RECONSTRUCTION OF SPARSE SIGNALS USING LIKELIHOOD MAXIMIZATION FROM COMPRESSIVE MEASUREMENTS WITH GAUSSIAN AND SATURATION NOISE

Shuvayan Banerjee^{††}, Radhendushka Srivastava^{††}, Ajit Rajwade[†]

^{††}Dept. of Mathematics, and [†]Dept. of Computer Science & Engg.; IIT Bombay, India

ABSTRACT

Most compressed sensing algorithms do not account for the effect of saturation in noisy compressed measurements, though saturation is an important consequence of the limited dynamic range of existing sensors. The few algorithms that handle saturation effects either simply discard saturated measurements, or impose additional constraints to ensure consistency of the estimated signal with the saturated measurements (based on a known saturation threshold) given uniform-bounded noise. In this paper, we instead propose a new data fidelity function which is directly based on ensuring a certain form of consistency between the signal and the saturated measurements, and can be expressed as the negative logarithm of a certain carefully designed likelihood function. Our estimator works even in the case of Gaussian noise (which is potentially unbounded) in the measurements. We prove that our data fidelity function is convex. Moreover, we show that it satisfies the condition of Restricted Strong Convexity and thereby derive an upper bound on the reconstruction error of the estimator. We also show that our technique experimentally yields results superior to the state of the art under a wide variety of experimental settings, for compressive signal recovery from noisy and saturated measurements.

Index Terms— Compressed sensing, Noisy and Saturated measurements

1. INTRODUCTION

Compressed sensing (CS) aims to recover the true signal $x^* \in \mathbb{R}^n$ from its 'compressive measurements' of the form $y = Ax^* + \eta$ where $A \in \mathbb{R}^{m \times n}$ with $m \ll n$ is a sensing matrix representing the forward model of the compressive device, and $y \in \mathbb{R}^m$ is a vector of (possibly noisy) compressive measurements. The noise vector is $\eta \in \mathbb{R}^m$. Although this problem is ill-posed for most vectors in \mathbb{R}^n , CS theory states that it is well-posed and that the signal x^* can be recovered with high accuracy [4], if x^* is a sparse (or weakly-sparse) vector, and A obeys the so-called restricted isometry property (RIP). A sensing matrix A is said to obey the RIP of order s, if for any s-sparse vector x^* , we have $||Ax^*||_2^2 \approx ||x^*||_2^2$.

Here, the degree of approximation is given by the so-called *s*-order restricted isometry constant (RIC) of *A*. There exist precise error bounds for the recovery of x^* [4]. Moreover, most of the algorithms for CS recovery are also efficient in terms of computation speed, a well-known example being the LASSO [6], which seeks to minimize the objective function $J(x) \triangleq ||y - Ax||_2^2 + \lambda ||x||_1$, given a regularization parameter λ .

However, the vast majority of the literature assumes a zero mean i.i.d. Gaussian distribution (with known variance) as the noise model. Many practical sensing systems, on the other hand, innately enforce noise of other distributions. Almost all sensors have a fixed (and usually known) dynamic range [a,b], a < b. However the underlying signal may be such that not all measurements $A^i x^*$ (where A^i is the *i*th row of A) can be accommodated within this range. Such measurements then get 'clipped' to the value a if $A^i x^* < a$, or to the value b if $A^i x^* > b$. This is called the 'saturation effect', and is common in all sensing systems (not only the compressive ones).

Problem statement: In this paper, we consider the following forward model for the measurements y for a compressive device with dynamic range $[-\tau, \tau]$:

$$\forall i \in \{1, 2, ..., m\}, y_i = \mathcal{C}(\mathbf{A}^i \mathbf{x}^* + \eta_i; -\tau, \tau).$$
(1)

Here the noise values are i.i.d., with $\eta_i \sim \mathcal{N}(0, \sigma^2)$ with known standard deviation σ . Also $\mathcal{C}(q; a, b)$ is a saturation operator defined as follows:

$$\mathcal{C}(q; a, b) = \begin{cases} a \text{ if } q < a, \\ b \text{ if } q > b, \\ q \text{ if } q \in [a, b]. \end{cases}$$
(2)

If q < a, the value of q is clipped to a, which is called 'negative saturation'. If q > b, the value of q is clipped to b, which is called 'positive saturation'. Given the forward model in Eqn. 1 with known A and τ , we seek to recover a sparse/weakly-sparse vector x from its compressive measurements y.

