Almost Lossless Analog Signal Separation

David Stotz[1], Erwin Riegler[2], and Helmut Bölcskei[1]

[1]ETH Zurich, Switzerland
[2]Vienna University of Technology, Austria

Abstract—We propose an information-theoretic framework for analog signal separation. Specifically, we consider the problem of recovering two analog signals from a noiseless sum of linear measurements of the signals. Our framework is inspired by the groundbreaking work of Wu and Verdú (2010) on almost lossless analog compression. The main results of the present paper are a general achievability bound for the compression rate in the analog signal separation problem, an exact expression for the optimal compression rate in the case of signals that have mixed discrete-continuous distributions, and a new technique for showing that the intersection of generic subspaces with subsets of sufficiently small Minkowski dimension is empty. This technique can also be applied to obtain a simplified proof of a key result in Wu and Verdú (2010).

I. INTRODUCTION

We consider the following signal separation problem: Reconstruct the vectors $\boldsymbol{y}$ and $\boldsymbol{z}$ from the noiseless observation

$$\boldsymbol{w} = \boldsymbol{A}\boldsymbol{y} + \boldsymbol{B}\boldsymbol{z} \qquad (1)$$

where $\boldsymbol{A}$ and $\boldsymbol{B}$ are (measurement) matrices. As detailed in [1, Sec. 1] this problem has numerous applications such as inpainting, super-resolution, and the recovery of clipped signals and of signals that are corrupted by impulse noise or narrowband interference.

The sparse signal recovery literature [1]–[10] provides separation guarantees under sparsity constraints on the vectors $\boldsymbol{y}$ and $\boldsymbol{z}$. The sparsity thresholds in [1], [5], [9], [10] are functions of the coherence parameters [1] of the matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ and hold for *all* $\boldsymbol{y}$ and $\boldsymbol{z}$ meeting these thresholds, but suffer from the square-root bottleneck [6]. For random signals, the asymptotic results in [2], [3], [8] overcome the square-root bottleneck, but hold "only" with overwhelming probability. For $\boldsymbol{B}$ the identity matrix and $\boldsymbol{A}$ a random orthogonal matrix it is shown in [7] that the probability of failure of a certain reconstruction procedure decays exponentially in the dimension of the ambient space.

Contributions: Inspired by the recent work of Wu and Verdú [11], we derive asymptotic recovery results for the analog signal separation problem with the vectors $\boldsymbol{y}$ and $\boldsymbol{z}$ random, possibly dependent, and of general distributions. Our results hold for deterministic $\boldsymbol{B}$ and for almost all (a.a.) matrices $\boldsymbol{A}$, but do not depend on coherence parameters. However, since we assume $\boldsymbol{y}$ and $\boldsymbol{z}$ to be random, the statements are in terms of probability of separation error with respect to the source distributions, and hence do not provide worst-case guarantees like the coherence-based results in [1], [5], [9], [10].

Specifically, we study the asymptotic setting $\ell, n \to \infty$ where the vectors $\boldsymbol{y} \in \mathbb{R}^{n-\ell}$ and $\boldsymbol{z} \in \mathbb{R}^{\ell}$ are realizations of random processes; for each n, we let $\ell = \lfloor \lambda n \rfloor$ and $k = \lfloor Rn \rfloor$ for parameters $R, \lambda \in [0,1]$ and measurement matrices $\boldsymbol{A} \in \mathbb{R}^{k \times (n-\ell)}$ and $\boldsymbol{B} \in \mathbb{R}^{k \times \ell}$, with $k \geqslant \ell$. We refer to the parameter R as the compression rate as it equals (approximately) the ratio between the number of measurements and the total number of parameters in $\boldsymbol{y}$ and $\boldsymbol{z}$. In Theorem 1, we show that for each (deterministic) full-rank matrix $\boldsymbol{B}$, recovering $\boldsymbol{y}$ and $\boldsymbol{z}$ from the measurement $\boldsymbol{w}$ is possible with arbitrarily small probability of separation error for a.a. matrices $\boldsymbol{A}$, provided that n is sufficiently large and the compression rate R is larger than the Minkowski dimension compression rate (see Definition 4) of the concatenated random vector $[\boldsymbol{y}^{\mathrm{T}} \ \boldsymbol{z}^{\mathrm{T}}]^{\mathrm{T}}$. Since the technique used to prove the related result [11, Thm. 18] in the context of almost lossless analog compression can not be adapted to our setting, we develop a new proof method. The foundation of our approach, inspired by [12], is a new technique for showing that the intersection of generic subspaces with subsets of sufficiently small Minkowski dimension is empty (Proposition 1). A novel concentration of measure result, developed in Lemma 3, turns out to be an essential ingredient of this technique. Applying our method to the setting in [11] leads to a significant simplification of the proof of [11, Thm. 18, 1)].

For $\boldsymbol{y}$ and $\boldsymbol{z}$ mixed discrete-continuously distributed with mixing parameters ρ_1 and ρ_2, respectively, we show that the Minkowski dimension compression rate can be evaluated explicitly to

$$(1-\lambda)\rho_1 + \lambda\rho_2. \qquad (2)$$

What is more, this threshold is tight in the sense that there is a converse if the compression rate R is smaller than (2).

