
MINIMIZING ACQUISITION MAXIMIZING INFERENCE - A DEMONSTRATION ON PRINT ERROR DETECTION

A PREPRINT

Suyash Shandilya

Defence Research and Development Organisation
New Delhi - 110054
su.sh2396@gmail.com

August 25, 2019

ABSTRACT

Is it possible to detect a feature in an image without ever being able to look at it? Images are known to be very redundant in spatial domain. When transformed to bases like Discrete Cosine Transform (DCT) or wavelets, they acquire a sparser (more effective) representation. Compressed Sensing is a technique which proposes simultaneous acquisition and compression of any signal by taking very few random linear measurements (M) instead of uniform samples at more than twice the bandwidth frequency (Shannon-Nyquist theorem). The quality of reconstruction directly relates with M , which should be above a certain threshold (determined by the level of sparsity, k) for a reliable recovery. Since these measurements can non-adaptively reconstruct the signal to a faithful extent using purely analytical methods like Basis Pursuit, Matching Pursuit, Iterative thresholding, etc., we can be assured that these compressed samples contain enough information about any relevant macro-level feature contained in the (image) signal. Thus if we choose to deliberately acquire an even lower number of measurements - in order to thwart the possibility of a comprehensible reconstruction, but high enough to infer whether a relevant feature exists in an image - we can achieve accurate image classification while preserving its privacy. Through the print error detection problem, it is demonstrated that such a novel system can be implemented in practise.

Keywords Compressive Sensing · Smashed filters · compressive classification · Single Pixel Camera · Data Privacy · Reprography · Print error detection · Support Vector Machines · Target Classification · Feature Detection

1 Introduction

1.1 The Print error detection problem

Printing errors are a very pervasive problem in any large scale printing mill. Generally any conspicuously erroneous material is reprinted while innocuous ones are left as is. But when the data being printed is sensitive (mathematical equations, numbers, encryption keys, etc.) or if the user is cautious about aesthetics, a single ambiguity in reading can impair an entire document. The cost of reprinting is usually infinitesimal in comparison to the jeopardy an unusable data might land the user in. Thus a system to proofread every sheet of every document is imperative in such a case.

A simple automation solution to the problem could be as follows. After the final print, each sheet is scanned. Simple algorithms can be used to extract the text regions from the scans. These individual word segments will constitute the dataset (ref 1.4) which can be used to train any standard classifier - Decision trees, Support Vector Machines (SVMs), Convolutional Neural Networks (CNN) etc.

The catch in this solution is very basic - many users do not want their classified data to be digitised. From cyberattacks to system compatibility, there are many reasons why many organizations avoid digitisation of sensitive material. Thus

proofreading in such cases needs to be done manually. This taxes all categories of resources like time, money, accuracy, reliability, human resource, etc. This seemingly makes the problem impossible to automate unless there could be a system which can detect these errors without digitally 'looking' at the data

1.2 Compressed Sensing

Compressed sensing has been a prolific research topic in applied mathematics and statistics over the last few years. Basics of compressed sensing can be found in studies by Candes, Romberg, Tao [1] and Donoho [2]. Readers may refer to [3] as an excellent introductory tutorial. A very brief introduction is as follows.

A typical signal acquisition process conventionally begins with sampling the signal in spatial or time domain at a frequency much higher than the Nyquist sampling rate [4]. Most of the signals (specially images) have a lot of redundancy at this initial stage. As a practical and pertinent example, a 10 megapixel RGB camera with a bitdepth of 16 bits per pixel should produce a 20 MB image. But the JPEG file format compresses the image to about a tenth of its size, if not more. For almost all practical purposes, this compressed file is as good as the raw image with the added mathematical and portability convenience. This depicts an innate inefficiency in the system wherein we first deploy expensive hardware to acquire data at a high sampling frequency and subsequently discard most of the samples.