1.1. Previous Work

There exists a moderate-sized literature on the problem of CS recovery from saturated measurements, which we summarize here. Right through this paper, we use S^-, S^+ to

AR thanks SERB Matrics Grant #10013890. The authors also thank the anonymous reviewers whose suggestions greatly improved this paper.

denote the sets that respectively consist of indices of negatively and positively saturated measurements. Also, S denotes the set of indices of all measurements, and the set of indices of non-saturated measurements is $S_{ns} \triangleq S - S^+ - S^-$. We assume S_{ns} , S^+ and S^- to be known beforehand. The work in [7] proposes two types of estimators for CS recovery from measurements with saturation effects and uniform quantization (i.e., bounded) noise: (1) 'saturation rejection' (SR), which weeds out saturated measurements and performs recovery only from the non-saturated measurements via the estimator: min $\|\boldsymbol{x}\|_1$ s. t. $\sum_{i \in S_{ns}} (y_i - \boldsymbol{A}^i \boldsymbol{x})^2 \le \epsilon_{ns}^2$; and (2) 'saturation consistency' (SC), which imposes the added constraint in the SR estimator that $\forall i \in S^-, A^i x \leq -(\tau - \Delta)$ and $\forall i \in S^+, A^i x \geq \tau - \Delta$, where Δ denotes quantization width. The SR method potentially ignores many useful measurements (depending on the relation between τ and $||x||_2$), and in the worst case the remaining part of the sensing matrix may not obey the RIP due to an insufficient number of measurements. The SC method is hard to adapt to saturation effects with Gaussian noise, which is unbounded in nature. The work in [8, 9] seeks to optimize the following cost function:

$$J_{ss}(\boldsymbol{x}) \triangleq \lambda(\|\boldsymbol{x}\|_{1} + \|\boldsymbol{r}\|_{1}) + \|\boldsymbol{y} - (\boldsymbol{A}\boldsymbol{x} + \boldsymbol{r})\|_{2}^{2} \\ = \lambda\|\boldsymbol{x}; \boldsymbol{r}\|_{1} + \|\boldsymbol{y} - [\boldsymbol{A}|\boldsymbol{I}](\boldsymbol{x}; \boldsymbol{r})\|_{2}^{2}.$$
(3)

When adapted to handle saturation, r would refer to the error due to saturation effects. Also, (x; r) is the concatenation of column vectors x, r; I is the $n \times n$ identity matrix; and the $||r||_1$ term promotes sparsity on the vector r. In this paper, we term this approach 'saturation sparsity' (Ss). Although [8, 9] prove RIP of [A|I], that property is true only in an asymptotic sense as $m \to \infty$ (with $n \to \infty$ and $m/n \to 0$). In the realistic regime when m, n is small, we have observed that such a technique has a tendency to estimate r to be a vector of all zeroes, due to the penalty on $||r||_1$. Recent work in [13] proposes a greedy approximation algorithm to minimize the following cost function, designed to be resilient to measurement outliers:

$$J_{\alpha}(\boldsymbol{x}) \triangleq \|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_{p}^{p} + \lambda \|\boldsymbol{x}\|_{0}; 0
(4)$$

An approximation algorithm to minimize such a cost function is essential, as the ℓ_0 pseudo-norm otherwise renders this problem to be NP-hard. Note that the approaches in [8, 9, 13] were designed for *general impulse noise* and *not* for saturation effects, and hence these methods do not use knowledge of the saturation threshold τ . Very recent work in [5] provides theoretical bounds for the following interesting estimator, termed 'noise-cognizant ℓ_1 -minimization' (NCLM):

$$\operatorname{argmin}_{\boldsymbol{x},\boldsymbol{r}} \|\boldsymbol{x}\|_{1} \text{ such that } (i)\mathcal{C}(\boldsymbol{A}\boldsymbol{x}+\boldsymbol{r};-\tau,\tau) = \boldsymbol{y}, \quad (5)$$

$$(ii)\|\boldsymbol{r}\|_{2} \leq \gamma_{1}\epsilon; (iii)\|\boldsymbol{x}\|_{2} \leq \gamma_{2}\mu\sqrt{m}.$$