Notation: For a relation $\# \in \{<, >, \leqslant, \geqslant, =, \neq, \in, \notin\}$, we write $f(n) \overset{\cdot}{\#} g(n)$ if there exists an $N \in \mathbb{N}$ such that $f(n) \# g(n)$ holds for all $n > N$. Leb^n denotes the n-dimensional Lebesgue measure and $\mathcal{B}^{\otimes n}$ the Borel σ-algebra on $\mathbb{R}^n$. We write $\| \cdot \|$ for the ℓ_2-norm on $\mathbb{R}^n$. Matrices are denoted by capital boldface and vectors by lowercase boldface letters. $B^n(\boldsymbol{x}, \varepsilon)$ is the ball centered at $\boldsymbol{x} \in \mathbb{R}^n$ of radius ε with respect to $\| \cdot \|$, and its volume is $\alpha(n, \varepsilon) = \mathrm{Leb}^n(B^n(\boldsymbol{x}, \varepsilon))$. The closure of a set $\mathcal{U} \subset \mathbb{R}^n$ is denoted by $\overline{\mathcal{U}}$. We use sans-serif letters, e.g. x, for random quantities and roman letters, e.g. x, for deterministic quantities. For a random variable X or a random vector x, μ_{X} and μ_{x} denote the respective distribution. We write $\mathbf{1}_{\mathsf{X} \in \mathcal{A}}$ for the characteristic function associated with the event $\mathsf{X} \in \mathcal{A}$.

II. MAIN RESULT

We start by noting that (1) can be rewritten as

$$\boldsymbol{w} = [\boldsymbol{A}\ \boldsymbol{B}] \begin{bmatrix} \boldsymbol{y} \\ \boldsymbol{z} \end{bmatrix}$$

which shows that, formally, the separation problem we consider can be cast as an almost lossless analog compression problem [11] with measurement matrix $\boldsymbol{H} = [\boldsymbol{A}\ \boldsymbol{B}]$ and random source vector $[\mathbf{y}^{\mathrm{T}}\ \mathbf{z}^{\mathrm{T}}]^{\mathrm{T}}$, where $\mathbf{y}$ and $\mathbf{z}$ are possibly dependent. As we shall see below in Remark 3, the results in [11] can, however, not be applied to our setting, as $\boldsymbol{B}$ is deterministic here, whereas the results in [11] hold for a.a. matrices $\boldsymbol{H}$.

Definition 1: Let $0 \leqslant \lambda \leqslant 1$. Suppose that $(\mathsf{Y}_i)_{i \in \mathbb{N}}$ and $(\mathsf{Z}_i)_{i \in \mathbb{N}}$ are stochastic processes on $(\mathbb{R}^{\mathbb{N}}, \mathcal{B}^{\otimes \mathbb{N}})$. Then, for $n \in \mathbb{N}$, the source vector $\mathbf{x}$ of length n is given by $\mathbf{x} = [\mathsf{X}_1 \ldots \mathsf{X}_n]^{\mathrm{T}}$ with

$$\mathsf{X}_i = \mathsf{Y}_i, \quad \text{for } i \in \{1, \ldots, n-\ell\}$$
$$\mathsf{X}_{n-l+i} = \mathsf{Z}_i, \quad \text{for } i \in \{1, \ldots, \ell\}$$

where $\ell = \lfloor \lambda n \rfloor$.

The rate Definitions 2 and 4 below are adapted from the corresponding definitions in [11].

Definition 2: (Analog compression - *linear* measurements/ *measurable* separator). For $\mathbf{x}$ as in Definition 1 and $\varepsilon > 0$, an (n, k) code consists of

(i) linear measurements $[\boldsymbol{A}\ \boldsymbol{B}] : \mathbb{R}^{n-\ell} \times \mathbb{R}^{\ell} \to \mathbb{R}^k$;
(ii) a separator $g : \mathbb{R}^k \to \mathbb{R}^{n-\ell} \times \mathbb{R}^{\ell}$ that is measurable with respect to $\mathcal{B}^{\otimes k}$ and $\mathcal{B}^{\otimes n}$.

We call R with $0 \leqslant R \leqslant 1$ an ε-achievable rate if there exists a sequence of $(n, \lfloor Rn \rfloor)$ codes such that

$$\mathbb{P}[g([\boldsymbol{A}\ \boldsymbol{B}]\mathbf{x}) \neq \mathbf{x}] \stackrel{\cdot}{<} \varepsilon.$$

We define the optimal linear compression rate $R_{\mathrm{L}}(\varepsilon)$ as the infimum over all ε-achievable rates. Here, the name "linear" reflects the restriction to linear measurements, employed throughout the paper.

Next, we define the Minkowski dimension. This quantity is sometimes also referred to as box-counting dimension, which is the origin for the subscript B in the notation $\dim_{\mathrm{B}}(\cdot)$ used below.

Definition 3: (Minkowski dimension, [13]). Let $\mathcal{S}$ be a nonempty bounded set in $\mathbb{R}^n$. Define the lower and upper Minkowski dimension of $\mathcal{S}$ as

$$\underline{\dim}_{\mathrm{B}}(\mathcal{S}) = \liminf_{\varepsilon \to 0} \frac{\log N_{\mathcal{S}}(\varepsilon)}{\log \frac{1}{\varepsilon}}$$
$$\overline{\dim}_{\mathrm{B}}(\mathcal{S}) = \limsup_{\varepsilon \to 0} \frac{\log N_{\mathcal{S}}(\varepsilon)}{\log \frac{1}{\varepsilon}}$$

where $N_{\mathcal{S}}(\varepsilon)$ is the covering number of $\mathcal{S}$ given by

$$N_{\mathcal{S}}(\varepsilon) = \min \left\{ m \in \mathbb{N} \mid \mathcal{S} \subset \bigcup_{i \in \{1, \ldots, m\}} B^n(\boldsymbol{x}_i, \varepsilon),\ \boldsymbol{x}_i \in \mathbb{R}^n \right\}.$$

If $\underline{\dim}_{\mathrm{B}}(\mathcal{S}) = \overline{\dim}_{\mathrm{B}}(\mathcal{S})$, we simply write $\dim_{\mathrm{B}}(\mathcal{S})$.

