

Average recovery performances of non-perfectly informed compressed sensing: With applications to multiclass encryption

Original

Average recovery performances of non-perfectly informed compressed sensing: With applications to multiclass encryption / Cambareri, Valerio; Mangia, Mauro; Pareschi, Fabio; Rovatti, Riccardo; Setti, Gianluca. - STAMPA. - 2015-:(2015), pp. 3651-3655. (40th IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2015 South Brisbane, QLD, Australia April 19-24, 2015) [10.1109/ICASSP.2015.7178652].

Availability:

This version is available at: 11583/2696673 since: 2022-02-08T17:11:57Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/ICASSP.2015.7178652

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

AVERAGE RECOVERY PERFORMANCES OF NON-PERFECTLY INFORMED COMPRESSED SENSING: WITH APPLICATIONS TO MULTICLASS ENCRYPTION

Valerio Cambareri^{1,2}, Mauro Mangia^{1,2}, Fabio Pareschi^{2,3}, Riccardo Rovatti^{1,2}, Gianluca Setti^{2,3}

¹ Department of Electrical, Electronic and Information Engineering, University of Bologna, Italy

² Advanced Research Center on Electronic Systems, University of Bologna, Italy

³ Engineering Department, University of Ferrara, Italy

ABSTRACT

The sensitivity of recovery algorithms with respect to a perfect knowledge of the encoding matrix is a general issue in many application scenarios in which compressed sensing is an option to acquire or encode natural signals. Quantifying this sensitivity in order to predict the result of signal recovery is therefore valuable when no *a priori* information can be exploited, e.g., when the encoding matrix is randomly perturbed without any exploitable structure. We tackle this aspect by means of a simplified model for the signal recovery problem, which enables the derivation of an average performance estimate that depends only on the interaction between the sensing and perturbation matrices.

The effectiveness of the resulting heuristic is demonstrated by numerical exploration of signal recovery under three simple perturbation matrix models. Finally, we show how this estimate matches very well the degradation experienced by non-perfectly informed decoders in applications of compressed sensing to protecting the acquired information content in ECG tracks and sensitive images.

Index Terms— Compressed Sensing, Matrix Perturbations, Average Recovery Performances, Encryption

1. INTRODUCTION

Compressed Sensing (CS) [1] is an acquisition technique in which the original signal is encoded into a set of *measurements* by means of a dimensionality-reducing linear transformation. These measurements can then be fed into a recovery algorithm (or *decoder*) that, by exploiting suitable prior information, is able to recover the signal in its original, high-dimensional domain. For this procedure to work, however, the decoder requires very accurate knowledge of the linear mapping used in the encoder i.e., an *encoding matrix* representing the above transformation.

In some applications such knowledge may be imperfect, and perturbations often arise either due to the nature of the physical mechanism by which the encoding matrix is applied [2,3] or to intentionally missing information at the decoder, i.e., when the decoder only knows the encoding matrix up to a certain degree of accuracy. In particular, while calibration [4] may be attempted in the first case, the latter case can be designed to minimize the possibility of recovering missing information by systematically using randomly varying encoding matrices and perturbations [5]. Fundamental results exist [6] extending the established theoretical signal recovery guarantees [7] to such perturbed cases. Yet, as it often happens with worst-case analyses based on the restricted isometry property [8], the resulting performance bounds are quite far from the average recovery quality attained by decoding algorithms.

In this paper, we propose a heuristic that allows a prediction of the typical recovery quality of a decoder that is subject to imperfect encoding matrix knowledge represented by an additive perturbation; while not general, the estimate effectively applies to typical configurations of CS-based acquisition systems. In particular, we will apply it to predicting the effect of some random perturbation models in the case of non-perfectly informed CS.