Compressed sensing proposes that if the signal is sparse in some domain (natural images are sparse in wavelets, DCT) then a sufficient number of incoherent linear non-adaptive measurements of the signal can be used to recover the signal as accurately as done by any conventional sampling system. This simplifies the acquisition process and greatly reduces the hardware complexity. The reconstruction however, is more computationally expensive. Essentially we are simply trying to obtain the sparsest possible representation of our signal. This can be mathematically expressed as:

$$\begin{aligned} & \underset{x}{\text{minimize}} && \|x\|_0 \\ & \text{subject to} && y = \phi x \end{aligned}$$

Where $y \in \mathbb{R}^M$ is the obtained compressed sample vector, $x \in \mathbb{R}^N$ is the original **sparse** signal and $\phi \in \mathbb{R}^{M \times N}$ is called as the *sensing matrix* ($M \ll N$). Here $\|\cdot\|_0$ denotes the l_0 -norm which is equal to the number of non-zero components in a vector. This is known to be an np-hard problem [2]. The real genius of compressive sensing lies in proving that if we relax the objective constraint by using l_1 instead of l_0 -norm, the problem changes to

$$\begin{aligned} & \underset{x}{\text{minimize}} && \|x\|_1 \\ & \text{subject to} && y = \phi x \end{aligned}$$

which is a convex problem. This problem has been termed as Basis Pursuit. More practical version of the problems like Basis Pursuit Denoising (BPDN) and LASSO are also popular in literature. Readers are advised to refer to [5] for more details on its variants.

Often x is not sparse in the original domain but has a sparse representation in a transformed domain like wavelets, etc. In that case, x is replaced with $\psi'x$ where ψ is the transformation matrix of the domain in which x is sparse.

The sensing matrix ϕ is very crucial to the concept here. It describes how the linear measurements are obtained from x . The number of columns in ϕ should be equal to the length of the signal $x (= N)$ while the number of rows of ϕ determine the number of compressed samples (M) the user wants. Higher the M better the reconstruction accuracy. Two important properties are associated with the sensing matrix: *incoherence* and *Restricted Isometry Property* (RIP). Simply stated, they ensure that the sensing matrix takes incoherent (well-spread) measurements, and disparate signals should not map to similar measurement values. This sensing matrix is expected to be maximally incoherent with the transformation matrix ψ . It was found that random matrices like Gaussian, Bernoulli matrix satisfy both the aforementioned properties with a very high probability. Here binary Gaussian matrix has been used for reasons better described in 1.3.

Greedy algorithms like Orthogonal Matching pursuit [6] and many of its variants like ROMP [7], StOMP [8], CoSaMP [9], are also used to recover the sparse signal. They are much faster than Basis Pursuit methods but have poor theoretical guarantees. Since the objective of the paper is to ensure minimal signal recovery in the worst possible case, I adhered to

evaluating the recovered signals via Basis Pursuit only.

1.3 Single Pixel Camera and Smashed Filters

The first practical implementation of compressed sensing was demonstrated by Duarte, et al. when they created a single pixel camera (SPC) [10]. They used a Digital Micromirror Device (DMD) which is an array of small mirrors which can be digitally aligned along two binary states (-10° or $+10^\circ$ from the axis). It is set in a way that one of the states aligns with the single photon detector while other reflects away from it. This essentially simulates a binary matrix. Since the system proposed in this work proposes this hardware, hence only binary matrix was chosen for all the experimentation in this paper. The reader may refer to [10] for a more elaborate understanding of the working of the model device.

In a similar work, Davenport et al. introduced the concept of *smashed filters* [11]. Therein the authors posit that operations like target classification can be performed on the compressed samples themselves by applying maximum likelihood classifier. This has been one of the primary motivation for this work. More on this in section 2.

1.4 Dataset

No pre-existing digital repertoire of common print errors could be found. So the experiments here have been conducted on simulated print errors. They have broadly been binned into three major types of errors as shown in figure 1. These artifacts have been chosen to be just big enough to cause an ambiguity in the text.

The text images were generated using the PIL library of python. All possible 3 letter words were made as grayscale images of pixel dimension 35×100 . Arial bold font was used. They were labelled as ‘good’ images. The three types of errors were distributed equally on the entire dataset of $17576 (= 26 \times 26 \times 26)$ images. These were labelled as bad images.



Figure 1: The 3 major type of print errors found in printing press. Other types of errors in most part are subsumed in either of these categories. Note that these are simulations of zoomed-in scans of text segments. A typical sheet may contain more than 500 such words. In terms of pixel ratio, each word shown here is about 0.1% of the full sheet.

