

Compressed Sensing with Quantized Measurements

Argyrios Zymnis, Stephen Boyd, and Emmanuel Candès

April 28, 2009

Abstract

We consider the problem of estimating a sparse signal from a set of quantized, Gaussian noise corrupted measurements, where each measurement corresponds to an interval of values. We give two methods for (approximately) solving this problem, each based on minimizing a differentiable convex function plus an ℓ_1 regularization term. Using a first order method developed by Yin et al, we demonstrate the performance of the methods through numerical simulation. We find that, using these methods, compressed sensing can be carried out even when the quantization is very coarse, *e.g.*, 1 or 2 bits per measurement.

1 Introduction

We consider the problem of estimating a sparse vector $x \in \mathbf{R}^n$ from a set of m noise corrupted quantized measurements, where the quantizer gives us an interval for each noise corrupted measurement. We give two methods for solving this problem, each of which reduces to solving an ℓ_1 regularized convex optimization problem of the form

$$\text{minimize } f(Ax) + \lambda \|x\|_1, \tag{1}$$

where f is a separable convex differentiable function (which depends on the method and the particular measurements), $A \in \mathbf{R}^{m \times n}$ is the measurement matrix, and λ is a positive weight chosen to control the sparsity of the estimated value of x .

We describe the two methods below, in decreasing order of sophistication. Our first method is ℓ_1 -regularized maximum likelihood estimation. When the noise is Gaussian (or any other log-concave distribution), the negative log-likelihood function for x , given the measurements, is convex, so computing the maximum likelihood estimate of x is a convex optimization problem; we then add ℓ_1 regularization to obtain a sparse estimate. The second method is quite simple: We simply use the midpoint, or centroid, of the interval, as if the measurement model were linear.

We will see that both methods work surprisingly well, with the first method sometimes outperforming the second.

The idea of ℓ_1 regularization to encourage sparsity is now well established in the signal processing and statistics communities. It is used as a signal recovery method from incomplete measurements, known as compressed (or compressive) sensing [CRT05, Tro06, CRT06, Don06]. The earliest documented use of ℓ_1 based signal recovery is in deconvolution of seismic data [CM73, TBM79]. In statistics, the idea of ℓ_1 regularization is used in the well known Lasso algorithm [Tib96] for feature selection. Other uses of ℓ_1 based methods include total variation denoising in image processing [ROF92, BC98], circuit design [VBG97, VBG98], sparse portfolio optimization [LFB07], and trend filtering [KKBG09].

Several recent papers address the problem of quantized compressed sensing. In [BB08], the authors consider the extreme case of sign (*i.e.*, 1-bit) measurements, and propose an algorithm based on minimizing an ℓ_1 -regularized one-sided quadratic function. Quantized compressed sensing, where quantization effects dominate noise effects, is considered in [JHF09]; the authors propose a variant of basis pursuit denoising, based on using an ℓ_p norm rather than an ℓ_2 norm, and prove that the algorithm performance improves with larger p . In [DPM09], an adaptation of basis pursuit denoising and subspace sampling is proposed for dealing with quantized measurements. In all of this work, the focus is on the effect of quantization; in this paper, we consider the combined affect of quantization and noise. Still, some of the methods described above, in particular the use of a one-sided quadratic penalty function, are closely related to the methods we propose here. In addition, several of these authors observed very similar results to ours, in particular, that compressed sensing can be successfully done even with very coarsely quantized measurements.

2 Setup

We assume that $z = Ax + v$, where $z \in \mathbf{R}^m$ is the noise corrupted but unquantized measurement vector, $A \in \mathbf{R}^{m \times n}$, and v_i are IID $\mathcal{N}(0, \sigma_i^2)$ noises. The quantizer for z_i is given by a function $\mathcal{Q}_i : \mathbf{R} \rightarrow \mathcal{Y}_i$, where \mathcal{Y}_i is a finite set of codewords. The quantized noise corrupted measurements are

$$y_i = \mathcal{Q}_i(z_i), \quad i = 1, \dots, m.$$

This is the same as saying that $z_i \in \mathcal{Q}_i^{-1}(y_i)$.

We will consider the case when the quantizer codewords correspond to intervals, *i.e.*, $\mathcal{Q}_i^{-1}(y_i) = [l_i, u_i)$. (Here we include the lower limit but not the upper limit; but whether the endpoints are included or not will not matter.) The values l_i and u_i are the lower and upper limits, or thresholds, associated with the particular quantized measurement y_i . We can have $l_i = -\infty$, or $u_i = \infty$, when the interval is infinite.

Thus, our m measurements tell us that

$$l \leq Ax + v \leq u,$$

where l and u are the lower and upper limits for the observed codewords. This model is very similar to the one used in [ZBG09] for quantized measurements in the context of fault estimation.