Definition 4: (Minkowski dimension compression rate). For $\mathbf{x}$ from Definition 1 and $\varepsilon > 0$, we define the Minkowski dimension compression rate as

$$R_{\mathrm{B}}(\varepsilon) = \limsup_{n \to \infty} a_n(\varepsilon), \quad \text{where}$$
$$a_n(\varepsilon) = \inf \left\{ \frac{\underline{\dim}_{\mathrm{B}}(\mathcal{S})}{n} \ \middle|\ \mathcal{S} \subset \mathbb{R}^n,\ \mathbb{P}[\mathbf{x} \in \mathcal{S}] \geqslant 1 - \varepsilon \right\}. \tag{3}$$

Remark 1: Note that in (3) the infimum is taken with respect to the *lower* Minkowski dimension, whereas the corresponding definition in [11] is based on the *upper* Minkowski dimension. Our main result, Theorem 1 below, when specialized to the setting of [11], i.e., $\lambda = 0$, therefore constitutes an improvement of the general achievability result in [11].

The following theorem states that for *every* full-rank matrix $\boldsymbol{B} \in \mathbb{R}^{k \times \ell}$, with $k \geqslant \ell$, every rate R with $R > R_{\mathrm{B}}(\varepsilon)$ is ε-achievable for a.a. $\boldsymbol{A}$.

Theorem 1: Let $\mathbf{x}$ be as in Definition 1 with $\varepsilon > 0$ and let $R > R_{\mathrm{B}}(\varepsilon)$. Then, for every full-rank matrix $\boldsymbol{B} \in \mathbb{R}^{k \times \ell}$, with $k \geqslant \ell$, and for a.a. (with respect to $\mathrm{Leb}^{k(n-\ell)}$) matrices $\boldsymbol{A} \in \mathbb{R}^{k \times (n-\ell)}$, where $k = \lfloor Rn \rfloor$, there exists a measurable separator g such that

$$\mathbb{P}[g([\boldsymbol{A}\ \boldsymbol{B}]\mathbf{x}) \neq \mathbf{x}] \stackrel{\cdot}{<} \varepsilon. \tag{4}$$

Proof: See Section V. ∎

Remark 2: The proof of Theorem 1 reveals that the minimum $N \in \mathbb{N}$ for (4) to hold for all $n > N$ depends only on the distribution of $\mathbf{x}$ and is independent of the matrices $\boldsymbol{A}$ and $\boldsymbol{B}$.

Remark 3: In [11, Thm. 18, 1)] it was shown that every rate R with $R > R_{\mathrm{B}}(\varepsilon)$ is ε-achievable in almost lossless analog compression for a.a. measurement matrices $\boldsymbol{H} \in \mathbb{R}^{k \times n}$. This result is generalized in Theorem 1 above to hold for $\boldsymbol{H} = [\boldsymbol{A}\ \boldsymbol{B}]$ for a given full-rank matrix $\boldsymbol{B}$, with $k \geqslant \ell$, for a.a. matrices $\boldsymbol{A} \in \mathbb{R}^{k \times (n-\ell)}$. Since in a concrete separation problem we often encounter a particular structure for $\boldsymbol{B}$, for example a certain dictionary under which the corresponding signal is sparse, it is important to have the statement hold for *all* matrices $\boldsymbol{B}$, instead of only for a.a. $\boldsymbol{H} = [\boldsymbol{A}\ \boldsymbol{B}]$. The proof of [11, Thm. 18, 1)] relies on intricate properties of measures on Grassmanian manifolds that are invariant under the action of the orthogonal group. These arguments can not be applied to our setting as the overall measurement matrix $\boldsymbol{H} = [\boldsymbol{A}\ \boldsymbol{B}]$ has a deterministic $k \times \ell$ block $\boldsymbol{B}$. This forced us to find an alternative proof, which is based on two key elements, a concentration of measure result stated in Lemma 3, and a dimension counting argument provided in Proposition 1. The dimension counting argument says that the $(n-k)$-dimensional nullspace of $\boldsymbol{H}$ and the approximate support set $\mathcal{S}$ in (3) of the source vector $\mathbf{x}$ will not intersect, if the Minkowski dimension of $\mathcal{S}$ is smaller than k. Underlying this argument is the basic idea that two objects whose dimensions do not add up to at least the dimension of their ambient space, in general, do not intersect. Our proof strategy also applies to the compression problem [11] and leads to a significant

simplification of the proof of [11, Thm. 18, 1)], as detailed in Section VI.