The parameters γ_1, γ_2, μ need to be selected based on properties of the sensing matrix, ϵ is a bound on $\|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_2$, and the vector r plays the same role as in Eqn. 3. Our method presented in this paper does not require the choice of parameters γ_1, γ_2 or an upper bound on $||\mathbf{x}||_2$. An algorithm that deals with saturation in CS measurements explicitly is presented in [12] along with a dictionary learning framework. They propose a data fidelity term that is different from our probabilistically motivated one. The difference in the behavior of the two data fidelity terms will be prominent if \mathbf{y} contains saturated measurements which would have otherwise attained very large magnitude in the absence of saturation. Moreover, their work does not present any error bound analysis unlike our work here.

The rest of this paper is organized as follows. The main objective function and its properties are presented in Sec. 2. Several numerical results are presented and discussed in Sec. 3. We conclude in Sec. 4 with a discussion of avenues for future work.

2. MAIN METHOD

In this section, we first present the cost function which we seek to optimize, for CS recovery under saturated measurements. Although we consider the signal x to be sparse in the canonical basis, our method is easily extensible to a signal that in sparse/weakly sparse in any known orthonormal basis (see Sec. 3). In the following, $\Phi(.)$ denotes the cumulative distribution function (CDF) of a standard normal random variable, and $\phi(.)$ denotes its probability density function (PDF).

2.1. Cost function and its properties

Our cost function $J_{our}(x)$ is given below:

$$J_{our}(\boldsymbol{x}) = \lambda \|\boldsymbol{x}\|_1 + L(\boldsymbol{y}, \boldsymbol{A}\boldsymbol{x}; \tau),$$
(6)

where

$$L(\boldsymbol{y}, \boldsymbol{A}\boldsymbol{x}; \tau) \triangleq \frac{1}{2} \sum_{i \in S_{ns}} \left(\frac{y_i - \boldsymbol{A}^i \boldsymbol{x}}{\sigma} \right)^2$$
$$- \sum_{i \in S^+} \log \left(1 - \Phi((\tau - \boldsymbol{A}^i \boldsymbol{x}) / \sigma) \right) - \sum_{i \in S^-} \log \left(\Phi((-\tau - \boldsymbol{A}^i \boldsymbol{x}) / \sigma) \right)$$

The first term in $L(\mathbf{y}, \mathbf{Ax}; \tau)$ is due to the Gaussian noise in the unsaturated measurements; the second (third) term encourages the values of $\mathbf{A}^{i}\mathbf{x}$, i.e., the members of S^{+} (likewise S^{-}) to be much greater than τ (likewise much less than $-\tau$). To understand the behaviour of the second term of $L(\mathbf{y}, \mathbf{Ax}; \tau)$, consider a measurement y_i such that $i \in S^{+}$. Referring to Eqn. 1, we have $P(y_i \ge \tau) = P(\eta_i \ge$ $\tau - \mathbf{A}^{i}\mathbf{x}) = 1 - \Phi((\tau - \mathbf{A}^{i}\mathbf{x})/\sigma)$. The last equality is due to the Gaussian nature of η_i . Given such a measurement, we seek to find \mathbf{x} such that $\mathbf{A}^{i}\mathbf{x} > \tau$, which will push $\tau - \mathbf{A}^{i}\mathbf{x}$ toward $-\infty$, i.e., push $\Phi((\tau - \mathbf{A}^{i}\mathbf{x})/\sigma)$ toward 0, and thus reduce the cost function. A similar argument can be made for the third term involving S^- . Consider that $P(y_i < -\tau) = P(\eta_i < -\tau - \mathbf{A}^i \mathbf{x}) = \Phi((-\tau - \mathbf{A}^i \mathbf{x})/\sigma).$ We seek to find x, which will tend to push $-\tau - A^i x$ toward $+\infty$, i.e., push $\Phi((-\tau - A^i x)/\sigma)$ toward 1, and thereby reduce the cost function. Assuming independence of the measurements, note that $L(y, Ax; \tau)$ is essentially the negative log of the following likelihood function:

$$\tilde{L}(\boldsymbol{y}, \boldsymbol{A}\boldsymbol{x}; \tau) \triangleq \prod_{i \in S_{ns}} \frac{e^{-(y_i - \boldsymbol{A}^i \boldsymbol{x})^2 / (2\sigma^2)}}{\sigma \sqrt{2\pi}} \quad (7)$$
$$\prod_{i \in S^+} [1 - \Phi((\tau - \boldsymbol{A}^i \boldsymbol{x}) / \sigma)] \prod_{i \in S^-} \Phi((-\tau - \boldsymbol{A}^i \boldsymbol{x}) / \sigma).$$