1.5 Paper Layout

There was a good amount of skepticism between the conception and testing of this concept. As a fail-safe another application was conceived - It would be proposed as a safer way of obtaining any image analysis as only a user with the exact key will be able to reconstruct the meaningful image. Fortunately (and mildly surprisingly), that became a secondary application as even extreme levels of compression gave remarkable results.

Section 2 discusses similar works which have preceded this paper. Sections 3,4 explain how the data was compressed and reconstructed (only for testing) respectively. Section 5 discusses the chosen models for classification and compares their performance. Section 6 describes other groundbreaking applications where this concept should be tested.

1.6 A note on terminology

The words compression and encryption are used very interchangeably throughout the text.

To be technically precise, the concept is to draw inference from a very heavily (lossy) compressed data. But the novel idea of this work is the application of such a system to data privacy. Thus, not technically accurate, the more application oriented description of the work would be the idea to draw inference from a severely *encrypted* data.

- Key \equiv sensing matrix (ϕ).
- Encryption \equiv compression ($x \rightarrow \phi x = y$)
- Decryption \equiv Reconstruction ($\phi x = y \rightarrow x$)

A similar terminology has also been described in [12] (II-C).

2 PREVIOUS WORK

Probably the first work to demonstrate Compressive Sensing’s capabilities in target classification was the Single Pixel Camera developed by Davenport et al. [11]. Their results bespeak that effective classification could be achieved from only compressed samples without any need of reconstruction. Therein the authors introduce smashed filters which uses maximum likelihood for classification (MLE)¹. The concept is demonstrated on images of toy tanker, school bus and truck and the accuracy results are calculated using leave-one-out validation². Since there were not privacy concern in their objective, hence there is no analysis in that regard. In fact there have been several other demonstrations of the application of smashed filters in different domains ([13][14][15] to cite a few) but none seem to have designed or evaluated a system with privacy concerns as priority.

Compressed sensing has been extensively studied from an encryption system standpoint. [16] was the first work where it was analysed as a symmetric encryption system. The authors prove that it does not provide information theoretic secrecy but argue that it is computationally secure, as long as the sensing matrix is used only once. Authors of [12] assert its effectiveness as an obfuscation layer in a cryptosystem. In [17], the authors prove the computational security of CS against a systematic search of the sensing matrix, even when the signal sparsity is known. A practical CS system with two confidentiality levels based on Bernoulli sensing matrices has been proposed and analyzed in [18], [19] and its security has been additionally investigated in [20]. In all such analyses the sensing matrix is considered to be available for the intended receiver and the data needs to be recovered reliably. It should be noted that both these assumptions are inconsequential here as we intentionally *crumble* the data to ensure its unusability.

Another noteworthy contribution has been made by Zhou [21] where the authors use compressed sensing as a matrix masking [22] tool on an entire statistical dataset to achieve “privacy preserving regression analysis”. To the best of my knowledge, this is the first work to demonstrate the simultaneous efficacy of compressed sensing in ensuring image security and high accuracy in image classification.

3 Data Encryption

Encrypted data ($y \in \mathbb{R}^M$) is achieved by the multiplication of a binary sensing matrix ($\phi \in \{0, 1\}^{M \times N}$) with the original data ($x \in \mathbb{R}^N$). This is a form of lossy encryption as $M \ll N$, thus the data is not likely to be fully recovered. Nevertheless, any level of accurate decryption requires ϕ to be known exactly, thus the sensing matrix is essentially the *key* for this encryption here. For a successful recovery, we need to choose a basis in which the data is represented sparsely enough. Daubechies – 10 wavelets [23] gave much better sparsity than other explored wavelets and DCT bases. For the sake of clarity and transparency, it should be mentioned that the image was first transformed to the wavelet domain, and then multiplied with the sensing matrix. In the hardware implementation however, we obtain ϕx (i.e. apply the compression on the original image) as y , we then find an x with minimum l_1 norm, subject to $y = \phi \psi' x$. Both the methods produce the exact same result. The former was relatively faster and slightly more convenient, thus was chosen as the preferred way of implementation.