Moreover, we illustrate the use of the developed estimate as a design tool for some encryption schemes based on CS [5,9] that allow the embedding of some security properties directly in the acquisition process, with the only condition that each encoding matrix is used in a one-time pad fashion [10,11]. In fact, low-cost encryption of analog signal sources is obtained as follows: if the encoder transmits its measurements to the decoder, a prior agreement on a *private key* must exist so that they are both aligned on an equal, very long sequence of (pseudo)random encoding matrices.

In a slightly more sophisticated version of this scheme, decoders knowing the true encoding matrix with no errors (generated from a *first-class* key) are able to retrieve the original signal at full quality, while *second-class* decoders are given a (pseudo)randomly perturbed version of the encoding matrix (generated from a second-class key) and their recovery is therefore of limited quality. This enables different access levels to the protected content, as in other communication protocols such as global positioning data or digital multimedia broadcasting.

Since the difference in quality between users of different classes depends on how much the true encoding matrix differs with respect to (w.r.t.) its perturbed version provided by the second-class key, in general larger amounts of perturbation hint at gracefully decreasing recovery quality; yet, a more quantitative analysis is needed for a proper design of multiclass encryption. If the aforementioned recovery guarantees are modified to account for perturbations [6], the obtained error bounds are quite pessimistic w.r.t. actual decoding performances, therefore being of limited help when the errors are not an undesired effect to counter, but a design parameter to tune against typical decoding performances. Instead, we will see that the estimate we develop can be used to effectively anticipate quality losses due to missing information at second-class decoders, and therefore to design multiclass encryption schemes complying with recovery quality specifications depending on the users' class.

2. A BRIEF REVIEW OF COMPRESSED SENSING

We here consider that a signal x is represented by n Nyquist-rate samples collected in a vector $x = (x_0 \ \cdots \ x_{n-1})^\top \in \mathbb{R}^n$. The fundamental assumption of CS is that x is *sparse*, i.e., there exists an n -dimensional *sparsity basis* D (here considered orthonormal)

such that for any instance x we have $x = Ds$, with s having at most $k \ll n$ non-zero components. Due to this hypothesis, x can be recovered from a set of $m < n$ properly designed measurements (with the minimum value of $m = \mathcal{O}(k \log n/k)$) collected as $y = Ax = ADs$, with $A \in \mathbb{R}^{m \times n}$ the encoding matrix.

Remarkably, one can formally guarantee that x can be recovered from y even in the presence of noise [7] and despite the fact that A causes a dimensionality reduction. For this to occur, s must be sufficiently sparse so that the linear mapping AD acts as an approximate isometry w.r.t. all signals of sparsity k , i.e., the singularity of the mapping is mainly due to the way it transforms dense vectors, while the distances between sparse vectors are approximately preserved into their measurements.

If this holds, recovery is possible by enforcing the *a priori* knowledge that s is sparse. Sparse signal recovery has been a very active research field in recent years [12–15]. Many recovery algorithms solve convex optimization problems such as *basis pursuit*

$$\hat{s} = \arg \min_{\xi \in \mathbb{R}^n} \|\xi\|_1 \quad \text{s.t. } y = AD\xi \quad (1)$$

(denoted as $\text{BP}(y, A, D)$ to highlight its prior information) or *basis pursuit with denoising*,

$$\hat{s} = \arg \min_{\xi \in \mathbb{R}^n} \|\xi\|_1 \quad \text{s.t. } \|AD\xi - y\|_2 \leq \varepsilon \quad (2)$$

(denoted as $\text{BPDN}(y, A, D, \varepsilon)$) where $\|\cdot\|_1$ is the ℓ_1 norm, $\|\cdot\|_2$ is the usual ℓ_2 norm, and $\varepsilon \geq 0$ controls the fidelity with which noisy measurements are matched. This said, much of the practical interest in CS comes from the fact that the above requirements on AD are satisfied when A is a subgaussian random matrix [8] with independent and identically distributed (i.i.d.) entries, and the formal guarantees on the solution of (1) or (2) are largely outperformed by practical signal recovery performances.

