup_{n \rightarrow \infty} \frac{1}{n} \text{mmse}(X^n | (\mathbf{A}_n X^n + \sigma N^k, \mathbf{A}_n)) \\ &= \limsup_{n \rightarrow \infty} \text{mmse}(X_1 | (\mathbf{A}_n X^n + \sigma N^k, \mathbf{A}_n)) \end{aligned} \quad (44)$$

(45)

where (45) follows from symmetry and $\text{mmse}(\cdot|\cdot)$ is defined in (15).⁸

General formulae for $D_L(X, R, \sigma^2)$ and $D_L^*(X, R, \sigma^2)$ are yet unknown. One example where they can be explicitly computed is given in Section IV-D – the Gaussian source.

4) Properties:

Theorem 7. 1) For fixed σ^2 , $D^*(X, R, \sigma^2)$ and $D_L^*(X, R, \sigma^2)$ are both decreasing, convex and continuous in R on $(0, \infty)$.

2) For fixed R , $D^*(X, R, \sigma^2)$ and $D_L^*(X, R, \sigma^2)$ are both decreasing, convex and continuous in $\frac{1}{\sigma^2}$ on $(0, \infty)$.

3)

$$D^*(X, R, \sigma^2) \leq D_L^*(X, R, \sigma^2) \leq D_L(X, R, \sigma^2) \leq 1. \quad (46)$$

Proof:

1) Fix σ^2 . Monotonicity with respect to the measurement rate R is straightforward from the definition of D^* and

⁸The MMSE on the right-hand side of (44) and (45) can be computed by first fixing the sensing matrix \mathbf{A}_n then averaging with respect to its distribution.

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 $f: \mathbb{R}^n \rightarrow \mathbb{R}^k$, $\sigma^2 \mapsto \text{mmse}(X^n | f(X^n) + \sigma N^k)$ is increasing. This is a consequence of the infinite divisibility of the Gaussian distribution as well as the data processing lemma of MMSE [44]. Consequently, $\sigma^2 \mapsto D^*(X, R, \sigma^2)$ is also increasing. Since D^* can be equivalently defined as

$$D^*(X, R, \sigma^2) = \limsup_{n \rightarrow \infty} \frac{1}{n} \inf_f \left\{ \text{mmse}(X^n | f(X^n) + N^k) : \mathbb{E}[\|f(X^n)\|_2^2] \leq \frac{k}{\sigma^2} \right\}, \quad (47)$$

convexity in $\frac{1}{\sigma^2}$ follows from time-sharing. The results on D_L^* follows analogously.

3) The leftmost inequality in (46) 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 (46) is not quite trivial because the power constraint in (35) is not imposed on each matrix in the ensemble. The proof of this inequality can be found in Appendix A. ■

Remark 9. Alternatively, the convexity properties of D^* can be derived from (40). 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 $\sigma^{-2} \mapsto D^*(X, R, \sigma^2)$ [45, p. 84]. The convexity of $R \mapsto D^*(X, R, \sigma^2)$ can be similarly proved.

Remark 10. Note that the time-sharing proofs of Theorem 7.1 and 7.2 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 random matrices with i.i.d. entries considered in Definition 8. Therefore, proving the convexity of $R \mapsto D_L(X, R, \sigma^2)$ amounts to showing that 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.

C. Phase transition of noise sensitivity

One of the main objectives of noisy compressed sensing is to achieve robust reconstruction, obtaining a reconstruction error that is proportional to the noise variance. To quantify robustness, we analyze *noise sensitivity*, namely 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 9. The worst-case noise sensitivity of optimal encoding is defined as

$$\zeta^*(X, R) = \sup_{\sigma^2 > 0} \frac{D^*(X, R, \sigma^2)}{\sigma^2}. \quad (48)$$

For linear encoding, $\zeta_L^*(X, R)$ and $\zeta_L(X, R)$ are analogously defined with D^* in (48) replaced by D_L^* and D_L , respectively.

Remark 11. In the analysis of LASSO and the AMP algorithms [13], the noise sensitivity is defined in a minimax fashion where a further supremum is taken over all input distributions that have an atom at zero of mass at least $1 - \epsilon$. In contrast, the sensitivity in Definition 9 is a Bayesian quantity where we fix the input distribution. Similar notion of sensitivity has been defined in [16, Equation (49)].

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 [13], [1]:

Definition 10. Define

$$\mathcal{R}^*(X) \triangleq \inf \{R > 0: \zeta^*(X, R) < \infty\}. \quad (49)$$

For linear encoding, $\mathcal{R}_L^*(X)$ and $\mathcal{R}_L(X)$ are analogously defined with ζ^* in (49) replaced by ζ_L^* and ζ_L .

By (46), the phase transition thresholds in Definition 10 are ordered naturally as

$$0 \leq \mathcal{R}^*(X) \leq \mathcal{R}_L^*(X) \leq \mathcal{R}_L(X) \leq 1, \quad (50)$$

where the rightmost inequality is shown below (after Theorem 8).

Remark 12. In view of the convexity properties in Theorem 7.2, the three worst-case sensitivities in Definition 9 are all (extended real-valued) convex functions of R .

Remark 13. Alternatively, we can consider the *asymptotic noise sensitivity* by replacing the supremum in (48) with the limit as $\sigma^2 \rightarrow 0$, denoted by ξ^*, ξ_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 $\text{var} X = 1$, we have

$$\zeta^*(X, R) < \infty \Leftrightarrow \xi^*(X, R) < \infty. \quad (51)$$

Therefore $\mathcal{R}^*(X)$ can be equivalently defined as the infimum of all rates $R > 0$, such that

$$D^*(X, R, \sigma^2) = O(\sigma^2), \quad \sigma^2 \rightarrow 0. \quad (52)$$

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 (51) need not be achieved as $\sigma^2 \rightarrow 0$. An example is given by the Gaussian input analyzed in Section IV-D.

D. Least-favorable input: Gaussian distribution

In this section we compute the distortion-rate tradeoffs for the Gaussian input distribution. Although Gaussian input distribution is not directly relevant for compressed sensing due to its lack of sparsity, it is still interesting to investigate the distortion-rate tradeoff in the Gaussian case for the following reasons:

- 1) As the least-favorable input distribution, Gaussian distribution simultaneously maximizes all three distortion-rate

functions subject to the variance constraint and provides upper bounds for non-Gaussian inputs.

- 2) Connections are made to classical joint-source-channel coding problems in information theory about transmitting Gaussian sources over Gaussian channels and (sub)optimality of linear coding (e.g., [46], [47], [48]).
- 3) It serves as a concrete illustration of the phenomenon of coincidence of all thresholds defined in Definitions 6 – 8, which are fully generalized in Section IV-E to the mixture model.

Theorem 8. Let $X_G \sim \mathcal{N}(0, 1)$. Then for any R, σ^2 and X of unit variance,

$$D^*(X, R, \sigma^2) \leq D^*(X_G, R, \sigma^2) = \frac{1}{(1 + \sigma^{-2})^R} \quad (53)$$

$$D_L^*(X, R, \sigma^2) \leq D_L^*(X_G, R, \sigma^2) = 1 - \frac{R}{\sigma^2 + \max\{1, R\}} \quad (54)$$

$$D_L(X, R, \sigma^2) \leq D_L(X_G, R, \sigma^2) \quad (55)$$

$$= \frac{1}{2} \left(1 - R - \sigma^2 + \sqrt{(1 - R)^2 + 2(1 + R)\sigma^2 + \sigma^4} \right) \quad (56)$$

Proof: Since the Gaussian distribution maximizes the rate-distortion function pointwise under the variance constraint [49, Theorem 4.3.3], the inequality in (53) follows from (40). 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 (54) and (56) follow. The distortion-rate functions of X_G are computed in Appendix B. ■

The Gaussian distortion-rate tradeoffs in (53) – (56) are plotted in Figs. 3 and 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.,

$$D^*(X_G, 1, \sigma^2) = D_L^*(X_G, 1, \sigma^2), \quad (57)$$

which is a well-known fact [46], [47]. As a result of (56), the rightmost inequality in (50) follows.




Fig. 3. $D^*(X_G, R, \sigma^2)$, $D_L^*(X_G, R, \sigma^2)$, $D_L(X_G, R, \sigma^2)$ against $\text{snr} = \sigma^{-2}$.

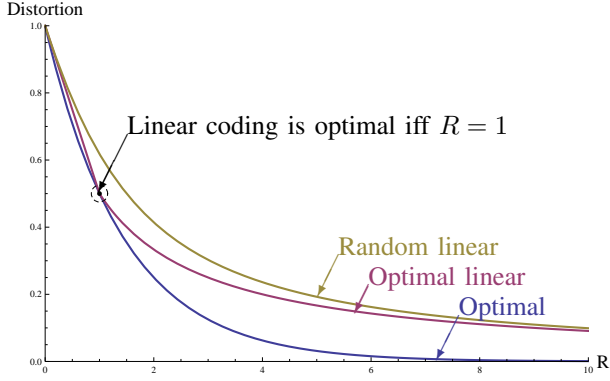


Fig. 4. $D^*(X_G, R, \sigma^2)$, $D_L^*(X_G, R, \sigma^2)$, $D_L(X_G, R, \sigma^2)$ against R when $\sigma^2 = 1$.

Next, using straightforward limits, we analyze the high-SNR asymptotics of (53) – (56). The smallest among the three, $D^*(X_G, R, \sigma^2)$ vanishes polynomially in σ^2 according to

$$D^*(X_G, R, \sigma^2) = \sigma^{2R} + O(\sigma^{2R+2}), \quad \sigma^2 \rightarrow 0 \quad (58)$$

regardless of how small $R > 0$ is. For linear encoding, we have

$$D_L^*(X, R, \sigma^2) = \begin{cases} 1 - R + R\sigma^2 + O(\sigma^4) & 0 \leq R < 1, \\ \sigma^2 + O(\sigma^4) & R = 1, \\ \frac{\sigma^2}{R} + O(\sigma^4) & R > 1. \end{cases} \quad (59)$$

$$D_L(X_G, R, \sigma^2) = \begin{cases} 1 - R + \frac{R}{1-R}\sigma^2 + O(\sigma^4) & 0 \leq R < 1, \\ \sigma - \frac{\sigma^2}{2} + O(\sigma^3) & R = 1, \\ \frac{\sigma^2}{R-1} + O(\sigma^4) & R > 1. \end{cases} \quad (60)$$

The weak-noise behavior of D_L^* and D_L are compared in different regimes of measurement rates:

- $0 \leq R < 1$: 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.
- $R = 1$: $D_L = \sigma(1 + o(1))$ is much worse than $D_L^* = \sigma^2(1 + o(1))$, which is achieved by choosing the encoding matrix to be identity. In fact, with nonnegligible probability, the optimal estimator that attains (56) 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. 3 ($R = 0.3$ and 5) and Fig. 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 worst-case noise sensitivity is concerned.

From (58) and (60), we observe that both D_L^* and D_L exhibit a sharp phase transition near the critical rate $R = 1$:

$$\lim_{\sigma^2 \rightarrow 0} D_L^*(X, R, \sigma^2) = \lim_{\sigma^2 \rightarrow 0} D_L(X, R, \sigma^2) \quad (61)$$

$$= (1 - R)^+. \quad (62)$$

where $x^+ \triangleq \max\{0, x\}$. Moreover, from (53) – (56) we obtain the worst-case and asymptotic noise sensitivity functions for the Gaussian input as follows:

$$\zeta^*(X_G, R) = \begin{cases} \exp(-Rh(\frac{1}{R})) & R \geq 1 \\ \infty & R < 1 \end{cases}, \quad (63)$$

$$\xi^*(X_G, R) = \begin{cases} 0 & R > 1 \\ 1 & R = 1 \\ \infty & R < 1 \end{cases} \quad (64)$$

and

$$\zeta_L^*(X_G, R) = \xi_L^*(X_G, R) = \begin{cases} \frac{1}{R} & R \geq 1 \\ \infty & R < 1 \end{cases}, \quad (65)$$

$$\zeta_L(X_G, R) = \xi_L(X_G, R) = \begin{cases} \frac{1}{R-1} & R > 1 \\ \infty & R \leq 1 \end{cases} \quad (66)$$

The worst-case noise sensitivity functions are plotted in Fig. 5 against the measurement rate R . Note that (64) provides an example for Remark 13: 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.




Fig. 5. Worst-case noise sensitivity ζ^* , ζ_L^* and ζ_L for the Gaussian input, which all become infinity when $R < 1$ (the unstable regime).

In view of (64) – (66), the phase-transition thresholds in the Gaussian signal case are:

$$\mathcal{R}^*(X_G) = \mathcal{R}_L^*(X_G) = \mathcal{R}_L(X_G) = 1. \quad (67)$$

The equality of the three phase-transition thresholds turns out to hold well beyond the Gaussian signal model. In the next subsection, we formulate and prove the existence of the phase thresholds for all three distortion-rate functions and discrete-continuous mixtures, which turn out to be equal to the information dimension of the input distribution.