We henceforth term our technique 'likelihood maximization' or LM. The tendency to push $\Phi((\tau - A^i x)/\sigma)$ toward 0 or to push $\Phi((-\tau - A^i x)/\sigma)$ toward 1, is counter-balanced by the sparsity-promoting term $||\mathbf{x}||_1$, with λ deciding the relative weightage.

2.2. Theoretical Analysis

We now state an important property of $L(y, Ax; \tau)$, proved in the supplemental material [1].

Theorem 1: $L(y, Ax; \tau)$ is a convex function of x.

For further theoretical analysis, we present an overview of the broad framework in [10] and then adapt it meticulously for the analysis of our estimator in Eqn. 6. At first, we state definitions/results L1, D1 and T1 from [10]. We then use them to prove our results: Theorems 2, 3, 4. In the following, we denote the true (unknown) signal as x^* . We also use the notation $m_1 \triangleq |S_{ns}|, m_2 \triangleq |S^+|, m_3 \triangleq |S^-|.$

Lemma L1: (Lemma 1 of [10]): Let $\widehat{x_{\lambda}}$ be the optimum of a convex cost function $L^{g}(\boldsymbol{y};\boldsymbol{A}\boldsymbol{x}) + \lambda \|\boldsymbol{x}\|_{1}$ with a regularization parameter $\lambda \geq 2 \| \nabla L^g(\boldsymbol{y}; \boldsymbol{Ax^*}) \|_{\infty}$. Then the error vector $\Delta \triangleq \widehat{x_{\lambda}} - x^*$ belongs to the set $\mathbb{C}(G; x^*) \triangleq$ $\{\Delta | \| (\boldsymbol{x}^* - \widehat{\boldsymbol{x}_{\lambda}})_{G^c} \|_1 \leq 3 \| (\boldsymbol{x}^* - \widetilde{\boldsymbol{x}_{\lambda}})_G \|_1$, where G is the set of indices of the s non-zero elements of x, G^c is its complement and $\forall i \in G, x_G(i) = x_i > 0; \forall i \notin G, x_G(i) = 0. \blacksquare$

Definition D1: A loss function L is said to obey the restricted strong convexity (RSC) property with curvature $\kappa_L > 0$ and tolerance function $\tau_L(\boldsymbol{x^*})$ if the Bregman divergence $\delta L^g(\boldsymbol{\Delta}, \boldsymbol{x^*}) \triangleq L^g(\boldsymbol{y}; \boldsymbol{A} \widehat{\boldsymbol{x_{\lambda}}}) - L^g(\boldsymbol{y}; \boldsymbol{A} \boldsymbol{x^*}) \nabla L^g(\boldsymbol{y}; \boldsymbol{Ax^*})^t(\boldsymbol{\Delta})$ (the error between the loss function value at $\widehat{x_{\lambda}}$ and its first order Taylor series expansion about x^*) satisfies $\delta L^g(\boldsymbol{\Delta}, \boldsymbol{x^*}) \geq \kappa_L \|\boldsymbol{\Delta}\|_2^2 - \tau_L^2(\boldsymbol{x^*})$ for every vector $\Delta \in \mathbb{C}(G; x^*)$.

Intuitively, a loss function that obeys RSC is sharply curved around x^* , so that any difference in the loss function

 $|L^{g}(\boldsymbol{y};\boldsymbol{A}\boldsymbol{x}^{*}) - L^{g}(\boldsymbol{y};\boldsymbol{A}\widehat{\boldsymbol{x}_{\lambda}})|$ will imply a proportional estimation error $\|x^* - \widehat{x_{\lambda}}\|_1$ for all error vectors $\widehat{x_{\lambda}} - x^* \in$ $\mathbb{C}(G; \boldsymbol{x}^*)$. This property is an important sufficient condition for the loss function to obey (in relation to the regularizer), to allow for establishment of performance upper bounds. We refer the reader to [10] for more details.