3. AN AVERAGE PERFORMANCE ESTIMATE IN THE PRESENCE OF PERTURBATIONS

In many applications the encoding matrix can be factored as $A = A' + \Delta A$ where A' is known to the decoder, while ΔA is a perturbation matrix so that the second term of $y = A'x + \Delta Ax$ is signal-dependent noise (ΔA being unavailable to the decoder). Such a decoder may either be *naive* and solve $\text{BP}(y, A', D)$, or attempt denoising by $\text{BPDN}(y, A', D, \varepsilon)$, albeit with an unknown ε that has to be chosen carefully. The relative sophistication of such convex optimization problems prevents an average analysis of the sensitivity w.r.t. the perturbation matrix in typical recovery problems. For this reason, in our simplified model we assume that (i) $(A', \Delta A)$ are random matrices with known distributions of entries, (ii) an approximation $x' = Ds'$ is obtained by solving $\text{BP}(y, A', D)$ to satisfy $y = A'x'$. Pairing this with the original $y = Ax$, as $\Delta A = A - A'$ we obtain $A'\Delta x = \Delta Ax$ where $\Delta x = x' - x$. Starting from this, we further assume that ΔA is indeed a *perturbation* (its entity is small w.r.t. A') so that the approximation error Δx is small in the least-squares sense, i.e.,

$$\Delta x = \arg \min_{\Delta \zeta \in \mathbb{R}^n} \|\Delta \zeta\|_2^2 \quad \text{s.t. } A'\Delta \zeta = \Delta Ax$$

whose solution is $\Delta x = (A')^+ \Delta Ax$ with $^+$ denoting the Moore-Penrose pseudo-inverse. To investigate the ratio between the energies of x and of Δx we may then indicate with $K = \mathbf{E}[\cdot^\dagger]$ the

correlation matrix of its column vector argument (\cdot^\dagger denotes the Hermitian transpose) and with $\text{tr}(\cdot)$ the matrix trace to write

$$\begin{aligned} \mathbf{E}[\|\Delta x\|_2^2] &= \text{tr}(K_{\Delta x}) = \\ &= \text{tr} \left\{ \mathbf{E}_{A', \Delta A, x} \left[(A')^+ \Delta Ax x^\dagger \Delta A^\dagger [(A')^+]^\dagger \right] \right\} \\ &= \text{tr} \left\{ \mathbf{E}_{A', \Delta A} \left[(A')^+ \Delta A K_x \Delta A^\dagger [(A')^+]^\dagger \right] \right\} \end{aligned}$$

so that the ratio

$$\frac{\mathbf{E}[\|\Delta x\|_2^2]}{\mathbf{E}[\|x\|_2^2]} = \text{tr} \left\{ \mathbf{E}_{A', \Delta A} \left[(A')^+ \Delta A \frac{K_x}{\text{tr}(K_x)} \Delta A^\dagger [(A')^+]^\dagger \right] \right\} \quad (3)$$

where the energy-normalized correlation matrix $K_x/\text{tr}(K_x)$ takes into account the second-order statistical properties of the signal to acquire. If the sparsity basis D is orthonormal we may adopt the widely employed sparsity model considering each of $\binom{n}{k}$ supports of s with the same probability, and its k non-null components as i.i.d. zero-mean random variables. With this, the correlation $K_s/\text{tr}(K_s) = n^{-1}I_n$ and $K_x = DK_s D^\dagger = n^{-1} \text{tr}(K_s) I_n$, where I_n is the n -dimensional identity matrix. In this case, a simplified evaluation of the Average Recovery Signal-to-Noise Ratio, $\text{ARSNR} = \mathbf{E}[\|x\|_2^2]/\mathbf{E}[\|\Delta x\|_2^2]$ due to perturbation of the encoding matrix is

$$\text{ARSNR} \simeq n \text{tr}^{-1} \left\{ \mathbf{E}_{A', \Delta A} \left[(A')^+ \Delta A \Delta A^\dagger [(A')^+]^\dagger \right] \right\} \quad (4)$$