E. Non-Gaussian inputs

This subsection 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 IV-D carries over well beyond the Gaussian case. Proofs are deferred to Section VI-C.

The phase transition threshold for *optimal encoding* is given by the upper *information dimension* of the input:

Theorem 9. *For any X that satisfies (5),*

$$\mathcal{R}^*(X) = \bar{d}(X) \quad (68)$$

Moreover, if P_X is a discrete-continuous mixture as in (3), then for any $R \geq \gamma$, as $\sigma \rightarrow 0$,

$$D^*(X, R, \sigma^2) = \frac{\exp\left(2H(P_d)\frac{1-\gamma}{\gamma} - 2\mathcal{D}(P_c)\right)}{(1-\gamma)^{\frac{2(1-\gamma)}{\gamma}}\gamma} \sigma^{\frac{2R}{\gamma}} (1+o(1)) \quad (69)$$

where $\mathcal{D}(\cdot)$ denotes the non-Gaussianity of a probability measure, defined as its relative entropy with respect to a Gaussian distribution with the same mean and variance. Consequently, the asymptotic noise sensitivity of optimal encoding is

$$\xi^*(X, R) = \begin{cases} \infty & R < \gamma \\ \frac{\exp\left(2H(P_d)\frac{1-\gamma}{\gamma} - 2\mathcal{D}(P_c)\right)}{(1-\gamma)^{\frac{2(1-\gamma)}{\gamma}}\gamma} & R = \gamma \\ 0 & R > \gamma. \end{cases} \quad (70)$$

The next result shows that random linear encoders with i.i.d. Gaussian coefficients also achieve information dimension for any discrete-continuous mixtures, which, in view of Theorem 9, implies that, at least asymptotically, (random) linear encoding suffices for robust reconstruction as long as the input distribution contains no singular component.

Theorem 10. *Assume that X has a discrete-continuous mixed distribution as in (3), where the discrete component P_d has finite entropy. Then*

$$\mathcal{R}^*(X) = \mathcal{R}_L^*(X) = \mathcal{R}_L(X) = \gamma. \quad (71)$$

Moreover,

- 1) (71) holds for any non-Gaussian noise distribution with finite non-Gaussianity.
- 2) For any $R > \gamma$, the worst-case noise sensitivity of Gaussian sensing matrices is upper bounded by

$$\xi_L(X, R) \leq \frac{R^2}{(R-\gamma)^2} \left(\frac{R}{\gamma}\right)^{\frac{2\gamma}{R-\gamma}} \times \exp\left(\frac{2H(P_d)(1-\gamma) + 2h(\gamma)}{R-\gamma} + 1\right). \quad (72)$$

Remark 14. The achievability proof of $\mathcal{R}_L(X)$ is a direct application of Theorem 6, 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 $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Then

$$\begin{aligned} & \mathbb{E} \left[\|g_n(\mathbf{A}_n X^n + \sigma N^k) - X^n\|^2 \right] \\ & \leq L^2(R) \sigma^2 \mathbb{E} [\|N^k\|^2] + \epsilon_n = kL^2(R) \sigma^2 \text{var} N + \epsilon_n, \end{aligned} \quad (73)$$

which implies that robust reconstruction is achievable at rate R and the worst-case noise sensitivity is upper bounded by $L_R^2 R$.

Notice that the above 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}_L^*) to hold, the non-Gaussian noise needs to have finite non-Gaussianity. The disadvantage of this approach is that currently it lacks an explicit construction because the extendability of Lipschitz functions (Kirschbraun's theorem) is only an existence result which relies on the Hausdorff maximal principle [40, Theorem 1.31, p. 21], which is equivalent to the axiom of choice. On Euclidean spaces it is possible to obtain an explicit construction by applying the results in [50], [51] to a countable dense subset of the domain. However, such a construction is far from being practical.

Remark 15. We emphasize the following “universality” aspects of Theorem 10:

- Gaussian random sensing matrices achieve the optimal transition threshold for any discrete-continuous mixture, as long as it is known at the decoder;
- 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 (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, or even its variance;
- The conclusion holds for non-Gaussian noise as long as it has finite non-Gaussianity.

F. Results relying on replica heuristics

Based on the statistical-physics approach in [52], [53], the decoupling principle results in [53] were imported into the compressed sensing setting in [14] to *postulate* the following formula for $D_L(X, R, \sigma^2)$. Note that this result is based on replica heuristics currently lacking a rigorous justification.

Replica Symmetry Postulate ([14, Corollary 1, p.5]).

$$D_L(X, R, \sigma^2) = \text{mmse}(X, \eta R \sigma^{-2}), \quad (74)$$

where $0 < \eta < 1$ satisfies the following equation [14, (12) – (13), pp. 4 – 5]:⁹

$$\frac{1}{\eta} = 1 + \frac{1}{\sigma^2} \text{mmse}(X, \eta R \sigma^{-2}). \quad (75)$$

When (75) has more than one solution, η is chosen to minimize the free energy

$$I(X; \sqrt{\eta R \sigma^{-2}} X + N) + \frac{R}{2}(\eta - 1 - \log \eta). \quad (76)$$

In view of the the I-MMSE relationship [54], the solutions to (75) are precisely the stationary points of the free energy (76) as a function of η . In fact it is possible for (74) to have arbitrarily many solutions. For an explicit example, see Remark 21 in Section VI-C.

Note that the solution in (74) does *not* depend on the distribution of the random measurement matrix \mathbf{A} , as long as its entries are i.i.d. with zero mean and variance $\frac{1}{n}$. Therefore it is possible to employ a random sparse measurement matrix so that each encoding operation involves only a relatively few signal components, for example,

$$A_{ij} \sim \frac{p}{2} \delta_{\frac{-1}{\sqrt{pn}}} + (1-p) \delta_0 + \frac{p}{2} \delta_{\frac{1}{\sqrt{pn}}} \quad (77)$$

for some $0 < p < 1$. In fact, in the special case of $p = \frac{\log n}{n}$, the replica symmetry postulate can be rigorously proved [14, Sec. IV] (see also [55], [56]).

Assuming the validity of the replica symmetry postulate, 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 (18) plays a key role in analyzing the minimizer of the free energy (76).¹⁰

Theorem 11. Assume that the replica symmetry postulate holds for X . Then for any i.i.d. random measurement matrix \mathbf{A} whose entries have zero mean and variance $\frac{1}{n}$,

$$\underline{\mathcal{D}}(X) \leq \mathcal{R}_L(X) \leq \overline{\mathcal{D}}(X). \quad (78)$$

Therefore if $\mathcal{D}(X)$ exists, we have

$$\mathcal{R}_L(X) = \mathcal{D}(X) = d(X), \quad (79)$$

and in addition,

$$D_L(X, R, \sigma^2) = \frac{d(X)}{R - d(X)} \sigma^2 (1 + o(1)). \quad (80)$$

The general result in Theorem 11 holds for any input distribution but relies on the conjectured validity of the replica symmetry postulate. For the special case of discrete-continuous mixtures in (3), in view of Theorem 2, Theorem 11 predicts (with the caveat of the validity of the replica symmetry

postulate) that the phase-transition threshold for Gaussian sensing matrices is γ , which agrees with the rigorously proven result in Theorem 10. Therefore, the only added benefit of Theorem 11 is to allow singular components in the input distribution.

Remark 16. In statistical physics, the phase transition near the threshold often behaves according to a power law with certain universal exponent, known as the *critical exponent* [57, Chapter 3]. According to (80), as the measurement rate R approaches the fundamental limit $d(X)$, the replica method suggests that the optimal noise sensitivity blows up according as the power law $\frac{1}{R - d(X)}$, where the unit exponent holds universally for all mixture distributions. It remains an open question whether this power law behavior can be rigorously proven and whether the optimal exponent is one. Note that by using the Lipschitz extension scheme in the proof Theorem 10, we can achieve the noise sensitivity in (72), which blows up exponentially as the $R - d(X)$ vanishes and is likely to be highly suboptimal.

Remark 17. In fact, the proof of Theorem 11 shows that the converse part (left inequality) of (78) holds in a much stronger sense: as long as there is no residual error in the weak-noise limit, that is, if $D_L(X, R, \sigma^2) = o(1)$ as $\sigma^2 \rightarrow 0$, then $R \geq \underline{\mathcal{D}}(X)$ must hold. Therefore, the converse part of Theorem 11 still holds even if we weaken the right-hand side of (52) from $O(\sigma^2)$ to $o(1)$.

Remark 18. Assume the validity of the replica symmetry postulate. Combining Theorem 9, Theorem 11 and (50) gives an *operational* proof for $\underline{\mathcal{D}}(X) \leq \overline{\mathcal{D}}(X)$, the fourth inequality in (18), which has been proven analytically in [29, Theorem 8].

V. COMPARISONS TO LASSO AND AMP ALGORITHMS

Widely popular in the compressed sensing literature, the LASSO [7], [6] and the approximate message passing (AMP) algorithms [9] are low-complexity reconstruction procedures, which are originally obtained as solutions to the conventional minimax setup in compressed sensing. In this section, we compare the phase transition thresholds of LASSO and AMP achieved in the Bayesian setting to the optimal thresholds derived in Sections III – IV. Similar Bayesian analysis has been performed in [25], [58], [9], [59], [60].

A. Signal models

The following three families of input distributions are considered [9, p. 18915], indexed by $\chi = \pm, +$ and \square respectively, which all belong to the family of input distributions of the mixture form in (3):

- \pm : sparse signals (2);
- $+$: sparse non-negative signals (2) with the continuous component P_c supported on \mathbb{R}_+ .
- \square : simple signals (Section I-C) [25, Section 5.2, p. 540]

$$P = (1 - \gamma) \left(\frac{1}{2} \delta_0 + \frac{1}{2} \delta_1 \right) + \gamma P_c \quad (81)$$

⁹In the notation of [14, (12)], γ and $\epsilon\mu$ correspond to $R\sigma^{-2}$ and R in our formulation.

¹⁰It can be shown that in the limit of $\sigma^2 \rightarrow 0$, the minimizer of (76) when $R > \overline{\mathcal{D}}(X)$ and $R < \underline{\mathcal{D}}(X)$ corresponds to the largest and the smallest root of the fixed-point equation (74) respectively.

where P_c is some absolutely continuous distribution supported on the unit interval.

B. Noiseless measurements

In the noiseless case, we consider linear programming (LP) decoders and the AMP decoder [9] and the phase transition threshold of error probability. Phase transitions of greedy reconstruction algorithms have been analyzed in [61], which derived upper bounds (achievability results) for the transition threshold of measurement rate. 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 [27, Section I]):

$$g_{\pm}(y) = \arg \min \{\|x\|_1 : x \in \mathbb{R}^n, \mathbf{A}x = y\}, \quad (82)$$

$$g_+(y) = \arg \min \{\|x\|_1 : x \in \mathbb{R}_+^n, \mathbf{A}x = y\}, \quad (83)$$

$$g_{\square}(y) = \{x : x \in [0, 1]^n, \mathbf{A}x = y\}. \quad (84)$$

For sparse signals, (82) – (83) are based on ℓ_1 -minimization (also known as Basis Pursuit [6], which is the noiseless limit of LASSO defined in Section V-C), while for simple signals, the decoder (84) solves an LP feasibility problem. In general the decoders in (82) – (84) 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 $\mathbb{P}\{g_{\chi}(\mathbf{A}X^n) \neq \{X^n\}\}$, evaluated with respect to the product measure $(P_X)^n \times P_{\mathbf{A}}$.

The phase transition thresholds of the reconstruction error probability for decoders (82) – (84) are derived in [11] using combinatorial geometry. For sparse signals and ℓ_1 -minimization decoders (82) – (83), the expressions of the corresponding thresholds $R_{\pm}(\gamma)$ and $R_+(\gamma)$ are quite involved, given implicitly in [11, Definition 2.3]. As observed in [9, Finding 1], $R_{\pm}(\gamma)$ and $R_+(\gamma)$ agree numerically with the following expressions:¹¹

$$R_{\pm}(\gamma) = \min_{\alpha \geq 0} \gamma(1 + \alpha^2) + 2(1 - \gamma)((1 + \alpha^2)\Phi(-\alpha) - \alpha\varphi(\alpha)) \quad (85)$$

$$R_+(\gamma) = \min_{\alpha \geq 0} \gamma(1 + \alpha^2) + (1 - \gamma)((1 + \alpha^2)\Phi(-\alpha) - \alpha\varphi(\alpha)) \quad (86)$$

which is now rigorously established in view of the results in [63]. For simple signals, the phase transition threshold is proved to be [25, Theorem 1.1]

$$R_{\square}(\gamma) = \frac{\gamma + 1}{2}. \quad (87)$$

Moreover, substantial numerical evidence in [9] suggests that the phase transition thresholds for the AMP decoder coincide with the LP thresholds for all three input distributions. The suboptimal thresholds obtained from (85) – (87) are plotted in

¹¹In the series of papers [25], [27], [9], [13], the phase diagrams 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 paper, the parameterization (γ, R) is used instead. The ratio $\frac{\gamma}{R\chi(\gamma)}$ is denoted by $\rho(\gamma; \chi)$ in [9]. The same parametrization is also used in [62].