Theorem T1: (Theorem 1 of [10]) If L^g is convex, differentiable and obeys RSC property with curvature κ_L and

tolerance $\tau_L^2(\boldsymbol{x}^*)$, if $\widehat{\boldsymbol{x}_{\lambda}}$ is as defined in Lemma L1 with $\lambda \geq 2 \|\nabla L(\boldsymbol{y}; \boldsymbol{A}\boldsymbol{x}^*)\|_{\infty}$, and if \boldsymbol{x}^* is an *s*-sparse vector, then we have: $\|\widehat{\boldsymbol{x}_{\lambda}} - \boldsymbol{x}^*\|_2^2 \le \frac{9\lambda^2 s}{\kappa_L^2} + \frac{2\lambda\tau_L^2(\boldsymbol{x})}{\kappa_L}$.

We now state the following theorems pertaining to the cost function in Eqn. 6 and prove them in [1]:

Theorem 2: $L(\boldsymbol{y}, \boldsymbol{A}\boldsymbol{x}^*; \tau)$ from Eqn. 6 follows RSC with curvature $\kappa_L = \frac{\gamma}{2\sigma^2}$ and tolerance function $\tau_L^2(\boldsymbol{x}^*) = 0$, where γ is the restricted eigenvalue constant (REC) for A. Here, we use the structure of $\delta L^g(\Delta, x^*)$ defined in **D1** to find the values of curvature and tolerance function for our cost function.

Theorem 3: For our noise model and with additional constraints on the signal that $\forall i, \alpha \leq x_i \leq \beta$, we have the lower bound $\|\nabla L\|_{\infty} \ge \frac{\sqrt{\rho \log(n)}}{\sigma \sqrt{m}} \left(\sqrt{m_3} + C_1 \sqrt{m_1 + m_2}\right)$ with probability $1 - 2 \exp{-\frac{1}{2}(\varrho - 2) \log(n)}$ where C_1 is a constant depending only on $\bar{\alpha}$ and β , and $\varrho > 2$ is a constant. We develop this lower bound for $\|\nabla L\|_{\infty}$ so that we can apply T1 to find the upper bound on the reconstruction error in Theorem 4, our main result.

Theorem 4: Let $\widehat{x_{\lambda}}$ be the minimizer of the cost function in Eqn. 6 with regularization parameter $\lambda \geq 2 \|\nabla L\|_{\infty}$ and with the signal constraints from Thm. 3. Let x^* be the true s-sparse signal which gave rise to the compressive measurements in y. Then we have the following upper bound with the same probability as in Thm. 3: $144 \operatorname{slop}(n) \sigma^2 a$

$$\widehat{x_{\lambda}} - x^* \|_2^2 \le \frac{1443 \log(n) \delta - p}{\gamma^2 m} (\sqrt{m_3} + C_1 \sqrt{m_1 + m_2})^2.$$

Observations related to the upper bound: The upper bound in Theorem 4 (the main theoretical result of this paper) is directly proportional to $s \log(n)$ which is equivalent to the proven upper bound in LASSO reconstruction [6, chapter 11] for Gaussian noise without saturation effects. So, the tightness of the upper bound on the reconstruction error of our cost function is relatively close to that of LASSO reconstruction. The bound is directly proportional to σ^2 as well as $s = \|\boldsymbol{x}^*\|_0$ and inversely proportional to $\gamma = \text{REC}(\boldsymbol{A}; s)$, all of which is very intuitive and similar to the analysis in [11, 6] with simpler noise models. The bound also becomes looser with increase in the number of saturated measurements m_1, m_2 . If there are no saturated measurements, i.e., $m_1 = m_2 = 0$, then the bound reduces to the normal LASSO bound [6], except that here we consider A with unit column norm as against column norm of m in [6]. The bound also increases with m_3 . However, it turns out that the constant factor C_1 for the $O(\sqrt{m_1 + m_2})$ term in the bounds, is very large. This is because it contains other factors of the form $\frac{\phi(z)}{\Phi(z)}$ or $\frac{\phi(z)}{1-\Phi(z)}$ where z stands for either α or β (see suppl. mat. [1]), which are both large in absolute value for large $|\alpha|, |\beta|$. Hence the $O(\sqrt{m_1 + m_2})$ term dominates over the $O(\sqrt{m_3})$ term, which is intuitive. We have made no attempts to optimize the constant factors in these bounds, but we note that they follow the overall empirical trends observed for our estimator, as seen in Sec. 3.