The expectation on A' and ΔA depends on the system we are considering and may be effectively computed by Monte Carlo simulations for any given perturbation policy. From this point of view, the more suggestive and equivalent

$$\text{ARSNR} \simeq \mathbf{E}_{A', \Delta A}^{-1} \left[\frac{1}{n} \sum_{j=0}^{n-1} (\sigma_j)_{(A')^+ \Delta A}^2 \right]$$

that links the expected performance to the average of the singular values $\{(\sigma_j)_{(A')^+ \Delta A}\}_{j=0}^{n-1}$ is much less attractive due to a higher computational need.

Note that such an estimate has clear limitations. First, since it focuses on non-denoising recovery (i.e., the solution of $\text{BP}(y, A', D)$) it underestimates the attained recovery quality when the disturbance due to the perturbation can be compensated by the relative abundance of information on the problem due to (i) the availability of a large number of measurements in excess of the minimum required for recovery (therefore allowing efficient denoising) and (ii) knowing each instance's error norm $\varepsilon^* = \|\Delta Ax\|_2$ (the so-called “genie” tuning) with which $\text{BPDN}(y, A', D, \varepsilon^*)$ may be solved.

Secondly, the estimate will lose its validity for small values of m that do not allow an effective recovery, i.e., when even $\text{BP}(y, A, D)$ fully informed on A fails. In this case it is not sensible to assume that either (1) or (2) identify a good approximation of the true signal; thus, in a small- m setting the assumption that $\|\Delta x\|_2^2$ is very small is violated, and the estimate will not yield a relevant prediction of the recovery quality. Overall, (4) and the more general (3) are expected to be most effective when m is so that the *phase transition* of $\text{BP}(y, A, D)$ has occurred [16], but not much larger than the minimum m required to achieve it. Actually, this is how CS-based acquisition systems are commonly designed and why (4) will match the examples presented below.

While a variety of algorithms and problem formulations tackle the general case of signal recovery under perturbations [17, 18], significant improvements are therein shown to be possible when some structure in ΔA can be leveraged. In the absence of this side-information, e.g., if ΔA is a random matrix with i.i.d. entries drawn




Fig. 1: Comparison of the average performance estimate in (4) (dashed) against $\text{BP}(y, A', D)$ (empty circles), $\text{BPDN}(y, A', D, \varepsilon^*)$ (filled circles), $\text{OLS}(y, A', D, T)$ (solid line)

at each instance of x , as noted in [18] the performances approach those of the non-perfectly informed *oracle least-squares estimator*, $\text{OLS}(y, A', D, T)$, i.e.,

$$\hat{s} = \begin{cases} \hat{s}_T = (A'D_T)^+ y, & T = \text{supp}(s), |T| = k \\ \hat{s}_{T^c} = 0_{n-k}, & T^c = [0, n-1] \setminus \text{supp}(s) \end{cases} \quad (5)$$

(\cdot_T denotes restriction to the basis vectors of index T) while, more realistically, the average recovery performances without support information will lie between $\text{BP}(y, A', D)$ and $\text{BPDN}(y, A', D, \varepsilon^*)$.

4. NUMERICAL EXPERIMENTS

In this numerical experiment¹ we consider a simple setting of dimensionality $n = 256$ and assume D is the Discrete Cosine Transform (DCT) basis; we generate s by assuming each of its $\binom{n}{k}$ possible supports equally probable, with its k non-null components being i.i.d. random variables distributed as $\mathcal{N}(0, 1/k)$, and let $k = 8, 16, 32$ as prototypes of high-, medium-, and low-sparsity signals.