Fig. 6 along with the optimal threshold obtained from Theorem 6 which is γ .¹² 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 and AMP decoders are severely suboptimal unless γ is close to one.




Fig. 6. Suboptimal thresholds (85) – (87) obtained with LASSO and AMP v.s. optimal threshold for the three signal models in Section V-A.

In the highly sparse regime which is most relevant to compressed sensing problems, it follows from [27, Theorem 3] that for sparse signals ($\chi = \pm$ or $+$),

$$R_{\chi}(\gamma) = 2\gamma \log_e \frac{1}{\gamma} (1 + o(1)), \text{ as } \gamma \rightarrow 0, \quad (88)$$

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 $2s \log_e \frac{n}{s}$ measurements to successfully recover the unknown vector, whose number of nonzero components is denoted by s . 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 $R_{\square}(\gamma)$ converges to $\frac{1}{2}$ instead of zero as $\gamma \rightarrow 0$. This suboptimality is due to the fact that the LP feasibility decoder (84) 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

$$\mathbf{T}(x^n) = \frac{1}{n} \sum_{i=1}^n (\mathbf{1}_{\{x_i=0\}}, \mathbf{1}_{\{x_i \notin \{0,1\}\}}, \mathbf{1}_{\{x_i=1\}}).$$

The decoder outputs the solution to $\mathbf{A}x^n = y^k$ such that $\mathbf{T}(x^n)$ is closest to $(\frac{1-\gamma}{2}, \gamma, \frac{1-\gamma}{2})$ (in total variation distance for example).

¹²A similar comparison between the suboptimal threshold $R_{\pm}(\gamma)$ and the optimal threshold γ has been provided in [17, Fig. 2(a)] based on a replica-heuristic calculation.

C. Noisy measurements

In the noisy case, we consider the AMP decoder [13] and the ℓ_1 -penalized least-squares (i.e. LASSO) decoder [7]:

$$\tilde{g}(y, \mathbf{A}; \lambda) = \underset{x \in \mathbb{R}^n}{\operatorname{argmin}} \frac{1}{2} \|y - \mathbf{A}x\|_2^2 + \lambda \|x\|_1, \quad (89)$$

where $\lambda > 0$ is a regularization parameter. Note that in the limit of $\lambda \rightarrow 0$, LASSO reduces to the ℓ_1 -minimization decoder defined in (82). For Gaussian sensing matrices and Gaussian observation noise, the asymptotic mean-square error achieved by LASSO for a fixed λ

$$D^{(\lambda)}(X, R, \sigma^2) \triangleq \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\|X^n - \tilde{g}(\mathbf{A}X^n + \sigma N^k; \lambda)\|^2 \right] \quad (90)$$

can be determined as a function of P_X, λ and σ by applying [59, Corollary 1.6].¹³ In Appendix C, we show that for any X distributed according to the mixture

$$P_X = (1 - \gamma)\delta_0 + \gamma Q, \quad (91)$$

where Q is an arbitrary probability measure with no mass at zero, the asymptotic noise sensitivity of LASSO with optimized λ is given by the following equation:

$$\tilde{\xi}(X, R) \triangleq \inf_{\lambda} \lim_{\sigma^2 \rightarrow 0} \frac{D^{(\lambda)}(X, R, \sigma^2)}{\sigma^2} \quad (92)$$

$$= \begin{cases} \frac{R_{\pm}(\gamma)}{R - R_{\pm}(\gamma)} & R > R_{\pm}(\gamma) \\ \infty & R \leq R_{\pm}(\gamma) \end{cases} \quad (93)$$

where $R_{\pm}(\gamma)$ is given in (86). By the same reasoning in Remark 13, the worst-case noise sensitivity of LASSO is finite if and only if $R > R_{\pm}(\gamma)$. Note that (93) does not depend on Q as long as $Q(\{0\}) = 0$. Therefore $R_{\pm}(\gamma)$ also coincides with the phase transition threshold in a minimax sense, obtained in [13, Proposition 3.1(1.a)] by considering the lease favorable Q . Analogously, the LASSO decoder (89) can be adapted to other signal structures (see for example [13, Sec. VI-A]), resulting in the phase-transition threshold $R_+(\gamma)$ and $R_-(\gamma)$ for sparse positive and simple signals, given by (86) and (87), respectively. Furthermore, these thresholds also apply to the AMP algorithm [64].

Next, focusing on sparse signals, we compare the performance of LASSO and AMP algorithms to the optimum. In view of (93), the phase transition thresholds of noise sensitivity for the LASSO and AMP decoder are both $R_{\pm}(\gamma)$ for any X distributed according to (91). We discuss the following two special cases:

- 1) Q is absolutely continuous, or alternatively, P_X is a discrete-continuous mixture given in (2). The optimal phase transition threshold is γ as a consequence of Theorem 10. Therefore the phase transition boundaries are identical to Fig. 6 and the same observation in Section V-B applies.

¹³It should be noted that in [13], [59], the entries of the sensing matrix is distributed according to $\mathcal{N}(0, \frac{1}{k})$ (column normalization). While in the present paper the sensing matrix has $\mathcal{N}(0, \frac{1}{n})$ entries (row normalization) in order for the encoded signal to have unit average power. Therefore the expression in (93) is equal to that in [13, Equation (1.9)] divided by the measurement R .

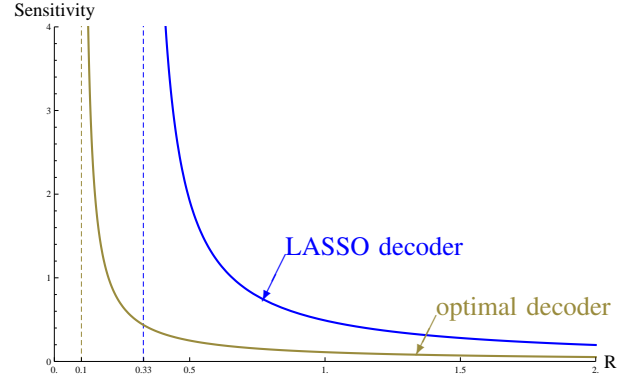


Fig. 7. Asymptotic noise sensitivity of the optimal decoder and the LASSO decoder exhibiting phase transitions: sparse signal model (2) with $\gamma = 0.1$.

- 2) Q is discrete with no mass at zero, e.g., $Q = \frac{1}{2}(\delta_1 + \delta_{-1})$. Since P_X is discrete with zero information dimension, the optimal phase transition threshold is equal to zero, while $R_{\pm}(\gamma)$ still applies to LASSO and AMP.

For sparse signals of the form (2) with $\gamma = 0.1$, Fig. 7 compares those expressions for the asymptotic noise sensitivity of LASSO (and AMP) algorithm to the optimal noise sensitivity predicted by Theorem 11 based on replica heuristics. Note that the phase transition threshold of LASSO is approximately 3.3 times the optimal.

VI. PROOFS

A. Auxiliary results

We need the following large-deviations result on Gaussian random matrices.

Lemma 1. Let $\sigma_{\min}(\mathbf{B}_k)$ denote the smallest singular value of the $k \times m_k$ matrix \mathbf{B}_k consisting of i.i.d. Gaussian entries with zero mean and variance $\frac{1}{k}$. For any $t > 0$, denote

$$F_{k, m_k}(t) \triangleq \mathbb{P} \{ \sigma_{\min}(\mathbf{B}_k) \leq t \}. \quad (94)$$

Suppose that $\frac{m_k}{k} \xrightarrow{k \rightarrow \infty} \alpha \in (0, 1)$. Then

$$\liminf_{k \rightarrow \infty} \frac{1}{k} \log \frac{1}{F_{k, m_k}(t)} \geq \frac{(1 - \alpha)^2}{2} \log \frac{(1 - \alpha)^2}{et^2} + \frac{\alpha}{2} \log \alpha. \quad (95)$$

Proof: For brevity let $\mathbf{H}_k = \sqrt{k}\mathbf{B}_k$ and suppress the dependence of m_k on k . Then $\mathbf{H}_k^T \mathbf{H}_k$ is an $m \times m$ Gaussian Wishart matrix. The minimum eigenvalue of $\mathbf{H}_k^T \mathbf{H}_k$ has a density, which admits the following upper bound [65, Proposition 5.1, p. 553].

$$f_{\lambda_{\min}(\mathbf{H}_k^T \mathbf{H}_k)}(x) \leq E_{k, m} x^{\frac{k-m-1}{2}} e^{-\frac{x}{2}}, \quad x \geq 0, \quad (96)$$

where

$$E_{k, m} \triangleq \frac{\sqrt{\pi} 2^{-\frac{k-m+1}{2}} \Gamma\left(\frac{k+1}{2}\right)}{\Gamma\left(\frac{m}{2}\right) \Gamma\left(\frac{k-m+1}{2}\right) \Gamma\left(\frac{k-m+2}{2}\right)}. \quad (97)$$

Then

$$\mathbb{P}\{\sigma_{\min}(\mathbf{B}_n) \leq t\} = \mathbb{P}\left\{\lambda_{\min}(\mathbf{H}_n^T \mathbf{H}_n) \leq kt^2\right\} \quad (98)$$

$$\leq E_{k,m} \int_0^{kt^2} x^{\frac{k-m-1}{2}} e^{-\frac{x}{2}} dx \quad (99)$$

$$\leq \frac{2k^{\frac{k-m+1}{2}} E_{k,m}}{k-m+1} t^{k-m+1}. \quad (100)$$

Applying Stirling's approximation to (100) yields (95). ■

Remark 19. More general non-asymptotic upper bound on $\mathbb{P}\{\sigma_{\min}(\mathbf{B}_k) \leq t\}$ is given in [66, Theorem 1.1], which holds universally for all sub-Gaussian distributions. Note that the upper bound in [66, Equation (1.10)] is of the form

$$\mathbb{P}\{\sigma_{\min}(\mathbf{B}_k) \leq t\} \leq (c_1 t)^{c_2 k} + \exp(-c_3 k) \quad (101)$$

where c_1, c_2, c_3 are constants. The second term in (101) is due to the fact that the least singular value for discrete ensembles (e.g. Rademacher) always has a mass at zero, which is exponentially small in k but independent of t . For Gaussian ensembles, however, we have $\sigma_{\min}(\mathbf{B}_k) > 0$ almost surely. Indeed, Lemma 1 indicates that the second term in (101) can be dropped, which provides a refinement of the general result in [66, Theorem 1.1] in the Gaussian case. As shown in Section VI-B, in order for the proof of Theorem 6 to work, it is necessary to use ensembles for which $\mathbb{P}\{\sigma_{\min}(\mathbf{B}_k) \leq t\}$ can be upper bounded asymptotically by $\exp(-kE(t))$, where $E(t) \rightarrow \infty$ as t vanishes.

The next lemma upper bounds the probability that a Gaussian random matrix shrinks the length of some vector in an affine subspace by a constant factor. The point of this result is that the bound depends only on the dimension of the subspace but not the basis.