Now comes the critical question of determining the number of compressed samples to be considered for inference. I simply began reconstructing random ‘good’ images for different values of M starting with $M = 1000$ (28.6% compression) and decreasing it with each step. It was realised that even for $M = 500$ (14.3% compression) the images were somewhat readable. When M was brought down to 200 (5.7%), it was seen that the text can no longer be reliably read. Thus the number of samples were limited to a maximum of 200. The inference results were then calculated for even lower values. As they began rendering incredible performance on a sample set, 50 (1.4%), 20 (0.57%) and 10 (0.28%) were additionally chosen for experimentation. The reconstruction result for various values of M can be seen in fig 2.

4 Data Decryption

The decryption part is sort of an adversary to the objective of data privacy. In general, the sensing matrix need never exist mathematically if the user decides not note it. But for the sake of assurance they may want to decipher the acquired data. This would require the exact knowledge of the key *i.e.* This encryption itself should be a sufficient level of security for most applications. However, in the event of an unsolicited disclosure of the key (or simply when the system is setup by an external agent who we may not expect to have deleted the key), the user may want to be assured that the

¹This is a primary difference between this work and [11]. In the case of AWGN, the generalised MLE is equivalent to *nearest neighbour* classification. We tried using the same in our problem only to find it sub-par in comparison to other classifiers like polynomial discriminants and SVMs

²All the results mentioned in this work are holdout validation results (the testing dataset was never used anywhere in evaluating the discriminant) unless stated otherwise

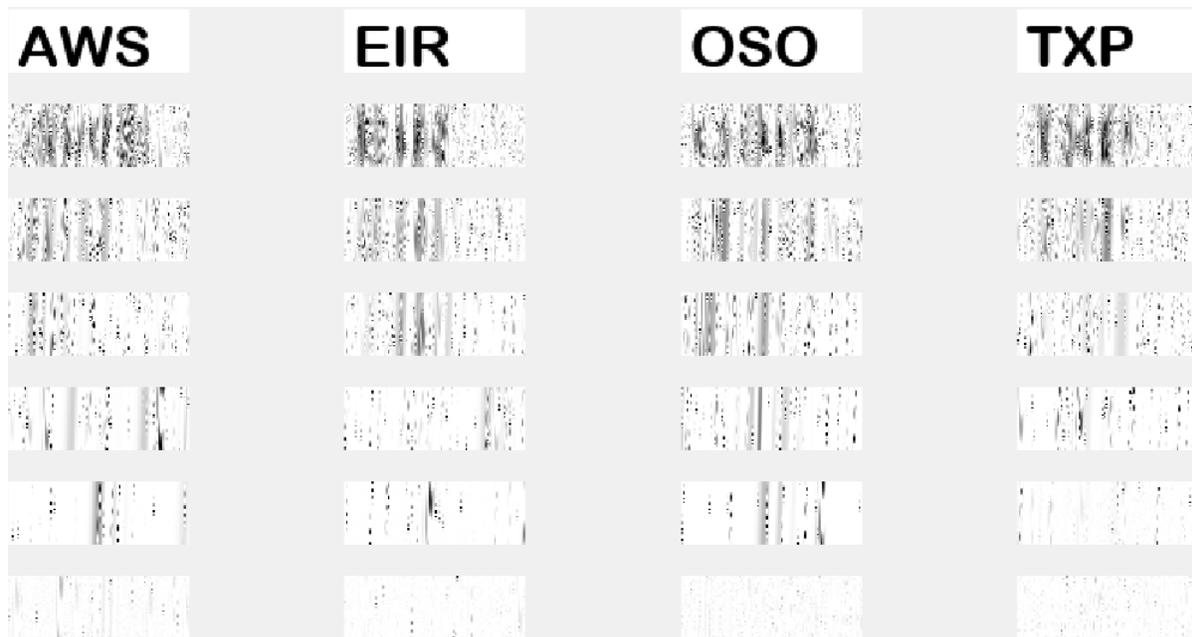


Figure 2: The top row corresponds to the original image. The second row corresponds to reconstruction using 500 compressed samples ($M = 500$, **only shown for reference, not used for classification**) where the characters are faintly visible. The subsequent rows correspond to $M = 200, 100, 50, 20, 10$ respectively.

decryption is futile if the motive is to acquire the original data.