The matrix $A' \in \mathbb{R}^{m \times n}$ is here considered a Gaussian random matrix with i.i.d. entries of unit variance. As noted in the previous Section, we expect the estimate (4) to apply after $\text{BP}(y, A, D)$ solves a problem with sufficiently large m . For a quantitative evaluation of this aspect, we generate 200 signal instances, encode them with no perturbation and then solve $\text{BP}(y, A, D)$ to measure the ARSNR with different values of m by means of SPGL_1 [19]. Given

¹The code to reproduce and extend them is available at <https://sites.google.com/site/ssigprocs/CS/avpert>

that the maximum achievable signal-to-noise ratio with this solver is ≈ 120 dB, by looking at the evidence in Fig. 1a we derive that an almost maximum target ARSNR level of 110 dB is reached when $m = 103$ for $k = 8$, $m = 138$ for $k = 16$ and $m = 184$ for $k = 32$, at which it is safe to assume that the decoder is operating after the phase transition. At these (m, k) pairs we explore the effect of perturbations and how closely it is predicted by (4): we choose random ΔA and introduce the *projection-perturbation ratio* $\text{PPR} = \mathbb{E}[\|\Delta A\|_F^2] / \mathbb{E}[\|A'\|_F^2]$, i.e., the relative average energy of A' w.r.t. ΔA to control its impact. In particular, the perturbation matrix ΔA is generated from one of three random models:

1. *Dense Gaussian Additive* (DGA): the entries of ΔA are i.i.d. with $\Delta A_{j,l} \sim \mathcal{N}(0, \sigma_{\Delta A}^2)$, with $\sigma_{\Delta A}^2 = \frac{1}{\text{PPR}}$;
2. *Dense Uniform Multiplicative* (DUM): $\Delta A = U \odot A'$, where \odot is the Hadamard product, U is independent of A' and its entries are i.i.d. with $U_{j,l} \sim \mathcal{U}(-\frac{\beta}{2}, \frac{\beta}{2})$ and $\beta = 2\sqrt{\frac{3}{\text{PPR}}}$;
3. *Sparse Sign Flipping* (SSF): a random set of index pairs J is independently generated so that each entry

$$\Delta A_{j,l} = \begin{cases} -2A'_{j,l}, & (j,l) \in J \\ 0, & (j,l) \notin J \end{cases}$$

corresponds to a sign flipping of an element of A' with probability η . The resulting sparse perturbation matrix has a density η which controls $\sigma_{\Delta A}^2 = 4\eta$ with $\eta = \frac{1}{4\text{PPR}}$.

The distribution parameters are chosen to obtain a given $\text{PPR} \in \{0, 5, \dots, 80\}$ dB. On these three models and for the chosen

(m, k) , we generate 200 instances of $(s, A', \Delta A)$, encode $x = Ds$ with $A = A' + \Delta A$ and attempt to recover $\hat{x} = D\hat{s}$ by the naive BP(y, A', D); BPDN(y, A', D, ε^*) where $\varepsilon^* = \|\Delta A x\|_2$ is “genie”-tuned for each instance; the non-perfectly informed OLS(y, A', D, T). These three results are compared with the outcome of a Monte Carlo simulation of our estimate in (4) averaged over 200 instances of $(A', \Delta A)$.

The results are depicted in Fig. 1b,1c,1d for fixed $k = 16$ and the three different perturbation models; the ARSNR of each decoder can be compared with the estimate as the PPR increases (i.e., the perturbation is progressively smaller). Moreover, since the estimate has negligible variations w.r.t. the perturbation model, we fix the latter to DGA and explore the effect of different sparsity levels at values for which the phase transition has occurred; the results are reported in Fig. 1e,1b and 1f. Note that, although it is only an estimate, (4) appears to be quite effective in anticipating the average performances right between BP(y, A', D) and BPDN(y, A', D, ε). This is coherent with its derivation that starts from a non-denoising, naive basis pursuit but assumes that the recovery has the ability of coming as close as possible to the true solution in the least-squares sense.