3 Methods

3.1 ℓ_1 -regularized maximum likelihood

The conditional probability of the measured codeword y_i given x is

$$p(y_i|x) = \Phi\left(\frac{-a_i^T x + u_i}{\sigma_i}\right) - \Phi\left(\frac{-a_i^T x + l_i}{\sigma_i}\right),$$

where a_i^T is the i th row of A and

$$\Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z \exp(-t^2/2) dt$$

is the cumulative distribution function of the standard normal distribution. The negative log-likelihood of y given x is given by

$$-\sum_{i=1}^m \log\left(\Phi\left(\frac{-a_i^T x + u_i}{\sigma_i}\right) - \Phi\left(\frac{-a_i^T x + l_i}{\sigma_i}\right)\right),$$

which we can express as $f_{\text{ml}}(Ax)$, where

$$f_{\text{ml}}(z) = \sum_{i=1}^m \log\left(\Phi\left(\frac{-z_i + u_i}{\sigma_i}\right) - \Phi\left(\frac{-z_i + l_i}{\sigma_i}\right)\right),$$

(This depends on the particular measurement observed through l and u .)

The negative log-likelihood function f_{ml} is a smooth convex function. This follows from concavity, with respect to the variable z , of

$$\log(\Phi(\alpha - z) - \Phi(\beta - z)) = \log \frac{1}{\sqrt{2\pi}} \int_{\beta-z}^{\alpha-z} \exp(-t^2/2) dt,$$

where $\alpha > \beta$. (This is the log of the probability that an $\mathcal{N}(z, 1)$ random variable lies in $[\beta, \alpha]$.) Concavity of φ follows from log-concavity of $\Phi(\alpha - z) - \Phi(\beta - z)$, which is the convolution of two log-concave functions (the Gaussian density and the function that is one between β and α and zero elsewhere); see, *e.g.*, [BV04, §3.5.2]. This argument shows that f_{ml} is convex for any measurement noise density that is log-concave.

We find the maximum likelihood estimate of x by minimizing $f_{\text{ml}}(Ax)$. To incorporate the sparsity prior, we add ℓ_1 regularization, and minimize $f_{\text{ml}}(Ax) + \lambda \|x\|_1$, adjusting λ to obtain the desired or assumed sparsity in x .

We can also add a prior on the vector x , and carry out maximum a posteriori probability estimation. The function

$$f_{\text{ml}}(Ax) - \log p_x(x),$$

where p_x is the prior density of x , is the negative log posterior density, plus a constant. Provided the prior density on x is log-concave, this function is convex; its minimizer gives the maximum a posteriori probability (MAP) estimate of x . Adding ℓ_1 regularization we can trade off posterior probability with sparsity in x .

3.2 ℓ_1 -regularized least squares

The second method we consider is simpler, and is based on ignoring the quantization. We simply use a real value for each quantization interval, and assume that the real value is the unquantized, but noise corrupted measurement. For the measurement y_i , we let $\tilde{y}_i \in \mathbf{R}$ be some value, independent of x , such as the midpoint or the centroid (under some distribution) of $[l_i, u_i]$. Assuming the distribution of z_i is $p_i(z)$, the centroid (or conditional mean value) is

$$\tilde{y}_i = \frac{\int_{l_i}^{u_i} wp(w) dw}{\int_{l_i}^{u_i} p(w) dw}.$$

We can then express the measurement as $z = \tilde{y} + q$, where $q \in \mathbf{R}^m$ denotes the quantization error.

Of course q is a function of $Ax + v$; but we use a standard approximation and consider q to be a random variable with zero mean and variance

$$\tilde{\sigma}^2 = \frac{\int_{l_i}^{u_i} (w - \tilde{y}_i)^2 p(w) dw}{\int_{l_i}^{u_i} p(w) dw}.$$

For the case of a uniform (assumed) distribution on z_i , we have $\tilde{\sigma}^2 = (u_i - l_i)^2/12$. See, *e.g.*, [Wid56]. Now we take the approximation one step further, and pretend that q is Gaussian. Under this approximation we have $\tilde{y} = Ax + \tilde{v}$, where $v_i \sim \mathcal{N}(0, \sigma_i^2 + \tilde{\sigma}_i^2)$. We can now use least-squares to estimate x , by minimizing the (convex quadratic) function $f_{\text{ls}}(Ax)$, where

$$f_{\text{ls}}(z) = \sum_{i=1}^m (1/2)(z_i - \tilde{y}_i)^2 / (\sigma_i^2 + \tilde{\sigma}_i^2).$$

To obtain a sparse estimate, we add ℓ_1 regularization, and minimize $f_{\text{ls}}(Ax) + \lambda \|x\|_1$. This problem is the same as the one considered in [KKL⁺07].

3.3 Penalty comparison

Figure 1 shows a comparison of the two different penalty functions used in our two methods, for a single measurement with $u = 0.2$, $l = -0.2$, $\tilde{y} = 0$, and $\sigma = 0.1$. We assume that the distribution of the unquantized measurement is uniform on $[-0.2, 0.2]$, which implies that the quantization noise standard deviation is about $\tilde{\sigma} = 0.11$. We can (loosely) interpret the penalty function for the second method as an approximation of the true maximum-likelihood penalty function.