3. EXPERIMENTAL RESULTS

Here we report results on CS recovery using our technique LM in comparison to the following existing approaches described in Sec. 1.1: (i) Saturation rejection (SR) from [7]; (ii) Saturation Consistency (SC) from [7] with the following constraint set designed to (approximately) handle Gaussian measurement noise: $\forall i \in S^-, A^i x \leq -\tau + 3\sigma$ and $\forall i \in S^+, A^i x \geq \tau - 3\sigma$; (iii) Saturation Sparsity (SS) from [9]; (iv) Saturation Ignorance (SI), a technique which recovers x pretending there was no saturation in y; and (v) NCLM from [5]. For *all techniques* including LM, we assume knowledge of τ and thereby that of sets S^+, S^- . For LM, we did not impose the constraints $\alpha \leq x_i \leq \beta$ from Thm. 3, due to negligible impact on the results.

Experiment description: All our experiments were performed on signals of dimension n = 256 that were sparse in the 1D-DCT (discrete cosine transform) basis. The supports of the DCT coefficient vectors were chosen randomly, and each signal had a different support. The elements of the sensing matrix A were drawn i.i.d. from $\mathcal{N}(0, 1/m)$ so that A would obey RIP with high probability [4]. Gaussian noise was added to the measurements, followed by application of the saturation operator C. We define $\zeta \triangleq \sum_{i=1}^{m} |\mathbf{A}^{i} \mathbf{x}|/m$, the average absolute value of noiseless unsaturated measurements. Keeping all other parameters fixed, we studied the variation in the performance of these six techniques with regard to change in the following factors, keeping other factors constant: (A) number of measurements m; (B) signal sparsity s expressed as fraction $f_{sp} \in [0, 1]$ of signal dimension n; (C) noise standard deviation σ expressed as a fraction $f_{\sigma} \in [0,1]$ of ζ ; and (D) the percentage $f_{sat} \in [0,1]$ of the m measurements that were saturated. For the measurements experiment (i.e., (A)), m was varied in $\{30, 40, 50, ..., 250\}$ with $s = 25, f_{sat} = 0.15, f_{\sigma} = 0.1$. For the sparsity experiment (i.e., (<u>B</u>)), f_{sp} was varied in {0.05, 0.1, 0.15, 0.2} with $m = 150, f_{sat} = 0.15, f_{\sigma} = 0.1$. For the noise experiment (i.e., (<u>C</u>)), we varied f_{σ} in {0.01, 0.02, 0.04, ..., 0.2} with $m = 150, f_{sp} = 25/256, f_{sat} = 0.15$. For the saturation experiment (i.e., (D)), f_s was varied in $\{0, 5, 10, ..., 50\}/150$ with $m = 150, f_{sp} = 25/256, f_{\sigma} = 0.1$. The performance was measured using relative root-mean squared error (RRMSE) (defined as $\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2 / \|\boldsymbol{x}\|_2$ where $\hat{\boldsymbol{x}}$ is an estimate of the signal x), computed over reconstructions from 10 noise trials.

Parameter settings: For the proposed LM technique and for Ss, the regularization parameter λ was chosen using cross-validation on a set of unsaturated measurements, following the method in [14]. The size of the cross-validation set was 0.3 times the number of measurements used for reconstruction. For SR and SC, we set $\epsilon_{ns} = \sigma \sqrt{|S_{ns}|}$. For SI, we used the estimator min $||\boldsymbol{x}||_1$ s. t. $||\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}||_2 \leq \sigma \sqrt{m}$. For NCLM, the bound on $||\boldsymbol{x}||_2$ was set to be the ℓ_2 -norm of the true signal (omnisciently), and that on $||\boldsymbol{r}||_2$ was set to be a statistical estimate of the magnitude of the pre-saturated noise vector. The well-known FISTA algorithm [2] was used for



Fig. 1. Comparison of NCLM and LM for $s = 15, m = 150, n = 256, f_{sat} = 0.35, f_{sig} = 0.1.$

LM, whereas CVX was used for SS, SC, SR and SI.