Since we are considering the worst case scenario, we intend to have the best possible reconstruction regardless of the computational complexity. We also assume that the compressed samples have zero error in measurement. The objective is to ascertain that even with the best resources and maximum information at hand, the reconstruction acquired samples will render no meaningful result. To that end, it was decided to use only basis pursuit. Thus the mathematical formulation was formalised simply as

The convex optimisation package CVX [24][25] was used to solve it. A typical reconstruction with $M = 1000$ took about 200 seconds while it only took 2 seconds for $M = 10$.

5 CLASSIFICATION METHODS AND RESULT ANALYSIS

Once the maximum number of compressed samples was established, the dataset was bifurcated into training and testing sets. 15000 images from each label were chosen for training while 2576 were chosen for testing from each label. The $30000 \times M$ matrix of training data was augmented with a column of categorical label vector (0 for good, 1 for bad). Several different classification models were tested with various different parameter values. 5-fold validation was used to compare the results of the different models. It was realised that Support Vector Machines were the best performing among all. The choice of the kernel function and scale significantly affected the performance. In general, cubic kernel function gave the best performance. The kernel scale varied largely for different values of M . The performance shown in figure 3 and table 1 can be considered to have the most optimised parameter values.

k -Nearest Neighbour did not perform as good as other models. Performance by different Decision trees and ensemble classifiers like Boosted trees, bagged trees, subspace discriminant, too performed relatively sub par. Simple classifiers like Linear and quadratic discriminant, logistic regression, performed reasonably better than them. Though the difference was not huge, these classifiers provide a more intuitive understanding of how the different categories are separated in the M -dimensional hyperspace.

Looking at table 1, the general order of performance would be Cubic SVM > Gaussian SVM > Quadratic discriminant > Logistic Regression > Linear discriminant. Random subsets of misclassified images were analysed for all these models. In general the mix contained all types of ‘bad’ images in about equal proportion. Thus we can be assured that the classifier did not fixate on any particular measurement. However there are two observations that need to be stated. The most pertaining observation was the consistent asymmetry between the types of misclassification errors – The bad images were classified as good more often than vice versa. The chance exception was in the case of the quadratic

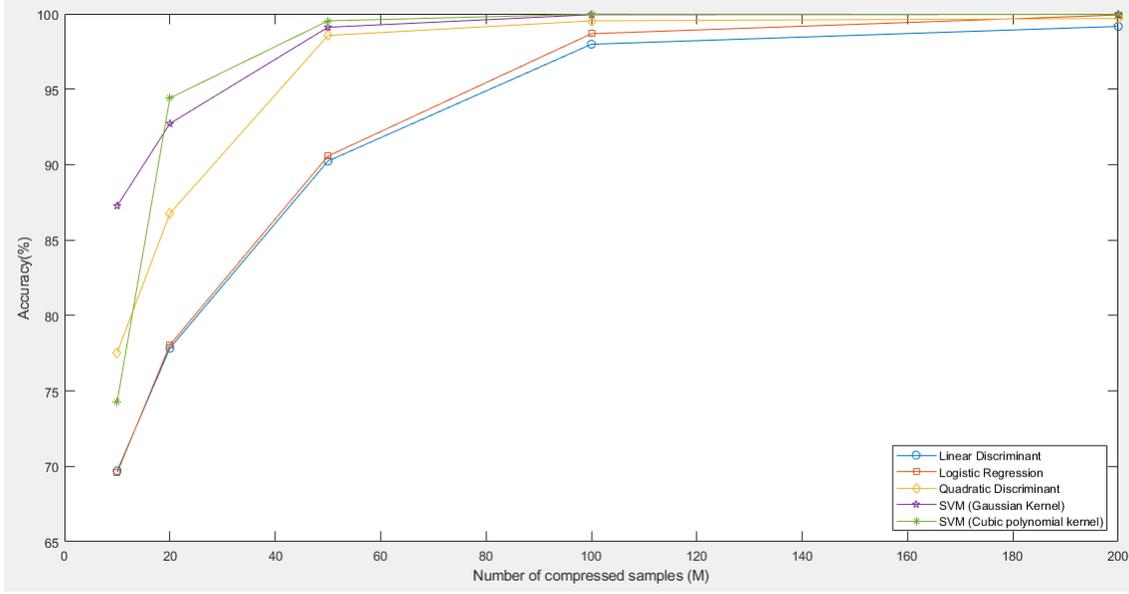


Figure 3: Performance of various classification models versus the number of compressed samples. Refer 1 for the precise values.