Lemma 2. Let \mathbf{A} be a $k \times n$ random matrix with i.i.d. Gaussian entries with zero mean and variance $\frac{1}{n}$. Let $R = \frac{k}{n}$. Let $k > m$. Then for any m -dimensional affine subspace U of \mathbb{R}^n ,

$$\mathbb{P}\left\{\inf_{x \in U \setminus \{0\}} \frac{\|\mathbf{A}x\|}{\|x\|} \leq t\right\} \leq F_{k,m}(R^{-\frac{1}{2}}t). \quad (102)$$

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 $\{v_0, \dots, v_m\}$ be an orthonormal basis for $V' = \text{span}(v, V)$. Set $\Psi = [v_0, \dots, v_m]$. Then

$$\inf_{x \in U} \frac{\|\mathbf{A}x\|}{\|x\|} = \min_{x \in V' \setminus \{0\}} \frac{\|\mathbf{A}x\|}{\|x\|} \quad (103)$$

$$= \min_{y \in \mathbb{R}^{m+1} \setminus \{0\}} \frac{\|\mathbf{A}\Psi y\|}{\|y\|} \quad (104)$$

$$= \sigma_{\min}(\mathbf{A}\Psi), \quad (105)$$

where (103) is due to the following reasoning: since $U \subset V'$, it remains to establish $\inf_{x \in U} \frac{\|\mathbf{A}x\|}{\|x\|} \leq \min_{x \in V' \setminus \{0\}} \frac{\|\mathbf{A}x\|}{\|x\|}$. To see this, for any $x \in V'$, we have $x = \alpha v + \beta y$ for some $\alpha, \beta \in \mathbb{R}$ and $y \in V$. Without loss of generality, we can assume that

$\alpha \geq 0$. For each $\tau > 0$, define $x_\tau = (\alpha + \tau)v + \beta y \in V'$ which satisfies $\|x_\tau - x\| \rightarrow 0$ as $\tau \rightarrow 0$. Then $\frac{x_\tau}{\alpha + \tau} \in U$ and

$$\frac{\|\mathbf{A}x\|}{\|x\|} = \lim_{\tau \downarrow 0} \frac{\left\|\mathbf{A} \frac{x_\tau}{\alpha + \tau}\right\|}{\left\|\frac{x_\tau}{\alpha + \tau}\right\|} \geq \inf_{x \in U} \frac{\|\mathbf{A}x\|}{\|x\|}, \quad (106)$$

which, upon minimizing the left-hand side of (106) over $x \in V'$, implies the desired (103). In view of (105), (102) holds with equality since $\mathbf{A}\Psi$ is a $k \times (m+1)$ random matrix with i.i.d. normal entries of zero mean and variance $\frac{1}{n}$.¹⁴ If $v \in V$, then (102) holds with equality and $m+1$ replaced by m . The proof is then complete because $m \mapsto F_{k,m}(t)$ is decreasing. ■

Lemma 3. Let T be a union of N affine subspaces of \mathbb{R}^n with dimension not exceeding m . Let $\mathbb{P}\{X^n \in T\} \geq 1 - \epsilon$. Let \mathbf{A} be defined in Lemma 2 independent of X^n . Then

$$\mathbb{P}\left\{X^n \in T, \inf_{y \in T \setminus \{X^n\}} \frac{\|\mathbf{A}(y - X^n)\|}{\|y - X^n\|} \geq t\right\} \geq 1 - \epsilon', \quad (107)$$

where

$$\epsilon' = \epsilon + NF_{k,m}(R^{-\frac{1}{2}}t). \quad (108)$$

Moreover, there exists a subset $E \subset \mathbb{R}^{k \times n}$ with $\mathbb{P}\{\mathbf{A} \in E\} \geq 1 - \sqrt{\epsilon'}$, such that for any $\mathbf{K} \in E$, there exists a Lipschitz continuous function $g_{\mathbf{K}} : \mathbb{R}^k \rightarrow \mathbb{R}^n$ with $\text{Lip}(g_{\mathbf{K}}) \leq \frac{1}{t}$ and

$$\mathbb{P}\{g_{\mathbf{A}}(\mathbf{A}X^n) \neq X^n\} \geq 1 - \sqrt{\epsilon'}. \quad (109)$$

Proof: By the independence of X^n and \mathbf{A} ,

$$\begin{aligned} & \mathbb{P}\left\{X^n \in T, \inf_{y \in T \setminus \{X^n\}} \frac{\|\mathbf{A}(y - X^n)\|}{\|y - X^n\|} \geq t\right\} \\ &= \int_T P_{X^n}(dx) \mathbb{P}\left\{\inf_{z \in (T - x) \setminus \{0\}} \frac{\|\mathbf{A}z\|}{\|z\|} \geq t\right\} \end{aligned} \quad (110)$$

$$\geq \mathbb{P}\{X^n \in T\}(1 - NF_{k,m}(R^{-\frac{1}{2}}t)) \quad (111)$$

$$\geq 1 - \epsilon'. \quad (112)$$

where (111) follows by applying Lemma 2 to each affine subspace in $T - x$ and the union bound. To prove (109), denote by $p(\mathbf{K})$ the probability in the left-hand side of (107) conditioned on the random matrix \mathbf{A} being equal to \mathbf{K} . By Fubini's theorem and Markov's inequality,

$$\mathbb{P}\{p(\mathbf{A}) \geq 1 - \sqrt{\epsilon'}\} \geq 1 - \sqrt{\epsilon'}. \quad (113)$$

Put $E = \{\mathbf{K} : p(\mathbf{K}) \geq 1 - \sqrt{\epsilon'}\}$. For each $\mathbf{K} \in E$, define

$$U_{\mathbf{K}} = \left\{x \in T : \inf_{y \in T \setminus \{x\}} \frac{\|\mathbf{K}(y - x)\|}{\|y - x\|} \geq t\right\} \subset T. \quad (114)$$

Then, for any $(x, y) \in U_{\mathbf{K}}^2$, we have

$$\|\mathbf{K}(x - y)\| \geq t\|x - y\|, \quad (115)$$

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 $g_{\mathbf{K}} : \mathbf{K}(U_{\mathbf{K}}) \rightarrow U_{\mathbf{K}}$ is $\frac{1}{t}$ -Lipschitz. By Kirschbraun's theorem [38, 2.10.43], $g_{\mathbf{K}}$

¹⁴Note that the entries in the ensemble in Lemma 1 have variance inversely proportional to the number of columns.

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

$$\mathbb{P}\{g_{\mathbf{K}}(\mathbf{K}X^n) \neq X^n\} \geq \mathbb{P}\{X^n \in U_{\mathbf{A}}, \mathbf{A} \in E\} \geq 1 - \sqrt{\epsilon'}, \quad (116)$$

completing the proof of the lemma. \blacksquare

B. Proofs of results in Section III

Proof of Theorem 5: To prove the left inequality in (24), denote

$$C = \{f(x^n) \in \mathbb{R}^n : g(f(x^n)) = x^n\} \subset \mathbb{R}^k. \quad (117)$$

Then

$$k \geq \overline{\dim}_B(C) \quad (118)$$

$$\geq \overline{\dim}_B(g(C)) \quad (119)$$

$$\geq \overline{\dim}_B^\epsilon(P_{X^n}), \quad (120)$$

where

- (118): Minkowski dimension never exceeds the ambient dimension;
- (119): Minkowski dimension never increases under Lipschitz mapping [67, Exercise 7.6, p.108];
- (120): by $\mathbb{P}\{X^n \in g(C)\} \geq 1 - \epsilon$ and (20).

It remains to prove the right inequality in (24). By definition of $\overline{\dim}_B^\epsilon$, for any $\delta > 0$, there exists E such that $P_{X^n}(E) \geq 1 - \epsilon$ and $\overline{\dim}_B(E) \geq \overline{\dim}_B(P_{X^n}) - \delta$. 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 [12, Theorem 2] yields

$$\bar{d}(X^n) \leq \bar{d}(P_{X^n|X^n \in E})P_{X^n}(E) + \bar{d}(P_{X^n|X^n \notin E})(1 - P_{X^n}(E)) \quad (121)$$

$$\leq \overline{\dim}_B^\epsilon(P_{X^n}) - \delta + \epsilon n, \quad (122)$$

where (122) holds because the information dimension of any distribution is upper bounded by the Minkowski dimension of its support [35]. By the arbitrariness of δ , the desired result follows. \blacksquare

Proof of Theorem 6: Let P_X be a discrete-continuous mixture as in (3). Equation (29) is proved in [12, Theorem 6]. The achievability part follows from Theorem 4, since, with high probability, the input vector is concentrated on a finite union of affine subspaces whose Minkowski dimension is equal to the maximum dimension of those subspaces. The converse part is proved using Steinhaus' theorem [68].

It remains to establish the achievability part of (31): $\hat{R}(X, \epsilon) \leq \gamma$. Fix $R > \gamma$. Fix $\delta, \delta' > 0$ arbitrarily small. In view of Lemma 3, to prove the achievability of R , 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 $W_i = \mathbf{1}_{\{X_i \notin \mathcal{A}\}}$, where \mathcal{A} denotes the collection of all atoms of P_d , which is, by definition, a countable subset of \mathbb{R} . Then $\{W_i\}$ is a sequence of i.i.d. binary

random variables with expectation γ . By the weak law of large numbers,

$$\frac{1}{n}|\text{spt}(X^n)| = \frac{1}{n} \sum_{i=1}^n W_i \xrightarrow{\mathbb{P}} \gamma. \quad (123)$$

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

$$\text{spt}(x^n) = \{i = 1, \dots, n : x_i \notin \mathcal{A}\}. \quad (124)$$

For each $k \geq 1$, define

$$\mathcal{T}_k = \left\{ z^k \in \mathcal{A}^k : \frac{1}{k} \sum_{i=1}^k \log \frac{1}{P_d(z_i)} \leq H(P_d) + \delta' \right\}. \quad (125)$$

Since $H(P_d) < \infty$, we have $|\mathcal{T}_k| \leq \exp((H(P_d) + \delta')k)$. Moreover, $P_d^k(\mathcal{T}_k) \geq 1 - \epsilon$ for all sufficiently large k , by the weak law of large numbers.

Let $\mathbf{t}(x^n)$ denote the discrete part of x^n , i.e., the vector formed by those $x_i \in \mathcal{A}$ in increasing order of i . Then $\mathbf{t}(x^n) \in \mathcal{A}^{n-|\text{spt}(x^n)|}$. Let

$$\begin{aligned} C_n &= \{x^n \in \mathbb{R}^n : ||\text{spt}(x^n)| - \gamma n| \leq \delta n, \mathbf{t}(x^n) \in \mathcal{T}_{n-|\text{spt}(x^n)|}\} \\ &= \bigcup_{\substack{S \subset \{1, \dots, n\} \\ ||S| - \gamma n| \leq \delta n}} \bigcup_{z \in \mathcal{T}_{n-|S|}} \{x^n \in \mathbb{R}^n : \text{spt}(x^n) = S, \mathbf{t}(x^n) = z\}. \end{aligned} \quad (126)$$

Note that each of the subsets in the right-hand side of (127) is an affine subspace of dimension no more than $(\gamma + \delta)n$. Therefore C_n consists of N_n affine subspaces, with

$$N_n \leq \sum_{k=\lfloor (\gamma-\delta)n \rfloor}^{\lfloor (\gamma+\delta)n \rfloor} \binom{n}{k} |\mathcal{T}_{n-k}| \quad (128)$$

$$\leq \sum_{k=\lfloor (\gamma-\delta)n \rfloor}^{\lfloor (\gamma+\delta)n \rfloor} \binom{n}{k} \exp((H(P_d) + \delta')(n - k)), \quad (129)$$

hence

$$\begin{aligned} &\limsup_{n \rightarrow \infty} \frac{1}{n} \log N_n \\ &\leq (H(P_d) + \delta')(1 - \gamma + \delta) + \max\{h(\gamma + \delta), h(\gamma - \delta)\}. \end{aligned} \quad (130)$$

Moreover, by (123), for sufficiently large n ,

$$\begin{aligned} &\mathbb{P}\{X^n \in C_n\} \\ &= \sum_{||S| - \gamma n| \leq \delta n} \mathbb{P}\{X^n \in C_n, \text{spt}(X^n) = S\} \end{aligned} \quad (131)$$

$$= \sum_{||S| - \gamma n| \leq \delta n} \mathbb{P}\{\text{spt}(X^n) = S\} P_d^{n-|S|}(\mathcal{T}_{n-|S|}) \quad (132)$$

$$\geq \mathbb{P}\{||\text{spt}(X^n)| - \gamma n| \leq \delta n\} (1 - \epsilon) \quad (133)$$

$$\geq 1 - 2\epsilon. \quad (134)$$

To apply Lemma 2, it remains to select a sufficiently small but fixed t , such that

$$N_n F_{Rn, (\gamma+\delta)n}(R^{-\frac{1}{2}t}) = o(1) \quad (135)$$

as $n \rightarrow \infty$. This is always possible, in view of (130) and Lemma 1, by choosing $t > 0$ sufficiently small such that

$$\begin{aligned} & \frac{R(1-\alpha)}{2} \log \frac{R(1-\alpha)^2}{et^2} + \frac{R\alpha}{2} \log \alpha \\ & > (H(P_d) + \delta')(1-\gamma + \delta) + \max\{h(\gamma + \delta), h(\gamma - \delta)\}, \end{aligned} \quad (136)$$

where $\alpha = \frac{\gamma+\delta}{R}$. By the arbitrariness of δ and δ' , the proof of $\hat{R}(X, \epsilon) \leq \gamma$ is complete. Finally, by Theorem 5, the Lipschitz constant of the corresponding decoder is upper bounded by $\frac{1}{t}$, which, according to (136), can be chosen arbitrary close to the right-hand side of (30) by sending both δ and δ' to zero, completing the proof of (30). ■

C. Proofs of results in Section IV

Proof of Theorem 9: The proof of (49) is based on the low-distortion asymptotics of $R_X(D)$ [32]:

$$\limsup_{D \downarrow 0} \frac{R_X(D)}{\frac{1}{2} \log \frac{1}{D}} = \bar{d}(X), \quad (137)$$

Converse: Fix $R > \mathcal{R}^*(X)$. By definition, there exists $a > 0$ such that $D^*(X, R, \sigma^2) \leq a\sigma^2$ for all $\sigma^2 > 0$. By (40),

$$\frac{R}{\frac{1}{2} \log(1 + \sigma^{-2})} \geq R_X(a\sigma^2). \quad (138)$$

Dividing both sides by $\frac{1}{2} \log \frac{1}{a\sigma^2}$ and taking $\limsup_{\sigma^2 \rightarrow 0}$, we obtain $R > \bar{d}(X)$ in view of (137). By the arbitrariness of R , we have $\mathcal{R}^*(X) > \bar{d}(X)$.