III. MIXED DISCRETE-CONTINUOUS DISTRIBUTIONS

In order to establish the connection to the traditional sparse signal separation problem considered, e.g., in [1]–[3], [5], [7]–[10], we next consider sources x with independent components, where each component of the constituent processes $(Y_i)_{i \in \mathbb{N}}$ and $(Z_i)_{i \in \mathbb{N}}$ has a mixed discrete-continuous distribution, with possibly different mixture parameters for (Y_i) and (Z_i).

Definition 5: We say that x from Definition 1 has a mixed discrete-continuous distribution if for each $n \in \mathbb{N}$ the random variables X_i for $i \in \{1, \dots, n\}$ are independent and distributed according to

$$\mu_{X_i} = \begin{cases} (1 - \rho_1)\mu_{d_1} + \rho_1 \mu_{c_1}, & i \in \{1, \dots, n - \ell\} \\ (1 - \rho_2)\mu_{d_2} + \rho_2 \mu_{c_2}, & i \in \{n - \ell + 1, \dots, n\}, \end{cases}$$ (5)

where $0 \leqslant \rho_i \leqslant 1$, the μ_{c_i} are distributions on $(\mathbb{R}, \mathcal{B})$, absolutely continuous with respect to Lebesgue measure, and the μ_{d_i} are discrete distributions.

Lemma 1: Suppose that x is distributed according to Definition 5. Then

$$R_B(\varepsilon) = \lambda\rho_1 + (1 - \lambda)\rho_2$$ (6)

for all ε satisfying $0 < \varepsilon < 1$.

Proof: The proof is largely similar to the proof of [11, Thm. 15]. A sketch of the part that is different is provided in Section VII. ∎

Theorem 1 shows that the optimal linear compression rate $R_L(\varepsilon)$ is lower-bounded by the Minkowski dimension compression rate $R_B(\varepsilon)$. In the mixed discrete-continuous case we can strengthen this result through the following converse.

Lemma 2: Suppose that x is distributed according to Definition 5 and let $\varepsilon > 0$ and $R > R_B(\varepsilon)$. Then, for each full-rank matrix $B \in \mathbb{R}^{k \times \ell}$, with $k \geqslant \ell$, and for Lebesgue a.a. (with respect to $\mathrm{Leb}^{k(n-\ell)}$) matrices $A \in \mathbb{R}^{k \times (n-\ell)}$, where $k = \lfloor Rn \rfloor$, there exists a measurable separator g such that

$$\mathbb{P}[g([A \ B]x) \neq x] \, \dot{<} \, \varepsilon.$$ (7)

Moreover, for every ε with $0 < \varepsilon < 1$, $R \geqslant R_B(\varepsilon)$ is also a necessary condition for (7) to hold, i.e., $R_L(\varepsilon) = R_B(\varepsilon)$.

Proof: Achievability: Follows from Theorem 1.
Converse: In the same spirit as the proof of the converse part of [11, Thm. 6]. ∎

Finally, we combine Lemmata 1 and 2 to get an analytical expression for the optimal linear compression rate.

Corollary 1: Suppose that x has a mixed discrete-continuous distribution accoding to Definition 5 and let $0 < \varepsilon < 1$. Then, we have

$$R_L(\varepsilon) = (1 - \lambda)\rho_1 + \lambda\rho_2.$$ (8)

Corollary 1 essentially states that the optimal linear compression rate is determined by the fraction of continuously distributed components in the concatenated source vector.

Interestingly, $R_L(\varepsilon)$ does not depend on coherence quantities of the measurement matrices A and B, which usually arise in recovery thresholds in the sparse signal separation problem, see, e.g., [1], [8]. In this respect, under the rate constraint $R > R_B(\varepsilon)$, a.a. matrices A are "incoherent" to a given matrix B. When the distribution of one of the signals is purely discrete, the optimal linear compression rate is determined solely by the distribution of the other signal. Finally, if the dimension of one of the signals is much larger than the dimension of the other, i.e., $\lambda \approx 0$ or $\lambda \approx 1$, then the characteristics of the higher-dimensional signal dominate the threshold in Corollary 1.

IV. TECHNICAL RESULTS

In this section, we collect the main technical results referred to earlier in the paper. These results are important ingredients of the proof of Theorem 1, detailed in Section V, and the simplification of the proof of [11, Thm. 18, 1)], described in Section VI. First, we present a concentration result that bounds the probability that the norm of the image of a deterministic vector under a random affine mapping is small.

Lemma 3: Let $A = [a_1 \ \dots \ a_k]^T$ be a random matrix in $\mathbb{R}^{k \times n}$ where the a_i are i.i.d. uniform on the set $B^n(\mathbf{0}, r)$. Then, for each $u \in \mathbb{R}^n \backslash \{\mathbf{0}\}$, each $v \in \mathbb{R}^k$, and $\delta > 0$, we have

$$\mathbb{P}[\|Au + v\| < \delta] \leqslant C(n, k, r)\frac{\delta^k}{\|u\|^k},$$

where $C(n, k, r)$ is a constant that depends on n, k, and r only.