5. APPLICATION TO MULTICLASS ENCRYPTION BY COMPRESSED SENSING

To embed a private-key security scheme into CS [5, 9–11, 20] we assume $A \in \{-1, +1\}^{m \times n}$, and partition the generation of $A', \Delta A$ into two pseudo-random number generators, PRNG $_{A'}$ and PRNG $_{\Delta A}$ which respectively expand the seeds $\omega_{A'}$ and $\omega_{\Delta A}$.

By assuming that the PRNGs have sufficiently long periods and statistical properties to mimic truly random bitstreams, we may consider A' an i.i.d. Bernoulli random matrix whose symbols have equal probability. Thus, if PRNG $_{\Delta A}$ is used to generate SSF perturbation matrices with J containing ηmn ($\eta \in (0, 1)$) index pairs (j, l) , the two matrices (A, A') share the same distribution (the sign flipping occurs independently and randomly for each instance of A') and it is not possible to derive one from the other, or even tell whether the true encoding matrix is either A or A' without some side-information.

The measurements are then produced by encoding x with A as $y = Ax$, with the first-class key being $K^{(1)} = (\omega_{A'}, \omega_{\Delta A})$ and allowing first-class decoders to construct A (the true encoding matrix), whereas the second-class key, $K^{(2)} = \omega_{A'}$ only allows second-class decoders to construct A' , i.e., a version of the encoding matrix affected by a perturbation whose energy is controlled by η . The amount of privacy provided by this scheme is investigated in [9]. We here address the problem of linking η with the average loss in recovery quality experienced by second-class decoders, which was previously tackled by quite loose upper- and lower-bounding [9]. This said, the heuristic in (4) is well-suited to anticipate the quality loss experienced by a second-class decoder in the multi-class encryption scheme described above. This is shown here by considering two exemplary signals: ECG tracks and images.

5.1. ECG Tracks

We elaborate this example by processing a large subset of ECG signals from the MIT PhysioNet database [21] sampled at $f_s = 256$ Hz. In particular we use 500 windows of $n = 256$ samples from the ECG track e0108 and encode them with $m = 144$ measurements; D is here assumed as the Symmlet-6 [22] orthonormal wavelet basis.

In this real-world setting, the ARSNR performances fall in a different range with respect to ideal, perfectly sparse synthetic sig-




Fig. 2: The effect of SSF on i.i.d. Bernoulli random matrices used to acquire real-world ECGs (η is the fraction of flipped matrix entries).




Fig. 3: The effect of SSF on i.i.d. Bernoulli random matrices used to acquire real-world images (η is the fraction of flipped matrix entries).

nals. Yet, the first-class decoder is still able to recover the original signal with ARSNR = 26.49 dB. The performances of a typical second-class decoder are exemplified by the BP(y, A', D) and BPDN(y, A', D, ε^*) tracks in Fig. 2. In the same Figure note that, as in the synthetic case, (3) is able to substantially anticipate the effect of SSF on the recovery quality of a real world-signal.

5.2. Sensitive text in images

Again, we expand this case from [9] in which we consider an image dataset of people holding printed identification text, and apply multiclass CS-based encryption to selectively hide this sensitive content to lower-class users. The 640×512 pixel images are encoded by CS in 10×8 blocks, each of 64×64 pixel while the two-class strategy is only applied to a sensitive image area of 3×4 blocks. The sparsity basis D is the 2D Daubechies-4 wavelet basis [22]. Each block of $n = 4096$ pixels is then encoded with $m = 1860$ measurements, and the performances are averaged over 100 instances of $(A', \Delta A)$. In an unperturbed case, this allows an ARSNR of 29.25 dB that is progressively reduced by perturbation as reported in Fig. 3 and effectively predicted by (3).

6. CONCLUSION

We have proposed an estimate of the average recovery performances attained by CS under random perturbation of the entries of its encoding matrix. This heuristic is simply calculated by estimating an expectation with Monte Carlo trials, and requires no prior information on the signal support; however, it applies only after the phase transition of the corresponding recovery problem. The estimate was shown to adhere with practical average recovery performances both by synthetic numerical experiments and with real-world signals.