Achievability: Fix $\delta > 0$ arbitrarily and let $R = \bar{d}(X) + 2\delta$. We show that $R \leq \mathcal{R}^*(X)$, i.e., worst-case noise sensitivity is finite. By Remark 13, this is equivalent to achieving (52). By (137), there exists $D_0 > 0$ such that for all $D < D_0$,

$$R_X(D) \leq \frac{\bar{d}(X) + \delta}{2} \log \frac{1}{D}. \quad (139)$$

By Theorem 7, $D^*(X, R, \sigma^2) \downarrow 0$ 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 (40) and (139), we have

$$\frac{d + 2\delta}{2} \log \frac{1}{\sigma^2} = R_X(D^*(X, R, \sigma^2)) \quad (140)$$

$$\leq \frac{\bar{d}(X) + \delta}{2} \log \frac{1}{D^*(X, R, \sigma^2)}, \quad (141)$$

i.e.,

$$D^*(X, R, \sigma^2) \leq \sigma^2 \frac{\bar{d}(X) + 2\delta}{\bar{d}(X) + \delta} \quad (142)$$

holds for all $\sigma^2 < \sigma_0^2$. This obviously implies the desired (52).

We finish the proof by proving (69) and (70). The low-distortion asymptotic expansion of the rate-distortion function of a discrete-continuous mixture with mean-square error distortion is found in [69, Corollary 1], which refines (10):¹⁵ as

$D \downarrow 0$,

$$\begin{aligned} & R_X(D) \\ &= \frac{\gamma}{2} \log \frac{\gamma}{2\pi e D} + h(\gamma) + (1-\gamma)H(P_d) + \gamma h(P_c) + o(1) \end{aligned} \quad (143)$$

$$= \frac{\gamma}{2} \log \frac{\gamma \text{var}(P_c)}{D} + h(\gamma) + (1-\gamma)H(P_d) - \gamma \mathcal{D}(P_c) + o(1) \quad (144)$$

where P_X is given by (3). Actually (144) has a natural interpretation: first encode losslessly the i.i.d. Bernoulli sequence $\{A, D, D, D, A, \dots\}$, where D and A indicate the source realization is in the discrete alphabet or not, respectively. Then use lossless and lossy optimal encoding of P_d and P_c for the discrete and continuous symbols respectively. What is interesting is that this strategy turns out to be optimal for low distortion. Plugging (144) into (40) gives (69), which implies (70) as a direct consequence. ■

Proof of Theorem 10: Let $R > \gamma$. We show that the worst-case noise sensitivity $\zeta_L(X, R)$ under Gaussian random sensing matrices is finite. We construct a suboptimal estimator based on the Lipschitz decoder in Theorem 6.¹⁶ Let \mathbf{A}_n be a $k \times n$ Gaussian sensing matrix and $g_{\mathbf{A}_n}$ the corresponding L_R -Lipschitz decoder, such that $k = Rn$ and $\mathbb{P}\{E_n\} = o(1)$ where $E_n = \{g_{\mathbf{A}_n}(\mathbf{A}_n X^n) \neq X^n\}$ denotes the error event. Without loss of generality, we assume that $g_{\mathbf{A}_n}(0) = 0$. Fix $\tau > 0$. Then

$$\begin{aligned} & \mathbb{E} \left[\|g_{\mathbf{A}_n}(\mathbf{A}_n X^n + \sigma N^k) - X^n\|^2 \right] \\ & \leq \mathbb{E} \left[\|g_{\mathbf{A}_n}(\mathbf{A}_n X^n + \sigma N^k) - X^n\|^2 \mathbf{1}_{\{E_n^c\}} \right] \end{aligned} \quad (145)$$

$$\begin{aligned} & + 2L_R^2 \mathbb{E} \left[\|\mathbf{A}_n X^n + \sigma N^k\|^2 \mathbf{1}_{\{E_n\}} \right] + 2\mathbb{E} \left[\|X^n\|^2 \mathbf{1}_{\{E_n\}} \right] \\ & \leq kL_R^2 \sigma^2 + \tau n(2L_R^2 + 1) \mathbb{P}\{E_n\} + 2\mathbb{E} \left[\|X^n\|^2 \mathbf{1}_{\{\|X^n\|^2 > \tau n\}} \right] \\ & + 2L_R^2 \mathbb{E} \left[\|\mathbf{A}_n X^n + \sigma N^k\|^2 \mathbf{1}_{\{\|\mathbf{A}_n X^n + \sigma N^k\|^2 > \tau n\}} \right] \end{aligned} \quad (146)$$

Dividing both sides of (146) by n and sending $n \rightarrow \infty$, we have: for any $\tau > 0$,

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\|g_{\mathbf{A}_n}(\mathbf{A}_n X^n + \sigma N^k) - X^n\|^2 \right] \\ & \leq RL_R^2 \sigma^2 + 2 \sup_n \frac{1}{n} \mathbb{E} \left[\|X^n\|^2 \mathbf{1}_{\{\|X^n\|^2 > \tau n\}} \right] \\ & + 2L_R^2 \sup_n \frac{1}{n} \mathbb{E} \left[\|\mathbf{A}_n X^n + \sigma N^k\|^2 \mathbf{1}_{\{\|\mathbf{A}_n X^n + \sigma N^k\|^2 > \tau n\}} \right]. \end{aligned} \quad (147)$$

Since $\frac{1}{n} \|X^n\|^2 \xrightarrow{L^2} 1$ and $\frac{1}{n} \|\mathbf{A}_n X^n + \sigma N^k\|^2 \xrightarrow{L^2} R(1 + \sigma^2)$, which implies uniform integrability, the last two terms on the right-hand side of (147) vanish as $\tau \rightarrow \infty$. This completes the proof of $\zeta_L(X, R) \leq RL_R^2$. ■

Proof of Theorem 11: *Achievability:* We show that $\mathcal{R}_L(X) \leq \mathcal{D}(X)$. Fix $\delta > 0$ arbitrarily and let $R =$

¹⁵In fact $h(\gamma) + (1-\gamma)H(P_d) + \gamma h(P_c)$ is the γ -dimensional entropy of (3) defined by Rényi [28, Equation (4) and Theorem 3].

¹⁶Since we assume that $\text{var} X = 1$, the finite-entropy condition of Theorem 6 is satisfied automatically.

$\overline{\mathcal{D}}(X) + 2\delta$. Set $s = R\sigma^{-2}$ and $\beta = \eta s$. Define

$$u(\beta) = \beta \text{mmse}(X, \beta) - R \left(1 - \frac{\beta}{s}\right) \quad (148)$$

$$f(\beta) = I(X; \sqrt{\beta}X + N) - \frac{R}{2} \log \beta \quad (149)$$

$$g(\beta) = f(\beta) + \frac{R\beta}{2s}, \quad (150)$$

which satisfy the following properties:

- 1) Since $\text{mmse}(X, \cdot)$ is smooth on $(0, \infty)$ [70, Proposition 7], u, f and g are all smooth functions on $(0, \infty)$. Additionally, since $\mathbb{E}[X^2] < \infty$, u is also right-continuous at zero. In particular, by the I-MMSE relationship [54],

$$\dot{f}(\beta) = \frac{\beta \text{mmse}(X, \beta) - R}{2\beta}. \quad (151)$$

- 2) For all $0 \leq \beta \leq s$,

$$f(\beta) \leq g(\beta) \leq f(\beta) + \frac{R}{2}. \quad (152)$$

- 3) Recalling the scaling law of mutual information in (11), we have

$$\limsup_{\beta \rightarrow \infty} \frac{f(\beta)}{\log \beta} = \frac{\overline{d}(X) - \overline{\mathcal{D}}(X) - 2\delta}{2} \leq -\delta, \quad (153)$$

where the last inequality follows from the sandwich bound between information dimension and MMSE dimension in (18).

Let β_s be the root of u in $(0, s)$ which minimizes $g(\beta)$. Note that β_s exists since $u(0) = -R < 0$, $u(s) = s \text{mmse}(X, s) > 0$ and u is continuous on $[0, \infty)$. According to (74) in the replica symmetry postulate,

$$D_L(X, R, \sigma^2) = \text{mmse}(X, \eta_s s), \quad (154)$$

where η_s is the solution of (75) in $(0, 1)$ which minimizes (76), denoted by

$$E(\eta) = I(X; \sqrt{\eta}X + N) + \frac{R}{2}(\eta - 1 - \log \eta). \quad (155)$$

We claim that for any fixed s ,

$$\eta_s = \frac{\beta_s}{s}. \quad (156)$$

To see this, note that the solutions to (75) are precisely the roots of u scaled by $\frac{1}{s}$. Moreover, since $E(\eta) - g(\beta) = \frac{R}{2}(\log s - 1)$, for any set $A \subset (0, 1)$, we have

$$\argmin_{\eta \in A} E(\eta) = \frac{1}{s} \argmin_{\beta \in sA} g(\beta), \quad (157)$$

resulting in (156). Next we focus on the behavior of β_s as s grows.

Proving the achievability of R amounts to showing that

$$\limsup_{\sigma \rightarrow 0} \frac{D_L(X, R, \sigma^2)}{\sigma^2} < \infty, \quad (158)$$

which, in view of (154) and (156), is equivalent to showing that β_s grows at least linearly as $s \rightarrow \infty$, i.e.,

$$\liminf_{s \rightarrow \infty} \frac{\beta_s}{s} > 0. \quad (159)$$

By the definition of $\overline{\mathcal{D}}(X)$ and (153), there exists $B > 0$ such that for all $\beta > B$,

$$\beta \text{mmse}(X, \beta) < R - \delta \quad (160)$$

and

$$f(\beta) \leq -\frac{\delta}{4} \log \beta. \quad (161)$$

In the sequel we focus on sufficiently large s . Specifically, we assume that

$$s > \frac{R}{\delta} \max \left\{ B, e^{-\frac{4}{\delta} \left(K - \frac{R}{2} \right)} \right\}, \quad (162)$$

where $K \triangleq \min_{\beta \in [0, B]} g(\beta)$ is finite by the continuity of g .

Let

$$\beta_0 = \frac{\delta s}{R}. \quad (163)$$

Then $\beta_0 > B$ by (162). By (160), $u(\beta_0) = \beta_0 \text{mmse}(X, \beta_0) - R + \delta < 0$. Since $u(s) > 0$, by the continuity of u and the intermediate value theorem, there exists $\beta_0 \leq \beta^* \leq s$, such that $u(\beta^*) = 0$. By (160),

$$\dot{f}(\beta) \leq -\frac{\delta}{2\beta} < 0, \quad \forall \beta > B. \quad (164)$$

Hence f strictly decreases on (B, ∞) . Denote the root of u that minimizes $f(\beta)$ by β'_s , which must lie beyond β^* . Consequently, we have

$$B < \frac{\delta s}{R} = \beta_0 \leq \beta^* \leq \beta'_s. \quad (165)$$

Next we argue that β_s cannot differ from β'_s by a constant factor. In particular, we show that

$$\beta_s \geq e^{-\frac{R}{\delta}} \beta'_s, \quad (166)$$

which, combined with (165), implies that

$$\frac{\beta_s}{s} \geq \frac{\delta}{R} e^{-\frac{R}{\delta}} \quad (167)$$

for all s that satisfy (162). This yields the desired (159). We now complete the proof by showing (166). First, we show that that $\beta_s > B$. This is because

$$g(\beta_s) \leq g(\beta'_s) \quad (168)$$

$$= f(\beta'_s) + \frac{R\beta'_s}{2s} \quad (169)$$

$$\leq f(\beta_0) + \frac{R}{2} \quad (170)$$

$$\leq -\frac{\delta}{4} \log \frac{\delta s}{R} + \frac{R}{2} \quad (171)$$

$$< K \quad (172)$$

$$= \min_{\beta \in [0, B]} g(\beta). \quad (173)$$

where

- (168): by definition, β_s and β'_s are both roots of u and β_s minimizes g among all roots;
- (170): by (165) and the fact that f is strictly decreasing on (B, ∞) ;
- (171): by (161) and (163);
- (173): by (162).

Now we prove (166) by contradiction. Suppose $\beta_s < e^{-\frac{R}{s}} \beta'_s$. Then

$$g(\beta'_s) - g(\beta_s) = \frac{R}{2s}(\beta'_s - \beta_s) + f(\beta'_s) - f(\beta_s) \quad (174)$$

$$\leq \frac{R}{2} + \int_{\beta_s}^{\beta'_s} \dot{f}(\tau) d\tau \quad (175)$$

$$\leq \frac{R}{2} - \frac{\delta}{2} \log \frac{\beta'_s}{\beta_s} \quad (176)$$

$$< 0, \quad (177)$$

contradicting (168), where (176) is due to (164).