Proof:

$$\alpha(n, r)^k \, \mathbb{P}[\|Au + v\| < \delta]$$
$$= \mathrm{Leb}^{kn}\{A \in B^n(\mathbf{0}, r) \times \dots \times B^n(\mathbf{0}, r) \mid \|Au + v\| < \delta\}$$
$$\leqslant \prod_{i=1}^{k} \mathrm{Leb}^n\{a_i \in B^n(\mathbf{0}, r) \mid |a_i^T u + v_i| < \delta\}$$
$$\overset{(a)}{=} \prod_{i=1}^{k} \mathrm{Leb}^n\left\{a_i \in B^n(\mathbf{0}, r) \mid \left|a_i^T e_1 + \frac{v_i}{\|u\|}\right| < \frac{\delta}{\|u\|}\right\}$$
$$\overset{(b)}{\leqslant} (2r)^{k(n-1)} \prod_{i=1}^{k} \mathrm{Leb}^1\left\{a_i \in \mathbb{R} \mid \left|a_i + \frac{v_i}{\|u\|}\right| < \frac{\delta}{\|u\|}\right\}$$
$$= \frac{(2r)^{k(n-1)}(2\delta)^k}{\|u\|^k},$$

where (a) follows from the fact that Leb^n is invariant under rotations and we consider a rotation that takes $u/\|u\|$ into $e_1 = [1 \, 0 \, \dots \, 0]^T$, and in (b) we denote by a_i the first component of the vector a_i and use the fact that the magnitudes of the remaining components of a_i are less than or equal to r. ∎

Proposition 1: Let $\mathcal{S} \subset \mathbb{R}^n$ be such that $d := \underline{\dim}_B(\mathcal{S}) < k$. Then

$$\{u \in \mathcal{S}\backslash\{\mathbf{0}\} \mid Au = \mathbf{0}\} = \emptyset,$$ (9)

for Lebesgue a.a. $A \in \mathbb{R}^{k \times n}$.

Proof: Suppose that A is distributed as specified in Lemma 3. In order to show that the Lebesgue measure of

matrices $\boldsymbol{A}$ for which (9) does not hold is zero, it suffices to prove that

$$\mathbb{P}[\exists \boldsymbol{u} \in \mathcal{S} \backslash \{\mathbf{0}\} : \mathbf{A}\boldsymbol{u} = \mathbf{0}] = 0, \tag{10}$$

for $r > 0$. We employ a union bound argument to lower-bound the norm of vectors in $\mathcal{S} \backslash \{\mathbf{0}\}$:

$$\mathbb{P}[\exists \boldsymbol{u} \in \mathcal{S} \backslash \{0\} : \mathbf{A}\boldsymbol{u} = \mathbf{0}]$$
$$\leqslant \sum_{j=1}^{\infty} \mathbb{P}[\exists \boldsymbol{u} \in \mathcal{S} \backslash B^n(\mathbf{0}, 1/j) : \mathbf{A}\boldsymbol{u} = \mathbf{0}]. \tag{11}$$

This allows us to conclude that it suffices to prove (10) for sets $\mathcal{S}' \subset \mathcal{S}$ with $\min\{\|\boldsymbol{u}\| \mid \boldsymbol{u} \in \overline{\mathcal{S}'}\} > 0$, as this would show that each term in the series in (11) is zero. Using the definition of the Minkowski dimension (Definition 3) and the fact that $\mathcal{S}' \subset \mathcal{S}$ implies $d' := \underline{\dim}_{\mathrm{B}}(\mathcal{S}') \leqslant d$, we can find a sequence ε_m tending to zero such that

$$\frac{\log \mathcal{N}_{\mathcal{S}'}(\varepsilon_m)}{\log \frac{1}{\varepsilon_m}} \xrightarrow{m \to \infty} d'.$$

Let $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{\mathcal{N}_{\mathcal{S}'}(\varepsilon_m)}$ be the centers of the balls of radius ε_m that cover $\mathcal{S}'$ (cf. Definition 3). Since $\min\{\|\boldsymbol{u}\| \mid \boldsymbol{u} \in \overline{\mathcal{S}'}\} > 0$, we can assume m to be sufficiently large for $\min\{\|\boldsymbol{x}_i\| \mid 1 \leqslant i \leqslant \mathcal{N}_{\mathcal{S}'}(\varepsilon_m)\} > 0$ to hold. As the norm of each row of $\mathbf{A}$ is bounded, all realizations of $\mathbf{A}$ have a common Lipschitz constant, say L. Putting things together, we find that

$$\mathbb{P}[\exists \boldsymbol{u} \in \mathcal{S}' : \mathbf{A}\boldsymbol{u} = \mathbf{0}]$$
$$\overset{(a)}{\leqslant} \sum_{i=1}^{\mathcal{N}_{\mathcal{S}'}(\varepsilon_m)} \mathbb{P}[\exists \boldsymbol{u} \in B^n(\boldsymbol{x}_i, \varepsilon_m) : \mathbf{A}\boldsymbol{u} = \mathbf{0}]$$
$$\leqslant \sum_{i=1}^{\mathcal{N}_{\mathcal{S}'}(\varepsilon_m)} \mathbb{P}[\exists \boldsymbol{u} \in B^n(\boldsymbol{x}_i, \varepsilon_m) : \|\mathbf{A}\boldsymbol{u}\| < \varepsilon_m]$$
$$\overset{(b)}{\leqslant} \sum_{i=1}^{\mathcal{N}_{\mathcal{S}'}(\varepsilon_m)} \mathbb{P}[\|\mathbf{A}\boldsymbol{x}_i\| < (L+1)\varepsilon_m]$$
$$\overset{(c)}{\leqslant} C(n, k, r, L) \, \mathcal{N}_{\mathcal{S}'}(\varepsilon_m) \, \varepsilon_m^k \xrightarrow[(d)]{m \to \infty} 0,$$

