Signal Processing 120 (2016) 71–87

Contents lists available at ScienceDirect

Signal Processing

journal homepage: www.elsevier.com/locate/sigpro

Compressive measurements generated by structurally random matrices: Asymptotic normality and quantization

Raziel Haimi-Cohen [a,*], Yenming Mark Lai [b]

[a] Alcatel-Lwucent Bell-Laboratories, Murray Hill, NJ 07974, USA
[b] University of Texas, Austin, TX 78712, USA

ARTICLE INFO

Article history:
Received 27 January 2015
Received in revised form
31 May 2015
Accepted 24 July 2015
Available online 15 August 2015

Keywords:
Compressed sensing
Quantization
Structurally Random Matrices

ABSTRACT

Structurally random matrices (SRMs) are a practical alternative to fully random matrices (FRMs) when generating compressive sensing measurements because of their computational efficiency and their universality with respect to the sparsifying basis. In this work we derive the statistical distribution of compressive measurements generated by various types of SRMs, as a function of the signal properties. We show that under a wide range of conditions, that distribution is a mixture of asymptotically multi-variate normal components. We point out the implications for quantization and coding of the measurements and discuss design considerations for measurements transmission systems. Simulations on real-world video signals confirm the theoretical findings and show that the signal randomization of SRMs yields a dramatic improvement in quantization properties.

1. Introduction

Compressive sensing [1] is concerned with determining a signal $\mathbf{x} \in \mathbb{R}^n$ from a vector of measurements

$$\mathbf{y} = \Phi \mathbf{x} \quad (1)$$

where $\Phi \in \mathbb{R}^{m \times n}$, $m \ll n$, is a *sensing matrix*, and $\mathbf{x}$ is k-sparse representation in the column space of a *sparsifyer* Ψ,

$$\mathbf{x} = \Psi \zeta, \quad \|\zeta\|_0 \le k, \quad (2)$$

where Ψ is an orthogonal or a tight frame matrix and $\|\zeta\|_0$ denotes the number of non-zero entries in ζ. If $\Phi\Psi$ meets certain conditions, ζ and hence $\mathbf{x}$ can be reconstructed from $\mathbf{y}$ by solving the constrained minimization problem

$$\min \|\zeta\|_1 \quad \text{s.t.} \quad \mathbf{y} = \Phi\Psi\zeta \quad (3)$$

Other results in the same vein extend these results to compressible signals (signals which can be approximated by sparse signals), or provide error bounds on the reconstructed solution when the measurements contain noise. (In this case (3) may also be modified to account for the noise.)

1.1. Sensing matrix design

Various design methods attempt to generate a sensing matrix Φ that enables correct reconstruction of $\mathbf{x}$ from a small number of measurements in a computationally efficient way. Generally this goal is achieved only with very high probability (w.h.p.): either Φ is a random matrix and w.h.p., the selected instance of Φ enables correct and efficient reconstruction of every possible $(\mathbf{x}, \zeta)$ pair which satisfy (2); or $\mathbf{x}$ and ζ are random signals which satisfy (2) and Φ is deterministic such that the pair $(\mathbf{x}, \zeta)$ can be reconstructed efficiently w.h.p. [2]. In this paper, we are interested in the first option.

* Corresponding author at: Alcatel-Lucent, Bell Laboratories Rm. 2B-221, 660 Mountain Ave., Murray Hill, NJ 07974, USA. Tel.: +1 908 582 4159.
E-mail addresses: razi@alcatel-lucent.com (R. Haimi-Cohen), mlai@ices.utexas.edu (Y. Mark Lai).

http://dx.doi.org/10.1016/j.sigpro.2015.07.023
0165-1684/

A *fully random matrix* (FRM) is a matrix whose entries are independent, identically distributed (IID) Gaussian or Bernoulli random variables (RVs) [3,4]. If $m \geq O(k \log(n/k))$, then for any given Ψ, w.h.p., $\Phi\Psi$ is such that every $\mathbf{x}$ and ξ which satisfy (2) can be reconstructed by solving (3). FRMs are *universal*, that is, the design of Φ is independent of Ψ, hence the choice of sparsifier can be deferred to the reconstruction stage, which is of significant practical importance. However, because of their completely unstructured nature, FRMs are computationally unwieldy in large scale applications since the random matrix needs to be both computed and stored.