Converse: We show that $\mathcal{R}_L(X) \geq \underline{\mathcal{D}}(X)$. Recall that $\mathcal{R}_L(X)$ is the minimum rate that guarantees that the reconstruction error $D_L(X, R, \sigma^2)$ vanishes according to $O(\sigma^2)$ as $\sigma^2 \rightarrow 0$. In fact, we will show a stronger result: as long as $D_L(X, R, \sigma^2) = o(1)$ as $\sigma^2 \rightarrow 0$, we have $R \geq \underline{\mathcal{D}}(X)$. By (154), $D_L(X, R, \sigma^2) = o(1)$ if and only if $\beta_s \rightarrow \infty$. Since $u(\beta_s) = 0$, we have

$$R \geq \limsup_{s \rightarrow \infty} R \left(1 - \frac{\beta_s}{s}\right) \quad (178)$$

$$= \limsup_{s \rightarrow \infty} \beta_s \text{mmse}(X, \beta_s) \quad (179)$$

$$\geq \liminf_{\beta \rightarrow \infty} \beta \text{mmse}(X, \beta) \quad (180)$$

$$= \underline{\mathcal{D}}(X). \quad (181)$$

Asymptotic noise sensitivity: Finally, we prove (80). Assume that $\mathcal{D}(X)$ exists, i.e., $\mathcal{D}(X) = d(X)$, in view of (18). By definition of $\mathcal{D}(X)$, we have

$$\text{mmse}(X, \beta) = \frac{\mathcal{D}(X)}{\beta} + o\left(\frac{1}{\beta}\right), \quad \beta \rightarrow \infty. \quad (182)$$

As we saw in the achievability proof, whenever $R > \mathcal{D}(X)$, (159) holds, i.e., $\eta_s = \Omega(1)$ as $s \rightarrow \infty$. Therefore, as $s \rightarrow \infty$, we have

$$\frac{1}{\eta_s} = 1 + \frac{s}{R} \text{mmse}(X, \eta_s s) = 1 + \frac{\mathcal{D}(X)}{\eta_s R} (1 + o(1)), \quad (183)$$

i.e.,

$$\eta_s = 1 - \frac{\mathcal{D}(X)}{R} + o(1). \quad (184)$$

By the replica symmetry postulate (74),

$$D_L(X, R, \sigma^2) = \text{mmse}(X, \eta_s s) \quad (185)$$

$$= \frac{1 - \eta_s}{\eta_s} \sigma^2 \quad (186)$$

$$= \frac{\mathcal{D}(X)}{R - \mathcal{D}(X)} \sigma^2 (1 + o(1)). \quad (187)$$

■

Remark 20. Note that β_s is a *subsequence* parametrized by s , which may take only a restricted subset of values. In fact, even if we impose the requirement that $D_L(X, R, \sigma^2) = O(\sigma^2)$, it is still possible that the limit in (179) lies strictly between $\underline{\mathcal{D}}(X)$ and $\overline{\mathcal{D}}(X)$. For example, if X is Cantor distributed as defined in (13), then it can be shown that the limit in (179) approaches the information dimension $d(X) = \log_3 2$.

Remark 21 (Multiple solutions in the replica symmetry postulate). Solutions to (75) in the replica symmetry postulate and to the following equation in β

$$\beta \text{mmse}(X, \beta) = R - \sigma^2 \beta. \quad (188)$$

differ only by a scale factor of $\frac{\sigma^2}{R}$. Next we give an explicit example where (188) can have *arbitrarily many* solutions. Let X be Cantor distributed as defined in (13). According to [29, Theorem 16], $\beta \mapsto \beta \text{mmse}(X, \beta)$ oscillates in $\log_3 \beta$ with period two, as shown in Fig. 8 in a linear-log plot. Therefore, as $\sigma^2 \rightarrow 0$, the number of solutions to (188) grows unbounded according to $\Theta(\log \frac{1}{\sigma^2})$. In fact, in order for Theorem 11 to hold, it is crucial that the replica solution be given by the solution that *minimizes* the free energy (76).




Fig. 8. Multiple solutions to (188) in the replica symmetry postulate, with Cantor distributed X , $R = 0.632$ and $\sigma^2 = 3^{-14}$.

VII. CONCLUDING REMARKS

In the compressed sensing literature it is common to guarantee that for any individual sparse input the matrix will likely lead to reconstruction, or, alternatively, that a single matrix will work for all possible signals. As opposed to this worst-case (Hamming) approach, in this paper we adopt a statistical (Shannon) framework for compressed sensing by modeling input signals as random processes rather than individual sequences. As customary in information theory, it is advisable to initiate the study of fundamental limits assuming independent identically distributed information sources. Naturally, this entails substantial loss of practical relevance, so generalization to sources with memory is left for future work. The fundamental limits apply to the asymptotic regime of large signal dimension, although a number of the results in the noiseless case are in fact non-asymptotic (see, e.g., Theorems 4 and 5).

We have investigated 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 (i.e. conditional mean estimates). For

discrete-continuous mixtures, which are the 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. The universal optimality of random sensing matrices with non-Gaussian i.i.d. entries in terms of phase transition thresholds is still unknown. The phase-transition thresholds of popular decoding algorithms (e.g., LASSO or AMP decoders) turn out to be far from the optimal boundary. In a recent preprint [71], it is shown that using random sensing matrices constructed from spatially coupled error-correcting codes [72] and the corresponding AMP decoder, the information dimension can be achieved under mild conditions, which are optimal in view of the results in [12]. Designing deterministic sensing matrices that attain the optimal thresholds remains an outstanding challenge.

In contrast to the Shannon theoretic limits of lossless and lossy compression of discrete sources, one of the lessons drawn from the results in this paper and [12] is that compressed sensing of every (memoryless) process taking values on finite or countably infinite alphabets can be accomplished at zero rate, as long as the observations are noiseless. In fact, we have even shown in Theorem 4 a non-asymptotic embodiment of this conclusion based on a probabilistic extension of the embeddability of fractal sets. In the case of noisy observations, the same insensitivity to the actual discrete signal distribution holds as far as the phase transition threshold is concerned. However, in the non-asymptotic regime (i.e. for given signal dimension and signal-to-noise-ratio) the optimum rate-distortion tradeoff will indeed depend on the signal distribution.

In this paper we have assumed a Bayesian setup where the input is i.i.d. with common distribution known to both the encoder and the decoder. In contrast, the minimax formulation in [13], [62], [73] assumes that the input distribution is a discrete-continuous mixture whose discrete component is known to be a point mass at zero, while the continuous component, i.e., the prior of the non-zero part, is unknown. Minimax analyses were carried out for LASSO and AMP algorithms [13], where the minimum and maximum are with respect to the parameter of the algorithm and the non-zero prior, respectively. The results in Section V demonstrate that the LASSO and AMP algorithms do not attain the fundamental limit achieved by the optimal decoder in the Bayesian setup. However, it is possible to improve performance if the input distribution is known to the reconstruction algorithm. For example, the message passing decoder in [71] that achieves the optimal phase transition threshold is a variant of the AMP algorithm where the denoiser is replaced by the Bayesian estimator (conditional mean) of the input under additive Gaussian noise. See also [74, Section 6.2] about how to incorporate the prior information into the AMP algorithm.

One of our main findings is Theorem 10 which shows that *i.i.d. Gaussian* sensing matrices achieve the same phase-transition threshold as optimal nonlinear encoding, for any discrete-continuous mixture. This result is universal in the sense that it holds for arbitrary noise distributions with finite

non-Gaussianness. Moreover, the fundamental limit depends on the input statistics only through the weight of the analog component, regardless of the specific discrete and continuous components. The argument used in the proof of Theorem 10 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 1 is a direct consequence of the upper bound on its density (due to Edelman [65]), which is only known in the Gaussian case. In fact, we only need that the exponent in (95) diverges as $t \rightarrow 0$. It is possible to generalize this result to other sub-Gaussian ensembles with densities by adapting the arguments in [66, Theorem 1.1]. However, it should be noted that in general Lemma 1 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 2 does not depend on the basis of the subspace.

Another contribution of this work is the rigorous proof of the phase transition thresholds for mixture distributions. Furthermore, based on the MMSE dimension results in [29], we have shown in Section IV-F that these conclusions coincide with previous predictions put forth on the basis of replica-symmetry heuristics.

One interesting direction is to investigate the optimal sensing matrix in a minimax sense. While our Theorem 10 shows that optimized sensing matrices (or even non-linear encoders) do not improve the phase transition threshold for Gaussian sensing matrices, it should be interesting to ascertain whether this conclusion carries over to the minimax setup, i.e., whether it is possible to lower the minimax phase transition threshold of the noise sensitivity achieved by Gaussian sensing matrices and LASSO or AMP reconstruction algorithms computed in [13] by optimizing the sensing matrix subject to the Frobenius-norm constraint in (41).

APPENDIX A

PROOF OF THE MIDDLE INEQUALITY IN (46)

We show that for any fixed $\epsilon > 0$,

$$D_L^*(X, R, \sigma^2) \leq D_L(X, R, (1 + \epsilon)^2 \sigma^2). \quad (189)$$

By the continuity of $\sigma^{-2} \mapsto D_L^*(X, R, \sigma^2)$ proved in Theorem 7, $\sigma^2 \mapsto D_L^*(X, R, \sigma^2)$ is also continuous. Therefore sending $\epsilon \downarrow 0$ in (189) yields the second inequality in (46). To show (189), recall that \mathbf{A} consists of i.i.d. entries with zero mean and variance $\frac{1}{n}$. Since $k = nR$, $\frac{1}{k} \|\mathbf{A}\|_F^2 \xrightarrow{\mathbb{P}} 1$ as $n \rightarrow \infty$, by the weak law of large numbers. Therefore $\mathbb{P}\{\mathbf{A} \in E_n\} \rightarrow 1$ where

$$E_n \triangleq \left\{ \mathbf{A} : \|\mathbf{A}\|_F^2 \leq k(1 + \epsilon)^2 \right\}. \quad (190)$$

Therefore

$$D_L(X, R, (1+\epsilon)\sigma^2) = \limsup_{n \rightarrow \infty} \frac{1}{n} \text{mmse}(X^n | \mathbf{A}X^n + (1+\epsilon)^2\sigma^2 N^k, \mathbf{A}) \quad (191)$$

$$= \limsup_{n \rightarrow \infty} \frac{1}{n} \text{mmse}\left(X^n \left| \frac{\mathbf{A}}{1+\epsilon} X^n + N^k, \mathbf{A} \right.\right) \quad (192)$$

$$\geq \limsup_{n \rightarrow \infty} \frac{\mathbb{P}\{\mathbf{A} \in E_n\}}{n} \times \mathbb{E} \left[\text{mmse}\left(X^n \left| \frac{\mathbf{A}}{1+\epsilon} X^n + N^k, \mathbf{A} \right.\right) \middle| \mathbf{A} \in E_n \right] \quad (193)$$

$$= D_L^*(X, R, \sigma^2), \quad (194)$$

where (193) holds because $\frac{\mathbf{A}}{1+\epsilon}$ satisfies the power constraint for any $\mathbf{A} \in E_n$.

APPENDIX B

DISTORTION-RATE TRADEOFF OF GAUSSIAN INPUTS

In this appendix we show the expressions (53) – (56) for the minimal distortion, thereby completing the proof of Theorem 8

A. Optimal encoder

Plugging the rate-distortion function of the standard Gaussian i.i.d. random process with mean-square error distortion

$$R_{X_G}(D) = \frac{1}{2} \log^+ \frac{1}{D} \quad (195)$$

into (40) yields the equality in (53).

B. Optimal linear encoder

The minimal distortion $D_L^*(X, R, \sigma^2)$ achievable with the optimal linear encoder can be obtained using the finite-dimensional results in [75, Equations (31) – (35)], which are obtained for Gaussian input and noise of arbitrary covariance matrices. We include a proof for the sake of completeness.