Discussion: The results of these experiments are summarized in Fig. 2, and show that the proposed LM technique consistently outperforms the competing methods numerically. This behaviour is particularly observable for high f_{sat} or f_{σ} . We observed that SC outperformed SR for high f_{sat} or f_{σ} . We also note that our technique performed better than NCLM (our closest competitor) in the regime of high f_{σ} and high f_{sat} , as can be seen from Fig. 1. In Fig. 2, we also plot a scaled form of the upper bound on the reconstruction error from Theorem 4, to show that the empirical results for LM obey the broad trends predicted by Theorem 4.

4. CONCLUSION

We have presented a principled likelihood-based method for compressive signal recovery under Gaussian noise combined with saturation effects. We have proved the convexity of our estimator, derived an upper bound on its reconstruction error, and shown that it numerically outperforms competing methods. The recent work in [3] handles compressive inversion with Poisson-Gaussian-uniform quantization noise, a very realistic noise model for measurements in imaging systems. Extending the numerical simulations as well as the convexity proofs to handle saturation effects in conjunction with such a noise model is a potential avenue for future work. Another useful avenue of research would be to derive lower bounds on the reconstruction error for the presented penalized estimator.

5. REFERENCES

- S. Banerjee, R. Srivastava, and A. Rajwade. Reconstruction of sparse signals under gaussian noise and saturation. https://arxiv.org/abs/2102.03975, 2021.
- [2] A. Beck and M. Teboulle. A fast iterative shrinkage-



Fig. 2. RMSE comparison: SR (saturation rejection), SC (saturation consistency), SI (saturation ignorance), SS (saturation sparsity), the NCLM method and the proposed LM technique w.r.t. variation in number of measurements m (topmost, expmt. (<u>A</u>)), signal sparsity s (2nd from top, expmt. (<u>B</u>)), noise σ (3rd from top, expmt. (<u>C</u>)) and fraction f_{sp} of 1959 saturated measurements (bottom-most, expmt. (<u>D</u>)). A scaled form of upper bounds from Thm. 4 are also shown.

thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

- [3] P. Bohra, D. Garg, K. S. Gurumoorthy, and A. Rajwade. Variance-stabilization-based compressive inversion under poisson or poisson–gaussian noise with analytical bounds. *Inverse Problems*, 35(10), 2019.
- [4] E. Candes and M. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 2008.
- [5] S. Foucart and J. Li. Sparse recovery from inaccurate saturated measurements. *Acta Applicandae Mathematicae*, 158:49–66, 2018.
- [6] T. Hastie, R. Tibshirani, and M. Wainwright. Statistical Learning with Sparsity: The LASSO and Generalizations. CRC Press, 2015.
- [7] J. Laska, P. Boufounos, M. Davenport, and R. Baraniuk. Democracy in action: Quantization, saturation, and compressive sensing. *Applied and Computational Harmonic Analysis*, 31(3):429 – 443, 2011.
- [8] J. Laska, M. Davenport, and R. Baraniuk. Exact signal recovery from sparsely corrupted measurements through the pursuit of justice. In *IEEE Asilomar Conference on Signals, Systems and Computers*, 2009.
- [9] Z. Li, F. Wu, and J. Wright. On the systematic measurement matrix for compressed sensing in the presence of gross errors. In *IEEE Data Compression Conference*, 2010.
- [10] S. Negahban, P. Ravikumar, M. Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularisers. *Statistical Science*, 27(10), 2012.
- [11] G. Raskutti, M. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research*, 2010.
- [12] L. Reckner, F. Bach, W. Wang, and M.D. Plumbey. Sparse recovery and dictionary learning from non-linear compressive measurements. *IEEE Transactions on Signal Processing*, 67, 2019.
- [13] G. Tzagkarakis, J. Nolan, and P. Tsakalides. Compressive sensing using symmetric alpha-stable distributions for robust sparse signal reconstruction. *IEEE Transactions on Signal Processing*, 67(3), 2019.
- [14] J. Zhang, L. Chen, P. T. Boufounos, and Y. Gu. On the theoretical analysis of cross validation in compressive sensing. In *ICASSP*, 2014.