where (a) follows from a union bound argument, (b) is a consequence of $\|\mathbf{A}\boldsymbol{x}_i\| \leqslant \|\mathbf{A}(\boldsymbol{x}_i - \boldsymbol{u})\| + \|\mathbf{A}\boldsymbol{u}\| \leqslant L\varepsilon_m + \|\mathbf{A}\boldsymbol{u}\|$, (c) is by application of Lemma 3, and (d) is a consequence of

$$\frac{\log \mathcal{N}_{\mathcal{S}'}(\varepsilon_m)}{\log \frac{1}{\varepsilon_m}} - k = \frac{\log\left(\mathcal{N}_{\mathcal{S}'}(\varepsilon_m)\varepsilon_m^k\right)}{\log \frac{1}{\varepsilon_m}} \xrightarrow{m \to \infty} d' - k < 0.$$

We have therefore shown that $\mathbb{P}[\exists \boldsymbol{u} \in \mathcal{S}' : \mathbf{A}\boldsymbol{u} = \mathbf{0}] = 0$. ∎

Remark 4: The result in Proposition 1 is very intuitive as it says that a generic $(n - k)$-dimensional subspace will intersect a d-dimensional object with $d < k$ at most trivially. A statement similar to Proposition 1 was proven in [12, Lem. 4.3]. The result in [12, Lem. 4.3] applies to linear combinations of Lipschitz mappings, and also gives an upper bound on the lower Minkowski dimension of the set on the left hand side of (9) when $d \geqslant k$. The proof of [12, Lem. 4.3] is based on the singular-value decomposition of $\boldsymbol{A}$. Our proof

above is more direct, but applies to $d < k$ only, the case relevant here.

Next, we provide a generalization of Proposition 1, which will be needed in the proof of Theorem 1.

Proposition 2: Let $\mathcal{S} \subset \mathbb{R}^n$ be such that $d := \underline{\dim}_{\mathrm{B}}(\mathcal{S}) < k$, and let $\boldsymbol{B} \in \mathbb{R}^{k \times \ell}$ be a matrix with $\operatorname{rank}(\boldsymbol{B}) = \ell$. Then,

$$\{\boldsymbol{u} \in \mathcal{S} \backslash \{\mathbf{0}\} \mid [\boldsymbol{A} \ \boldsymbol{B}]\boldsymbol{u} = \mathbf{0}\} = \emptyset,$$

for Lebesgue a.a. $\boldsymbol{A} \in \mathbb{R}^{k \times (n-\ell)}$.

The proof of Proposition 2 is similar to that of Proposition 1 above, and will therefore be omitted.

V. PROOF OF THEOREM 1

Since $R > R_{\mathrm{B}}(\varepsilon)$ and $k = \lfloor Rn \rfloor$, we have

$$a_n(\varepsilon) \overset{\cdot}{<} \frac{k}{n}, \tag{12}$$

which, together with the definition of $a_n(\varepsilon)$, implies that there exists a sequence[1] $\mathcal{U} := \mathcal{U}_n \subset \mathbb{R}^n$ such that

$$\underline{\dim}_{\mathrm{B}}(\mathcal{U}) \overset{\cdot}{<} k \tag{13}$$
$$\mathbb{P}[\mathbf{x} \in \mathcal{U}] \overset{\cdot}{\geqslant} 1 - \varepsilon. \tag{14}$$

For the remainder of the proof we choose n to be sufficiently large for (13) and (14) to hold in the #-sense. For $\boldsymbol{A} \in \mathbb{R}^{k \times (n-\ell)}$ and $\boldsymbol{B} \in \mathbb{R}^{k \times \ell}$ define the separator[2]

$$g(\boldsymbol{v}) = \begin{cases} \boldsymbol{x}, & \text{if } \{\boldsymbol{u} \mid [\boldsymbol{A} \ \boldsymbol{B}]\boldsymbol{u} = \boldsymbol{v}\} \cap \mathcal{U} = \{\boldsymbol{x}\} \\ \text{error}, & \text{else.} \end{cases} \tag{15}$$

Then

$$\mathbb{P}[g([\boldsymbol{A} \ \boldsymbol{B}]\mathbf{x}) \neq \mathbf{x}]$$
$$= \mathbb{P}[g([\boldsymbol{A} \ \boldsymbol{B}]\mathbf{x}) \neq \mathbf{x}, \mathbf{x} \in \mathcal{U}]$$
$$\qquad + \mathbb{P}[g([\boldsymbol{A} \ \boldsymbol{B}]\mathbf{x}) \neq \mathbf{x}, \mathbf{x} \notin \mathcal{U}] \tag{16}$$
$$\overset{(a)}{\leqslant} \mathbb{P}[g([\boldsymbol{A} \ \boldsymbol{B}]\mathbf{x}) = \text{error}, \mathbf{x} \in \mathcal{U}] + \varepsilon$$
$$\overset{(b)}{\leqslant} \mathbb{P}[\exists \boldsymbol{u} \in \mathcal{U}_{\mathbf{x}} \backslash \{\mathbf{0}\} : [\boldsymbol{A} \ \boldsymbol{B}]\boldsymbol{u} = \mathbf{0}, \mathbf{x} \in \mathcal{U}] + \varepsilon,$$