M	Gaussian SVM	Cubic SVM	Quadratic Discriminant	Logistic Regression	Linear Discriminant
200	99.9612	99.9612	99.6894	99.9224	99.1654
100	99.9418	99.9806	99.5342	98.6801	97.9814
50	99.1071	99.5342	98.5637	90.5668	90.2368
20	92.7213	94.4099	86.7624	78.0474	77.8339
10	87.2671	74.3012	77.5427	69.5846	69.7011

Table 1: Performance of different classifiers. The accuracy's are in %. Figure 3 is a graphical representation of the same.

discriminant for $M = 200$ wherein all the erroneous classifications were for ‘good’ images which is quite an outstanding aberration. Another odd consistency was regarding the characters in the misclassified images. Despite all being equal in frequency, the character ‘W’ happened to be in most of them, regardless of the label. Other frequent characters were ‘N’, ‘M’ and ‘X’. It can be observed that these characters have the most prominent diagonal segments. But that is merely stating correlations without establishing causality. A concrete reason for neither of these oddities could be established. I heartedly welcome any analysis the reader would like to share which could shed more light on these observations.

6 Future Work

Throughout the course of developing this work, I was most fixated on drafting this section. As and when such a system is designed, using it in detecting print errors would be the least of its use. While advances in artificial intelligence has increased the intelligence quotient of our computers, incrementing Emotional intelligence is still a humongous and very coveted task. Although we now have a plethora of AI architectures, they all have limited scope of implementation where user/data privacy is a concern. If we can establish that human emotions can be inferred from digital portraits with a remarkable accuracy via such limited and incoherent measurement that the original image can never be reconstructed to reveal the user identity, it would be landmark in the field of emotion recognition. Having said that, the face database would be much less cooperative than the one used here as faces are much more diverse and have relatively small features that help in accurate detection. This idea can be extended to any field where data privacy is as important as accurate inference, if not more. Exploring its results in such fields would be my immediate future endeavour on this work. Many crucial analysis were left out in this work for the sake of accentuating the proof of concept. The SPC modelled Poisson noise in its measurements and pursued the reconstruction accordingly. Such an addition might affect accuracy or privacy either favourably or otherwise.

Deep learning empowered reconstruction algorithms like Reconnet [26] Stacked Denoising Auto-encoders (SDAs) [27] have lately shown remarkable efficacy in reconstructing images from very few compressed samples (as low as 1%). Testing the data privacy against such reconstruction algorithms should be an imperative course in the line of its advancement. On the other hand, deep neural networks are known to be much “deeper” in understanding the classifier of many datasets which are otherwise difficult to extricate. Such advancements should invigorate both objectives of this paper.

There are many successful variants of compressive sensing which make use of the structure of data to provide even better reconstructions. Among the most pertaining ones would be block sparsity based reconstruction [28]. Since our data is clearly block sparse in spatial as well as transformed domain, such an algorithm is likely to deliver better results than the basic basis pursuit used here.

7 CONCLUSION

Our mantra in this paper has been to “*maximise inference while minimising acquisition*”. We took the print error detection problem in reprography – where the sensitivity of the print data forbids the digitisation, or any unwarranted access in general – to demonstrate that using *compressive classification*, we can achieve remarkable accuracy in feature detection while ensuring that the data is practically unaccessed. Even at compression ratio as low as 0.3% and 0.6%, the results were very satisfactory. We tested multiple models to achieve the desired classification rate. Kernel SVMs (Cubic and Gaussian specifically) turned out to be most efficient in general. It was observed that even simple linear discriminants provided astounding accuracy when the compression ratio was relatively higher, but not high enough to provide meaningful reconstruction.