7. REFERENCES

- [1] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 21–30, 2008.
- [2] R. M. Willett, R. F. Marcia, and J. M. Nichols, "Compressed sensing for practical optical imaging systems: a tutorial," *Optical Engineering*, vol. 50, no. 7, pp. 072 601–072 601, 2011.
- [3] M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "Leakage compensation in analog random modulation pre-integration architectures for biosignal acquisition," in *Biomedical Circuits and Systems Conference (BioCAS), 2014 IEEE*. IEEE, 2014, pp. 432–435.
- [4] C. Bilen, G. Puy, R. Gribonval, and L. Daudet, "Convex optimization approaches for blind sensor calibration using sparsity," *Signal Processing, IEEE Transactions on*, vol. 62, no. 18, pp. 4847–4856, Sept. 2014.
- [5] V. Cambareri, J. Haboba, F. Pareschi, R. Rovatti, G. Setti, and K.-w. Wong, "A two-class information concealing system based on compressed sensing," in *Circuits and Systems (IS-CAS), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 1356–1359.
- [6] M. A. Herman and T. Strohmer, "General deviants: An analysis of perturbations in compressed sensing," *Selected topics in signal processing, IEEE Journal of*, vol. 4, no. 2, pp. 342–349, 2010.
- [7] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on pure and applied mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [8] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
- [9] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "Low-complexity multiclass encryption by compressed sensing," *Signal Processing, IEEE Transactions on*, To appear, 2014. [Online]. Available: arxiv.org/abs/1307.3360
- [10] Y. Rachlin and D. Baron, "The secrecy of compressed sensing measurements," in *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*. IEEE, 2008, pp. 813–817.
- [11] T. Bianchi, V. Bioglio, and E. Magli, "On the security of random linear measurements," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 3992–3996.
- [12] J. A. Tropp and S. J. Wright, "Computational methods for sparse solution of linear inverse problems," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 948–958, 2010.
- [13] P. L. Combettes and J.-C. Pesquet, "A douglas–rachford splitting approach to nonsmooth convex variational signal recovery," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 1, no. 4, pp. 564–574, 2007.
- [14] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18 914–18 919, 2009.
- [15] P. Maechler, C. Studer, D. E. Bellasi, A. Maleki, A. Burg, N. Felber, H. Kaeslin, and R. G. Baraniuk, "Vlsi design of approximate message passing for signal restoration and compressive sensing," *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on*, vol. 2, no. 3, pp. 579–590, 2012.
- [16] D. L. Donoho and J. Tanner, "Precise undersampling theorems," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 913–924, 2010.
- [17] H. Zhu, G. Leus, and G. B. Giannakis, "Sparsity-cognizant total least-squares for perturbed compressive sampling," *Signal Processing, IEEE Transactions on*, vol. 59, no. 5, pp. 2002–2016, 2011.
- [18] J. T. Parker, V. Cevher, and P. Schniter, "Compressive sensing under matrix uncertainties: An approximate message passing approach," in *Signals, Systems and Computers (ASILOMAR), 2011 Conference Record of the Forty Fifth Asilomar Conference on*. IEEE, 2011, pp. 804–808.
- [19] E. Van Den Berg and M. P. Friedlander, "Probing the pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008.
- [20] A. Orsdemir, H. O. Altun, G. Sharma, and M. F. Bocko, "On the security and robustness of encryption via compressed sensing," in *Military Communications Conference, 2008. MILCOM 2008. IEEE*. IEEE, 2008, pp. 1–7.
- [21] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000 (June 13), circulation Electronic Pages: <http://circ.ahajournals.org/cgi/content/full/101/23/e215> PMID:1085218; doi: 10.1161/01.CIR.101.23.e215
- [22] S. Mallat, *A wavelet tour of signal processing: the sparse way*. Access Online via Elsevier, 2008.