Randomly sampled transforms (RST) address the computational complexity problem by imposing structural constraints on the randomness. Let

$$\Phi = \sqrt{n/m}SW$$

where $W \in \mathbb{R}^{n \times n}$ is a square, orthonormal matrix having a fast transform, and $S \in \mathbb{R}^{m \times n}$ is a *random entries selection matrix*, that is, a matrix whose rows are selected randomly, with uniform distribution, from the rows of I_n, the $n \times n$ identity matrix. $\Phi\mathbf{x}$ can then be computed efficiently by calculating the fast transform $W\mathbf{x}$ and selecting a random subset of the transform coefficients. RSTs guarantee a correct solution, w.h.p., if

$$m \geq O(\mu^2(W, \Psi)k \log n) \quad (4)$$

where $\mu(W, \Psi)$, the mutual coherence of $W$and Ψ, is

$$\mu(W, \Psi) \triangleq \sqrt{n} \max_{1 \leq i \leq m,\ 1 \leq j \leq n} |\mathbf{w}_i \boldsymbol{\psi}_j| / (\|\mathbf{w}_i\|_2 \|\boldsymbol{\psi}_j\|_2)$$

where $\mathbf{w}_i$, $\boldsymbol{\psi}_j$ are the ith row and jth column of W, Ψ, respectively [5]. Since $1 \leq \mu(W, \Psi) \leq \sqrt{n}$, we can choose $m \ll n$ which satisfies (4) only if W is selected so that $\mu(W, \Psi)$ is small. Therefore, RSTs are not universal.

The universality issue was addressed by the introduction of *structurally random matrices* (SRM) [6,7]:

$$\Phi = \sqrt{n/m}SWR \quad (5)$$

where S, W are as above and $R \in \mathbb{R}^{n \times n}$, the *randomizer*, is a random matrix. Hence

$$\Phi\mathbf{x} = \sqrt{n/m}SW(R\mathbf{x}) = \sqrt{n/m}SW(R\Psi)\xi.$$

Therefore, a SRM with a given sparsifier Ψ behaves as the RST $\sqrt{n/m}SW$ with the random sparsifier $R\Psi$. If $R\Psi$ and W are mutually incoherent w.h.p., then SRMs are universal, and the known results for RSTs with incoherent sparsifiers (e.g. performance with compressible signals or noisy measurements) hold w.h.p.

Two types of randomization were proposed: *Local randomization* (LR), where each entry of $\mathbf{x}$ is multiplied by ± 1 with equal probability; and *global randomization* (GR), where the entries of $\mathbf{x}$ are randomly shuffled. Both forms are computationally simple and were shown, for a large class of transforms W, to be universal [7]. The universality of LR was extended in [9] to the more general case where $\sqrt{n/m}SW$ in (5) is replaced by any matrix with the restricted isometry property (RIP).

Other methods were also proposed for constructing universal and computationally efficient sensing matrices, such as Random convolution (RC) [8].

1.2. Quantization and coding of measurements

Many application of compressive sensing, e.g. video surveillance and streaming [10–16] involve sending the measurements for processing over a communication channel. The transmission of measurements requires a *coding scheme*, which entails source coding that is typically implemented by quantization followed by channel coding of the quantization codewords.

Conventional media coding standards are efficient over a wide range of input signals and operating conditions. One of the keys to this robustness is the usage of various signal-adaptive techniques in order to control the bit rate and improve performance. These techniques are applied before, during, and after quantization. For example, a linear prediction [17,18] model may be estimated for the signal and the quantization may be performed on the prediction error, which reduces the bit rates needed to achieve specific quantization accuracies; the granularity of the quantizer may be varied according the signal content; and one out of several possible variable length coding schemes may be selected to achieve low rate lossless coding of the quantization codewords. The parameters of the linear prediction model, the quantizer granularity, and the lossless coding scheme need to be shared with the decoder, and hence they are encoded and sent as side information. Since the side information is critical for the decoding of the signal as a whole, it is typically encoded with higher accuracy and, in noisy channels, with better error protection, than the rest of the data. The amount of data in the side information is very small, hence the bit rate overhead caused by sending it is usually negligible in comparison to the performance achieved by it.

The preferred coding scheme for compressive measurements depends on a variety of factors, but it is invariably based on assumptions about the probability distribution of the measurements, which is determined by the type of sensing matrix used. Furthermore, applying any of the signal-adaptive techniques described above requires having a parametric model where this distribution is characterized by parameters estimated from the signal and transmitted to the decoder as side information.

The quantization of compressive measurements has recently received significant attention. Dai et al. [19,20] studied the effect of quantization on reconstruction accuracy with various quantizer designs and provided asymptotic boundaries on the rate-distortion function when quantization is followed by Huffman coding [20]. The efficacy of uniform vs. non-uniform scalar measurement quantization was compared specifically for video signals in [14,21]. Unlike all other quantizer designs we reviewed, the quantizer of Venkatraman and Makur [14] is signal-adaptive: its operation is controlled by the variance of the measurements in each frame, which is sent to the decoder as side information. A quantizer optimized for compressed sensing reconstruction is presented in [22]. Laska et al. studied the effect of saturation [23], the trade-off between the number of measurements and quantization accuracy [24] and the extreme case of 1-bit quantizers [24–26]. Modifications to the reconstruction algorithms to address quantization effects were proposed in [23,26,27]. In all