References

- [1] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, Feb 2006.
- [2] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, April 2006.
- [3] E. J. Candes and M. B. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, March 2008.
- [4] P. P. Vaidyanathan. Generalizations of the sampling theorem: Seven decades after nyquist. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 48(9):1094–1109, Sep. 2001.
- [5] Ewout Van Den Berg and Michael P Friedlander. Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912, 2008.
- [6] L. Rebollo-Neira and D. Lowe. Optimized orthogonal matching pursuit approach. *IEEE Signal Processing Letters*, 9(4):137–140, April 2002.
- [7] Joel A Tropp and Anna C Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on information theory*, 53(12):4655–4666, 2007.
- [8] David Leigh Donoho, Iddo Drori, Yaakov Tsaig, and Jean-Luc Starck. *Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit*. Department of Statistics, Stanford University, 2006.
- [9] Deanna Needell and Joel A Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and computational harmonic analysis*, 26(3):301–321, 2009.
- [10] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):83–91, March 2008.
- [11] Mark A. Davenport, Marco Duarte, Michael B. Wakin, Jason N. Laska, Dharmpal Takhar, Kevin Kelly, and Richard Baraniuk. The smashed filter for compressive classification and target recognition - art. no. 64980h. *Proceedings of SPIE*, 6498, 02 2007.
- [12] T. Bianchi, V. Bioglio, and E. Magli. Analysis of one-time random projections for privacy preserving compressed sensing. *IEEE Transactions on Information Forensics and Security*, 11(2):313–327, Feb 2016.
- [13] H. Braun, P. Turaga, and A. Spanias. Direct tracking from compressive imagers: A proof of concept. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8139–8142, May 2014.
- [14] M. R. Balthasar, M. Leigsnering, and A. M. Zoubir. Compressive classification for through-the-wall radar imaging. In *2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*, pages 2288–2292, Aug 2012.

- [15] A. Amaravati, S. Xu, J. Romberg, and A. Raychowdhury. A 130 nm 165 nj/frame compressed-domain smashed-filter-based mixed-signal classifier for “in-sensor” analytics in smart cameras. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 65(3):296–300, March 2018.
- [16] Y. Rachlin and D. Baron. The secrecy of compressed sensing measurements. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 813–817, Sep. 2008.
- [17] A. Orsdemir, H. O. Altun, G. Sharma, and M. F. Bocko. On the security and robustness of encryption via compressed sensing. In *MILCOM 2008 - 2008 IEEE Military Communications Conference*, pages 1–7, Nov 2008.
- [18] V. Cambareri, J. Haboba, F. Pareschi, H. R. Rovatti, G. Setti, and K. Wong. A two-class information concealing system based on compressed sensing. In *2013 IEEE International Symposium on Circuits and Systems (ISCAS2013)*, pages 1356–1359, May 2013.
- [19] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti. Low-complexity multiclass encryption by compressed sensing. *IEEE Transactions on Signal Processing*, 63(9):2183–2195, May 2015.
- [20] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti. On known-plaintext attacks to a compressed sensing-based encryption: A quantitative analysis. *IEEE Transactions on Information Forensics and Security*, 10(10):2182–2195, Oct 2015.
- [21] Shuheng Zhou, John Lafferty, and Larry Wasserman. Compressed and privacy-sensitive sparse regression. *IEEE Transactions on Information Theory*, 55(2):846–866, 2009.
- [22] George T Duncan, Robert W Pearson, et al. Enhancing access to microdata while protecting confidentiality: Prospects for the future. *Statistical Science*, 6(3):219–232, 1991.
- [23] Ingrid Daubechies. *Ten lectures on wavelets*, volume 61. Siam, 1992.
- [24] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [25] Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- [26] Kuldeep Kulkarni, Suhas Lohit, Pavan Turaga, Ronan Kerviche, and Amit Ashok. Reconnet: Non-iterative reconstruction of images from compressively sensed measurements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 449–458, 2016.
- [27] Ali Mousavi, Ankit B Patel, and Richard G Baraniuk. A deep learning approach to structured signal recovery. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1336–1343. IEEE, 2015.
- [28] Yonina C Eldar and Helmut Bolcskei. Block-sparsity: Coherence and efficient recovery. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2885–2888. IEEE, 2009.