where $\mathcal{U}_{\mathbf{x}} = \mathcal{U} - \{\mathbf{x}\} = \{\boldsymbol{u} - \mathbf{x} \mid \boldsymbol{u} \in \mathcal{U}\}$, (a) follows from the definition of the separator (15) and from (14), and (b) again is by definition of the separator (15). Since $\underline{\dim}_{\mathrm{B}}(\mathcal{U}) = \underline{\dim}_{\mathrm{B}}(\mathcal{U}_{\boldsymbol{x}})$, we find, through application of Proposition 2, that

$$\operatorname{Leb}^{k(n-\ell)}\{\boldsymbol{A} \mid \exists \boldsymbol{u} \in \mathcal{U}_{\boldsymbol{x}} \backslash \{\mathbf{0}\} : [\boldsymbol{A} \ \boldsymbol{B}]\boldsymbol{u} = \mathbf{0}\} = 0, \tag{17}$$

for all $\boldsymbol{x}$. Therefore, the integral of (17) with respect to $\mu_{\mathbf{x}}(d\boldsymbol{x})$ is zero, and, noting that (17) can be written as an integral with respect to $d\boldsymbol{A}$, we can apply Fubini's Theorem to interchange the two integrals and obtain

$$\int_{\mathbb{R}^{k \times (n-\ell)}} \mathbb{P}[\exists \boldsymbol{u} \in \mathcal{U}_{\mathbf{x}} \backslash \{\mathbf{0}\} : [\boldsymbol{A} \ \boldsymbol{B}]\boldsymbol{u} = \mathbf{0}, \mathbf{x} \in \mathcal{U}]d\boldsymbol{A} = 0. \tag{18}$$

---

[1]The definition of $\mathcal{U}$ is to be understood in the sense that the sequence index n is dropped for simplicity of exposition.

[2]Taking "error" to be an arbitrary element of $\mathbb{R}^n \backslash \mathcal{U}$ we obtain a measurable map $g \colon \mathbb{R}^k \to \mathbb{R}^{n-\ell} \times \mathbb{R}^\ell$ as required in Definition 2.

Therefore, we have $\mathbb{P}[\exists \boldsymbol{u} \in \mathcal{U}_{\mathbf{x}} \setminus \{\mathbf{0}\} : [\boldsymbol{A}\ \boldsymbol{B}]\boldsymbol{u} = \mathbf{0}, \mathbf{x} \in \mathcal{U}] = 0$ for a.a. $\boldsymbol{A} \in \mathbb{R}^{k \times (n-\ell)}$. In summary, we have shown that

$$\mathbb{P}[g([\boldsymbol{A}\ \boldsymbol{B}]\mathbf{x}) \neq \mathbf{x}] \leqslant \varepsilon,$$

for a.a. $\boldsymbol{A}$, which completes the proof. ∎

VI. SIMPLIFYING THE PROOF OF [11, THM. 18, 1)]

In this section, we sketch how the technique developed in the proof of Proposition 1 can be applied to devise a simplified and elementary proof of [11, Thm. 18, 1)]. The framework of almost lossless analog compression in [11] for the case of linear measurements and a measurable decoder considers a general stochastic source process $\mathbf{x}$. The problem is to reconstruct $\mathbf{x}$ from $\boldsymbol{H}\mathbf{x}$, where $\boldsymbol{H}$ is the measurement matrix. The result in [11, Thm. 18, 1)] says that for $R > R_{\mathrm{B}}(\varepsilon)$, for a.a. $\boldsymbol{H} \in \mathbb{R}^{k \times n}$, there exists a measurable decoder g such that

$$\mathbb{P}[g(\boldsymbol{H}\mathbf{x}) \neq \mathbf{x}] \dot{<} \varepsilon,$$

where $k = \lfloor Rn \rfloor$.

Using Proposition 1, we can give an alternative, simplified proof of this result as follows. We choose a set $\mathcal{U} \subset \mathbb{R}^n$ such that (13) and (14) hold, and define the decoder according to

$$g(\boldsymbol{v}) = \begin{cases} \boldsymbol{x}, & \text{if } \{\boldsymbol{u} \mid \boldsymbol{H}\boldsymbol{u} = \boldsymbol{v}\} \cap \mathcal{U} = \{\boldsymbol{x}\} \\ \text{error}, & \text{else.} \end{cases} \quad (19)$$

The probability of a decoding error is then decomposed as in (16). Applying Proposition 1 we find that a.a. matrices $\boldsymbol{H}$ are injective on $\mathcal{U}$. Finally, invoking Fubini's Theorem as in the argument leading to (18) allows us to conclude that the probability of decoding error is zero when $\mathbf{x} \in \mathcal{U}$, leaving the total probability of decoding error to be smaller than ε and thus finishing the proof.

VII. SKETCH OF THE PROOF OF LEMMA 1

Recall the role of λ in Definition 1. The cases $\lambda \in \{0, 1\}$ are equivalent to the case $\lambda = 1/2$, $\rho_1 = \rho_2$, $\mu_{\mathrm{d}_1} = \mu_{\mathrm{d}_2}$, and $\mu_{\mathrm{c}_1} = \mu_{\mathrm{c}_2}$. Hence, we can assume, without loss of generality, that $0 < \lambda < 1$. This implies that we can take $\ell = \lfloor \lambda n \rfloor \notin \{0, n\}$.