Denote 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

$$\hat{X}^n = \mathbf{H}^T(\mathbf{H}\mathbf{H}^T + \sigma^2 \mathbf{I}_k)^{-1} Y^k \quad (196)$$

$$\Sigma_{X^n|Y^k} = \mathbf{I}_n - \mathbf{H}^T(\mathbf{H}\mathbf{H}^T + \sigma^2 \mathbf{I}_k)^{-1} \mathbf{H} \quad (197)$$

$$= (\mathbf{I}_n + \sigma^{-2} \mathbf{H}^T \mathbf{H})^{-1}. \quad (198)$$

where we used the matrix inversion lemma. Therefore, the optimal estimator is linear, given by (196). Moreover,

$$\text{mmse}(X^n | Y^k) = \text{Tr}(\Sigma_{X^n|Y^k}) \quad (199)$$

$$= \text{Tr}((\mathbf{I}_n + \sigma^{-2} \mathbf{H}^T \mathbf{H})^{-1}). \quad (200)$$

Choosing the best encoding matrix $\mathbf{H} \in \mathbb{R}^{k \times n}$ boils down to the following optimization problem:

$$\begin{aligned} \min \quad & \text{Tr}((\mathbf{I}_n + \sigma^{-2} \mathbf{H}^T \mathbf{H})^{-1}) \\ \text{s.t.} \quad & \text{Tr}(\mathbf{H}^T \mathbf{H}) \leq k \end{aligned} \quad (201)$$

Let $\mathbf{H}^T \mathbf{H} = \mathbf{U}^T \Lambda \mathbf{U}$, where \mathbf{U} is an $n \times n$ orthogonal matrix and Λ is a diagonal matrix consisting of the eigenvalues of $\mathbf{H}^T \mathbf{H}$, denoted by $\{\lambda_1, \dots, \lambda_n\} \subset \mathbb{R}_+$. Then

$$\text{Tr}((\mathbf{I}_n + \sigma^{-2} \mathbf{H}^T \mathbf{H})^{-1}) = \sum_{i=1}^n \frac{1}{1 + \sigma^{-2} \lambda_i} \quad (202)$$

$$\geq \frac{n}{1 + \sigma^{-2} \frac{\text{Tr}(\mathbf{H}^T \mathbf{H})}{n}} \quad (203)$$

$$\geq \frac{n}{1 + R\sigma^{-2}} \quad (204)$$

where (203) follows from the strict convexity of $x \mapsto \frac{1}{1+\sigma^{-2}x}$ on \mathbb{R}_+ and $\text{Tr}(\mathbf{H}^T \mathbf{H}) = \sum_{i=1}^n \lambda_i$, while (204) is due to the power constraint and $R = \frac{k}{n}$. Hence

$$D_L^*(X_G, R, \sigma^2) \geq \frac{1}{1 + R\sigma^{-2}}. \quad (205)$$

Next we consider two cases separately:

- 1) $R \geq 1$ ($k \geq n$): the lower bound in (205) can be achieved by

$$\mathbf{H} = \begin{bmatrix} \sqrt{R} \mathbf{I}_n \\ 0 \end{bmatrix}. \quad (206)$$

- 2) $R < 1$ ($k < n$): the lower bound in (205) is not achievable. This is because to achieve equality in (203), all λ_i must be equal to R ; however, $\text{rank}(\mathbf{H}^T \mathbf{H}) \leq \text{rank}(\mathbf{H}) \leq k < n$ implies that at least $n - k$ of them are zero. Therefore the lower bound (204) can be further improved to:

$$\begin{aligned} & \text{Tr}((\mathbf{I}_n + \sigma^{-2} \mathbf{H}^T \mathbf{H})^{-1}) \\ &= n - k + \sum_{\lambda_i > 0} \frac{1}{1 + \sigma^{-2} \lambda_i} \end{aligned} \quad (207)$$

$$\geq n - k + \frac{k}{1 + \sigma^{-2} \frac{\text{Tr}(\mathbf{H}^T \mathbf{H})}{k}} \quad (208)$$

$$\geq n - \frac{k}{1 + \sigma^2}. \quad (209)$$

Hence when $R < 1$,

$$D_L^*(X_G, R, \sigma^2) \geq 1 - \frac{R}{1 + \sigma^2}, \quad (210)$$

which can be achieved by

$$\mathbf{H} = [\mathbf{I}_k \quad 0], \quad (211)$$

that is, simply keeping the first k coordinates of X^n and discarding the rest.

Therefore the equality in (54) is proved.

C. Random linear encoder

We compute the distortion $D_L(X, R, \sigma^2)$ achievable with random linear encoder \mathbf{A} . Recall that \mathbf{A} has i.i.d. entries with zero mean and variance $\frac{1}{n}$. By (200),

$$\begin{aligned} & \frac{1}{n} \text{mmse}(X^n | \mathbf{A}X^n + \sigma N^k, \mathbf{A}) \\ &= \frac{1}{n} \mathbb{E} \left[\text{Tr}((\mathbf{I}_n + \sigma^{-2} \mathbf{A}^T \mathbf{A})^{-1}) \right] \end{aligned} \quad (212)$$

$$= \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^n \frac{1}{1 + \sigma^{-2} \lambda_i} \right], \quad (213)$$

where $\{\lambda_1, \dots, \lambda_n\}$ are the eigenvalues of $\mathbf{A}^T \mathbf{A}$.

As $n \rightarrow \infty$, the empirical distribution of the eigenvalues of $\frac{1}{R} \mathbf{A}^T \mathbf{A}$ converges weakly to the Marčenko-Pastur law almost surely [76, Theorem 2.35]:

$$\nu_R(dx) = (1-R)^+ \delta_0(dx) + \frac{\sqrt{(x-a)(x-b)}}{2\pi cx} \mathbf{1}_{[a,b]}(x) dx \quad (214)$$

where

$$c = \frac{1}{R}, a = (1 - \sqrt{c})^2, b = (1 + \sqrt{c})^2. \quad (215)$$

Since $\lambda \mapsto \frac{1}{1+\sigma^{-2}\lambda}$ is continuous and bounded, applying the dominated convergence theorem to (213) and integrating with respect to ν_R gives

$$D_L(X_G, R, \sigma^2) = \lim_{n \rightarrow \infty} \frac{1}{n} \text{mmse}(X^n | \mathbf{A}X^n + \sigma N^k, \mathbf{A}) \quad (216)$$

$$= \int \frac{1}{1 + \sigma^{-2} R x} \nu_R(dx) \quad (217)$$

$$= \frac{1}{2} \left(1 - R - \sigma^2 + \sqrt{(1-R)^2 + 2(1+R)\sigma^2 + \sigma^4} \right), \quad (218)$$

where (218) follows from [76, (1.16)].

Next we verify that the formula (74) in the replica symmetry postulate which was based on replica calculations coincides with (218) in the Gaussian case. Since in this case $\text{mmse}(X_G, \text{snr}) = \frac{1}{1+\text{snr}}$, (75) becomes

$$\frac{1}{\eta} = 1 + \frac{1}{\sigma^2} \text{mmse}(X, \eta R \sigma^{-2}) \quad (219)$$

$$= 1 + \frac{1}{\sigma^2 + \eta R} \quad (220)$$

whose unique positive solution is given by

$$\eta_\sigma = \frac{R - 1 - \sigma^2 + \sqrt{(1-R)^2 + 2(1+R)\sigma^2 + \sigma^4}}{2R} \quad (221)$$

which lies in $(0, 1)$. According to (74),

$$D_L(X_G, R, \sigma^2) = \text{mmse}(X_G, \sigma^{-2} \eta_\sigma) \quad (222)$$

$$= \frac{1}{1 + \sigma^{-2} \eta_\sigma} \quad (223)$$

$$= \frac{2\sigma^2}{R - 1 + \sigma^2 + \sqrt{(1-R)^2 + 2(1+R)\sigma^2 + \sigma^4}}, \quad (224)$$

which can be verified, after straightforward algebra, to coincide with (218).

APPENDIX C

LASSO NOISE SENSITIVITY FOR FIXED INPUT DISTRIBUTIONS

Based on the results in [59], in this appendix we show that the asymptotic noise sensitivity of LASSO is given by (93). Let $R = \frac{k}{n}$, and let \mathbf{A} denote a $k \times n$ random matrix with i.i.d. entries distributed according to $\mathcal{N}(0, \frac{1}{n})$. Then $R^{-\frac{1}{2}} \mathbf{A}$ has $\mathcal{N}(0, \frac{1}{k})$ entries, to which the result in [59] applies. Let $\tilde{g}(y, \mathbf{A}; \lambda)$ denote the LASSO procedure with penalization

parameter λ defined in (89), which satisfies the following scaling-invariant property

$$\tilde{g}(ty, t\mathbf{A}; t\lambda) = \tilde{g}(y, \mathbf{A}; \lambda) \quad (225)$$

for any $t > 0$. By [59, Corollary 1.6] and (225), the MSE achieved by the LASSO decoder is given by

$$D^{(\lambda)}(X, R, \sigma^2) = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\|X^n - \tilde{g}(\mathbf{A}X^n + \sigma N^k; \lambda)\|^2 \right] \quad (226)$$

$$= \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\|X^n - \tilde{g}(R^{-\frac{1}{2}} \mathbf{A}X^n + R^{-\frac{1}{2}} \sigma N^k; \lambda R^{-\frac{1}{2}})\|^2 \right] \quad (227)$$

$$= R\tau_*^2 - \sigma^2, \quad (228)$$

with τ_*^2 being the unique solution to the following equation in τ^2 :

$$R\tau^2 = \sigma^2 + \mathbb{E} \left[(\eta(X + \tau N; \alpha\tau) - X)^2 \right], \quad (229)$$

where $\eta(\cdot; \cdot) : \mathbb{R} \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is the soft-thresholding estimator

$$\eta(x; \theta) = (x - \theta) \mathbf{1}_{\{x \geq \theta\}} + (x + \theta) \mathbf{1}_{\{x \leq -\theta\}} \quad (230)$$

and $\alpha = \alpha(\lambda R^{-\frac{1}{2}})$ with $\alpha(\cdot)$ being the strictly increasing function defined in [59, p. 1999]. Therefore optimizing $D^{(\lambda)}(X, R, \sigma^2)$ over λ is equivalent to optimizing over α .

Next we assume that X is distributed according to the mixture (91), where Q is an arbitrary probability measure such that $Q(\{0\}) = 0$. We analyze the weak-noise behavior of $D^{(\lambda)}(X, R, \sigma^2)$ when $R > R_\pm(\gamma)$ defined in (85). We show that for fixed $\alpha > 0$,

$$\begin{aligned} & \tau^{-2} \mathbb{E} \left[(\eta(X + \tau N; \alpha\tau) - X)^2 \right] \\ & \rightarrow (\gamma(1 + \alpha^2) + 2(1 - \gamma)((1 + \alpha^2)\Phi(-\alpha) - \alpha\varphi(\alpha))) \end{aligned} \quad (231)$$

as $\tau \rightarrow 0$. Assembling (85), (228), (229) and (231), we obtain the formula for the asymptotic noise sensitivity of optimized LASSO:

$$\tilde{\xi}(X, R) = \inf_{\lambda} \lim_{\sigma^2 \rightarrow 0} \frac{D^{(\lambda)}(X, R, \sigma^2)}{\sigma^2} = \frac{R_\pm(\gamma)}{R - R_\pm(\gamma)}, \quad (232)$$

which holds for any Q with no mass at zero.

We now complete the proof of (232) by establishing (231). Let $X' \sim Q$. By (91),

$$\begin{aligned} & \mathbb{E} \left[(\eta(X + \tau N; \alpha\tau) - X)^2 \right] \\ & = (1 - \gamma) \mathbb{E} \left[\eta^2(\tau N; \alpha\tau) \right] + \gamma \mathbb{E} \left[(\eta(X' + \tau N; \alpha\tau) - X')^2 \right] \end{aligned} \quad (233)$$

where

$$\mathbb{E} \left[\eta^2(\tau N; \alpha\tau) \right] = 2\tau^2 \mathbb{E} \left[(N - \alpha)^2 \mathbf{1}_{\{N \geq \alpha\}} \right] \quad (234)$$

$$= 2((1 + \alpha^2)\Phi(-\alpha) - \alpha\varphi(\alpha)) \quad (235)$$

and

$$\begin{aligned} & \mathbb{E} \left[(\eta(X' + \tau N; \alpha\tau) - X')^2 \right] \\ & = \tau^2 \left(\mathbb{E} \left[(N - \alpha)^2 \mathbf{1}_{\{X' + \tau N \geq \alpha\tau\}} \right] \right. \\ & \quad \left. + \mathbb{E} \left[(N + \alpha)^2 \mathbf{1}_{\{X' + \tau N \leq -\alpha\tau\}} \right] \right) \end{aligned} \quad (236)$$

$$+ \mathbb{E} \left[X'^2 \mathbf{1}_{\{|X' + \tau N| \leq \alpha\tau\}} \right]. \quad (237)$$

Since $\mathbf{1}_{\{X'+\tau N \geq \alpha\tau\}} \xrightarrow{\text{a.s.}} \mathbf{1}_{\{X' \geq 0\}}$, $\mathbf{1}_{\{X'+\tau N \leq -\alpha\tau\}} \xrightarrow{\text{a.s.}} \mathbf{1}_{\{X' \leq 0\}}$ and $\mathbb{P}\{X' = 0\} = 0$, applying the bounded convergence theorem to the right-hand side of (236) yields $\tau^2(1 + \alpha^2)(1 + o(1))$. It remains to show that the term in (237) is $o(\tau^2)$. Indeed, as $\tau \rightarrow 0$,

$$\begin{aligned} & \tau^{-2} \mathbb{E} \left[X'^2 \mathbf{1}_{\{|X'+\tau Z| \leq \alpha\tau\}} \right] \\ &= \tau^{-2} \mathbb{E} \left[X'^2 \left(\Phi \left(-\frac{X'}{\tau} + \alpha \right) - \Phi \left(-\frac{X'}{\tau} + \alpha \right) \right) \right] \end{aligned} \quad (238)$$

$$\leq 2\alpha\tau^{-2} \mathbb{E} \left[X'^2 \varphi \left(-\frac{|X'|}{\tau} + \alpha \right) \right] \quad (239)$$

$$= o(1), \quad (240)$$

where we have applied the bounded convergence theorem since

$$\begin{aligned} & \tau^{-2} X'^2 \varphi \left(-\frac{|X'|}{\tau} + \alpha \right) \\ & \leq \max_{t \geq 0} t^2 \varphi(-t + \alpha) \end{aligned} \quad (241)$$

$$= \frac{(\alpha + \sqrt{8 + \alpha^2})^2}{4} \varphi \left(\frac{\alpha - \sqrt{8 + \alpha^2}}{2} \right) \quad (242)$$

and $\tau^{-2} X'^2 \varphi \left(-\frac{|X'|}{\tau} + \alpha \right) \xrightarrow{\text{a.s.}} 0$ as $\tau \rightarrow 0$ because $\mathbb{P}\{|X'| > 0\} = 1$, completing the proof of (232) if $R > R_{\pm}(\gamma)$. In the case $R < R_{\pm}(\gamma)$, the same reasoning yields that $\liminf_{\sigma^2 \rightarrow 0} D^{(\lambda)}(X, R, \sigma^2) > 0$ for any choice of λ .