Let $\mathcal{A}_i$ be the set of atoms of μ_{d_i}. Then

$$\mathbb{E}[\mathbf{1}_{X_j \notin \mathcal{A}_i}] = \mu_{X_j}(\mathcal{A}_i^c)$$
$$= \begin{cases} \rho_1, & \text{for } i = 1,\ j \in \{1, \dots, n-\ell\} \\ \rho_2, & \text{for } i = 2,\ j \in \{n-\ell+1, \dots, n\}. \end{cases}$$

By the weak law of large numbers,

$$\frac{1}{n-\ell} \sum_{j=1}^{n-\ell} \mathbf{1}_{X_j \notin \mathcal{A}_1} \xrightarrow{\mathbb{P}} \rho_1$$
$$\frac{1}{\ell} \sum_{j=n-\ell+1}^{n} \mathbf{1}_{X_j \notin \mathcal{A}_2} \xrightarrow{\mathbb{P}} \rho_2,$$

which yields

$$\frac{|\operatorname{spt}(\mathbf{x})|}{n} = \frac{n-\ell}{n} \frac{1}{n-\ell} \sum_{j=1}^{n-\ell} \mathbf{1}_{X_j \notin \mathcal{A}_1} + \frac{\ell}{n} \frac{1}{\ell} \sum_{j=n-\ell+1}^{n} \mathbf{1}_{X_j \notin \mathcal{A}_2}$$
$$\xrightarrow{\mathbb{P}} (1-\lambda)\rho_1 + \lambda\rho_2 \quad (20)$$

with the generalized support

$$\operatorname{spt}(\boldsymbol{x}) = \{i \in \{1, \dots, n-\ell\} \mid x_i \notin \mathcal{A}_1\}$$
$$\cup \{i \in \{n-\ell+1, \dots, n\} \mid x_i \notin \mathcal{A}_2\}.$$

Let $\kappa > 0$ be arbitrary and set

$$\mathcal{C} := \{\boldsymbol{x} \mid |\operatorname{spt}(\boldsymbol{x})| < ((1-\lambda)\rho_1 + \lambda\rho_2 + \kappa)n\}$$
$$\mathcal{D} := \{\boldsymbol{x} \mid |\operatorname{spt}(\boldsymbol{x})| > ((1-\lambda)\rho_1 + \lambda\rho_2 - \kappa)n\}.$$

Then, by convergence in probability in (20), we have

$$\mathbb{P}[\mathbf{x} \in \mathcal{C}] \dot{\geqslant} 1 - \varepsilon \quad (21)$$
$$\mathbb{P}[\mathbf{x} \in \mathcal{D}] \dot{\geqslant} 1 - \varepsilon. \quad (22)$$

The remaining steps of the proof are almost identical to the proof of [11, Thm. 15] and are therefore omitted. The idea is to decompose $\mathcal{C}$ and $\mathcal{D}$ into basic subsets, whose elements have certain components equal elements of the atomic sets $\mathcal{A}_1, \mathcal{A}_2$, and the remaining components arbitrary. This allows us to bound the Minkowski dimension of $\mathcal{C}$ and $\mathcal{T} \cap \mathcal{D}$, for arbitrary $\mathcal{T}$ with $\mathbb{P}[\mathbf{x} \in \mathcal{T}] \geqslant 1 - \varepsilon$, and thus to sandwich the Minkowski dimension compression rate according to

$$(1-\lambda)\rho_1 + \lambda\rho_2 - \kappa \leqslant R_{\mathrm{B}}(\varepsilon) \leqslant (1-\lambda)\rho_1 + \lambda\rho_2 + \kappa,$$

which yields the claim, since κ is arbitrary.

REFERENCES

[1] C. Studer, P. Kuppinger, G. Pope, and H. Bölcskei, "Recovery of sparsely corrupted signals," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3115–3130, May 2012.

[2] X. Li, "Compressed sensing and matrix completion with constant proportion of corruptions," *Constructive Approximation*, vol. 37, no. 1, pp. 79–99, Feb. 2013.

[3] J. Wright and Y. Ma, "Dense error correction via ℓ^1-minimization," *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3540–3560, Jul. 2010.

[4] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. Natl. Acad. Sci.*, vol. 106, no. 45, pp. 18 914–18 919, 2009.

[5] D. L. Donoho and G. Kutyniok, "Microlocal analysis of the geometric separation problem," *Communications on Pure and Applied Mathematics*, vol. 66, no. 1, pp. 1–47, Jan. 2013.

[6] J. A. Tropp, "On the conditioning of random subdictionaries," *Applied and Computational Harmonic Analysis*, vol. 25, pp. 1–24, 2008.

[7] M. B. McCoy and J. A. Tropp, "Sharp recovery bounds for convex deconvolution, with applications," *arXiv:1205.1580 [cs.IT]*, May 2012.

[8] G. Pope, A. Bracher, and C. Studer, "Probabilistic recovery guarantees for sparsely corrupted signals," *IEEE Trans. Inf. Theory*, vol. 59, no. 5, pp. 3104–3116, May 2013.

[9] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, Nov. 2001.

[10] D. L. Donoho and P. B. Stark, "Uncertainty principles and signal recovery," *SIAM Journal on Applied Mathematics*, vol. 49, no. 3, pp. 906–931, Jun. 1989.

[11] 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.

[12] T. Sauer, J. A. Yorke, and M. Casdagli, "Embedology," *Journal of Statistical Physics*, vol. 65, no. 3-4, pp. 579–616, Nov. 1991.

[13] K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, 2nd ed. John Wiley & Sons, 2004.