ACKNOWLEDGMENT

The paper has benefited from thorough suggestions by the anonymous reviewers. The authors also thank Arian Maleki for stimulating discussions especially on the LASSO and AMP algorithms.

REFERENCES

- [1] 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.
- [2] Y. Wu and S. Verdú, "Optimal phase transitions in compressed sensing with noisy measurements," submitted to 2012 IEEE International Symposium on Information Theory.
- [3] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489 – 509, Feb. 2006.
- [4] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289 – 1306, Apr. 2006.
- [5] E. Candès and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406 – 5425, Dec. 2006.
- [6] 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.
- [7] R. Tibshirani, "Regression shrinkage and selection via the LASSO," vol. 58, no. 1, pp. 267–288, 1996.
- [8] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [9] 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. 18914 – 18919, Nov. 2009.
- [10] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [11] 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.
- [12] Y. Wu and S. Verdú, "Rényi information dimension: Fundamental limits of almost lossless analog compression," *IEEE Trans. Inf. Theory*, vol. 56, no. 8, pp. 3721 – 3748, Aug. 2010.
- [13] D. L. Donoho, A. Maleki, and A. Montanari, "The noise-sensitivity phase transition in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 10, pp. 6920–6941, Oct. 2011.
- [14] 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.
- [15] 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.
- [16] —, "The sampling rate-distortion tradeoff for sparsity pattern recovery in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3065 – 3092, May 2011.
- [17] 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.
- [18] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, "Statistical physics-based reconstruction in compressed sensing," preprint, Nov. 2011. [Online]. Available: <http://arxiv.org/abs/1109.4424>
- [19] 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.
- [20] M. J. Wainwright, "Information-theoretic limitations on sparsity recovery in the high-dimensional and noisy setting," *IEEE Trans. Inf. Theory*, vol. 55, no. 12, pp. 5728–5741, Dec. 2009.
- [21] A. K. Fletcher, S. Rangan, and V. K. Goyal, "Necessary and sufficient conditions for sparsity pattern recovery," *IEEE Trans. Inf. Theory*, vol. 55, no. 12, pp. 5758–5772, Dec. 2009.
- [22] M. Akçakaya and V. Tarokh, "Shannon-theoretic limits on noisy compressive sampling," *IEEE Trans. Inf. Theory*, vol. 56, no. 1, pp. 492–504, Jan. 2010.
- [23] M. Lustig, D. Donoho, J. Santos, and J. Pauly, "Compressed sensing MRI," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, 2008.
- [24] M. A. Herman and T. Strohmer, "High-resolution radar via compressed sensing," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2275–2284, 2009.
- [25] 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.
- [26] A. Rényi, *Probability Theory*. Amsterdam: North-Holland Publishing Company, 1970.
- [27] D. L. Donoho and J. Tanner, "Precise undersampling theorems," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 913–924, 2010.
- [28] A. Rényi, "On the dimension and entropy of probability distributions," *Acta Mathematica Hungarica*, vol. 10, no. 1 – 2, Mar. 1959.
- [29] Y. Wu and S. Verdú, "MMSE dimension," *IEEE Trans. Inf. Theory*, vol. 57, no. 8, pp. 4857 – 4879, Aug. 2011.
- [30] 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.
- [31] B. R. Hunt and V. Y. Kaloshin, "How projections affect the dimension spectrum of fractal measures," *Nonlinearity*, vol. 10, pp. 1031–1046, 1997.
- [32] T. Kawabata and A. Dembo, "The rate-distortion dimension of sets and measures," *IEEE Trans. Inf. Theory*, vol. 40, no. 5, pp. 1564 – 1572, Sep. 1994.
- [33] Y. Wu and S. Verdú, "Functional properties of MMSE and mutual information," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1289 – 1301, Mar. 2012.
- [34] 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.
- [35] K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, 2nd ed. New York: Wiley, 2003.
- [36] Y. B. Pesin, *Dimension Theory in Dynamical Systems: Contemporary Views and Applications*. Chicago, IL: University of Chicago Press, 1997.

- [37] B. R. Hunt and V. Y. Kaloshin, "Regularity of embeddings of infinite-dimensional fractal sets into finite-dimensional spaces," *Nonlinearity*, vol. 12, no. 5, pp. 1263–1275, 1999.
- [38] H. Federer, *Geometric Measure Theory*. New York, NY: Springer-Verlag, 1969.
- [39] P. Feng and Y. Bresler, "Spectrum-blind minimum-rate sampling and reconstruction of multiband signals," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 3, 1996, pp. 1688–1691.
- [40] J. T. Schwartz, *Nonlinear Functional Analysis*. New York, NY: Gordon and Breach Science Publishers, 1969.
- [41] K. Do Ba, P. Indyk, E. Price, and D. P. Woodruff, "Lower bounds for sparse recovery," in *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, 2010, pp. 1190–1197.
- [42] C. E. Shannon, "Communication in the presence of noise," in *Proceedings of the IRE*, vol. 37, no. 1, Jan. 1949, pp. 10–21.
- [43] J. L. Massey, "Joint source channel coding," in *Communication Systems and Random Process Theory*, J. K. Skwirzynski, Ed. Sijthoff and Nordhoff.
- [44] R. Zamir, "A proof of the Fisher information inequality via a data processing argument," *IEEE Trans. Inf. Theory*, vol. 44, no. 3, pp. 1246–1250, May 1998.
- [45] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, United Kingdom: Cambridge University Press, 2004.
- [46] 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.
- [47] J. Ziv, "The behavior of analog communication systems," *IEEE Trans. Inf. Theory*, vol. 16, no. 5, pp. 587–594, 1970.
- [48] M. Gastpar, B. Rimoldi, and M. Vetterli, "To code, or not to code: Lossy source-channel communication revisited," *IEEE Trans. Inf. Theory*, vol. 49, no. 5, pp. 1147–1158, 2003.
- [49] T. Berger, *Rate Distortion Theory: A Mathematical Basis for Data Compression*. Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [50] U. Brehm, "Extensions of distance reducing mappings to piecewise congruent mappings on \mathbf{R}^m ," *Journal of Geometry*, vol. 16, no. 1, pp. 187–193, 1981.
- [51] A. V. Akopyan and A. S. Tarasov, "A constructive proof of Kirszbraun's theorem," *Matematicheskie Zametki*, vol. 84, no. 5, pp. 781–784, 2008.
- [52] T. Tanaka, "A statistical-mechanics approach to large-system analysis of CDMA multiuser detectors," *IEEE Trans. Inf. Theory*, vol. 48, pp. 2888–2910, Nov. 2002.
- [53] D. Guo and S. Verdú, "Randomly spread CDMA: Asymptotics via statistical physics," *IEEE Trans. Inf. Theory*, vol. 51, no. 6, pp. 1983–2010, Jun. 2005.
- [54] D. Guo, S. Shamai (Shitz), and S. Verdú, "Mutual Information and Minimum Mean-Square Error in Gaussian Channels," *IEEE Trans. Inf. Theory*, vol. 51, no. 4, pp. 1261 – 1283, Apr. 2005.
- [55] D. Guo and C. C. Wang, "Multiuser detection of sparsely spread cdma," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 3, pp. 421–431, Apr. 2008.
- [56] A. Montanari, "Estimating random variables from random sparse observations," *European Transactions on Telecommunications*, vol. 19, no. 4, pp. 385–403, 2008.
- [57] H. E. Stanley, *Introduction to phase transitions and critical phenomena*. London, United Kingdom: Oxford University Press, 1971.
- [58] S. Rangan, A. K. Fletcher, and V. K. Goyal, "Asymptotic analysis of map estimation via the replica method and applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1902–1923, Mar. 2012.
- [59] M. Bayati and A. Montanari, "The LASSO risk for Gaussian matrices," *IEEE Trans. Inf. Theory*, vol. 58, no. 4, pp. 1997–2017, Apr. 2012.
- [60] A. Maleki, L. Anitori, Z. Yang, and R. G. Baraniuk, "Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP)," submitted to *IEEE Trans. Inf. Theory*, 2011. [Online]. Available: <http://arxiv.org/abs/1108.0477>
- [61] 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.
- [62] D. L. Donoho, I. Johnstone, and A. Montanari, "Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising," submitted to *IEEE Trans. Inf. Theory*, 2011. [Online]. Available: <http://arxiv.org/abs/1111.1041>
- [63] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 764–785, 2011.
- [64] A. Montanari, "Graphical models concepts in compressed sensing," in *Compressed Sensing: Theory and Applications*. Cambridge Press, 2012, pp. 394–438.
- [65] 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.
- [66] 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.
- [67] P. Mattila, *Geometry of Sets and Measures in Euclidean Spaces: Fractals and Rectifiability*. Cambridge, United Kingdom: Cambridge University Press, 1999.
- [68] H. Steinhaus, "Sur les distances des points des ensembles de mesure positive," *Fundamenta Mathematicae*, vol. 1, pp. 93 – 104, 1920.
- [69] A. György, T. Linder, and K. Zeger, "On the rate-distortion function of random vectors and stationary sources with mixed distributions," *IEEE Trans. Inf. Theory*, vol. 45, pp. 2110 – 2115, Sep. 1999.
- [70] D. Guo, Y. Wu, S. Shamai (Shitz), and S. Verdú, "Estimation in Gaussian Noise: Properties of the Minimum Mean-square Error," *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 2371 – 2385, Apr. 2011.
- [71] D. L. Donoho, A. Javanmard, and A. Montanari, "Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing," submitted to *IEEE Trans. Inf. Theory*, Dec. 2011. [Online]. Available: <http://arxiv.org/abs/1112.0708>
- [72] S. Kudekar and H. D. Pfister, "The effect of spatial coupling on compressive sensing," in *Proceedings of the Forty-eighth Annual Allerton Conference on Communication, Control, and Computing*, 2010, pp. 347–353.
- [73] D. L. Donoho, I. Johnstone, A. Maleki, and A. Montanari, "Compressed sensing over ℓ_p -balls: Minimax mean square error," submitted to *Annals of Statistics*, 2011. [Online]. Available: <http://arxiv.org/abs/1103.1943>
- [74] D. L. Donoho, A. Maleki, and A. Montanari, "How to design message passing algorithms for compressed sensing," preprint, 2011. [Online]. Available: <http://www.ece.rice.edu/~mam15/bpist.pdf>
- [75] K. H. Lee and D. P. Petersen, "Optimal linear coding for vector channels," *IEEE Trans. Commun.*, vol. 24, no. 12, pp. 1283–1290, Dec. 1976.
- [76] 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.

PLACE
PHOTO
HERE

Yihong Wu (S'10) received the B.E. degree from Tsinghua University, Beijing, China, in 2006 and the M.A. and Ph.D. degrees from Princeton University, Princeton, NJ, in 2008 and 2011, all in electrical engineering. He is currently a postdoctoral fellow with the Statistics Department, The Wharton School, University of Pennsylvania, Philadelphia, PA.

His research interests are in information theory, statistical decision theory, high-dimensional inference, approximation theory and fractal geometry.

Dr. Wu was a recipient of the Princeton University Wallace Memorial Honorific Fellowship in 2010. He was a recipient of 2011 Marconi Society Paul Baran Young Scholar Award and the Best Student Paper Award at the 2011 IEEE International Symposiums on Information Theory (ISIT).

PLACE
PHOTO
HERE

Sergio Verdú (S'80–M'84–SM'88–F'93) is on the faculty of the School of Engineering and Applied Science at Princeton University.

A member of the National Academy of Engineering, Verdú is the recipient of the 2007 Claude E. Shannon Award and of the 2008 IEEE Richard W. Hamming Medal.