4 A first order method

Problems of the form (1) can be solved using a variety of algorithms, including interior point methods [BV04, KKL⁺07], projected gradient methods [FNW07], Bregman iterative regularization algorithms [OBG⁺06, YOGD08], homotopy methods [HRTZ04, EHJT04], and




Figure 1: Comparison of the two penalty functions for a single measurement with $u = 0.2$, $l = -0.2$, $\tilde{y} = 0$, $\sigma = 0.1$, and $\tilde{\sigma} = 0.11$.

a first order method based on Nesterov’s work [BBC09]. Some of these methods use a homotopy or continuation algorithm, and so efficiently compute a good approximation of the regularization path, *i.e.*, the solution of problem (1) as λ varies.

We describe here a simple first order method due to Hale, Yin, and Zhang [HYZ08], which is a special case of a forward-backward splitting algorithm for solving convex problems [Tse91, Tse00]. We start from the optimality conditions for (1). Using subdifferential calculus, we obtain the following necessary and sufficient conditions for x to be optimal for (1):

$$\nabla f(Ax) \in \begin{cases} \{-\lambda_i\} & x_i > 0, \\ [-\lambda_i, +\lambda_i] & x_i = 0, \\ \{+\lambda_i\} & x_i < 0, \end{cases} \quad i = 0, \dots, n. \quad (2)$$

These optimality conditions tell us in particular that $x = 0$ is optimal for (1) if and only if

$$\lambda \geq \lambda^{\max} = \|\nabla f(0)\|_{\infty}.$$

We make use of this fact when selecting the initial value of λ for our algorithm.

From (2) we deduce that for any $\tau > 0$, x is optimal if and only if

$$x = \text{sgn}(x - \tau A^T \nabla f(Ax)) \circ (|x - \tau A^T \nabla f(Ax)| - \tau \lambda)_+, \quad (3)$$

where $\text{sgn}(\cdot)$ and $(\cdot)_+$ are the elementwise sign and nonnegative part operators respectively, and \circ denotes the Hadamard (elementwise) product.

From equation (3) we see that x is optimal if and only if it is a fixed point of the following iteration:

$$x^{k+1} = \text{sgn}(x^k - \tau A^T \nabla f(Ax^k)) \circ (|x^k - \tau A^T \nabla f(Ax^{k+1})| - \tau \lambda)_+. \quad (4)$$

In [HYZ08], the authors prove that this iteration converges to an optimal point of problem (1), starting from an arbitrary point $x^0 \in \mathbf{R}^n$, as long as $\lambda_{\max}(A^T \nabla^2 f(Ax)A)$, *i.e.*, the largest eigenvalue of the Hessian of $f(Ax)$, is bounded.

The fixed point continuation method is summarized below:

given tolerance $\epsilon > 0$, parameters $\tau > 0$, $0 < \beta < 1$

initialize $x_0 := 0$, $\tilde{\lambda} := \lambda^{\max}$, $k := 0$

while $\tilde{\lambda} > \lambda$

$\tilde{\lambda} := \beta \tilde{\lambda}$

while $\|x_k - x_{k-1}\|_2 > \epsilon \|x_{k-1}\|_2$

$k := k + 1$

$x_k := x_{k-1} - \tau A^T \nabla f(Ax_{k-1})$

$x_k := \text{sgn}(x_{k-1}) \circ (|x_{k-1}| - \tau \lambda)_+$

For more details about this algorithm, as well as a convergence proof see [HYZ08].

For completeness we give $\nabla f(z)$ for each of the two penalties that we consider. For the negative log-likelihood we have

$$\nabla f_{\text{ml}}(z)_i = \frac{\exp(-\tilde{u}_i^2/2) - \exp(-\tilde{l}_i^2/2)}{\sigma_i \int_{\tilde{l}_i}^{\tilde{u}_i} \exp(-t^2/2) dt},$$

where $\tilde{u}_i = (z_i - u_i)/\sigma_i$, $\tilde{l}_i = (z_i - l_i)/\sigma_i$, $i = 1, \dots, m$. For the quadratic penalty we have

$$\nabla f_{\text{ls}}(z)_i = (z_i - \bar{y}_i)/(\sigma_i + \tilde{\sigma}_i), \quad i = 1, \dots, m.$$

We found that parameter values of $\tau = 1/\lambda_{\max}(A^T A)$, $\epsilon = 10^{-4}$, and $\beta = 0.5$ work well for a large number of problems.

5 Numerical results

We now look at a numerical example with $n = 250$ variables and $m = 500$ measurements. For all our simulations, we use a fixed matrix A whose elements are drawn randomly from a $\mathcal{N}(0, 0.004)$ distribution. For each individual simulation run we choose the elements of x randomly with

$$x_i = \begin{cases} 1, & \text{with probability } 0.05, \\ -1, & \text{with probability } 0.05, \\ 0, & \text{with probability } 0.90. \end{cases}$$

Thus $a_i^T x$ has zero mean and standard deviation 0.316. The noise standard deviation is $\sigma_i = 0.1$ for all i , so the signal to noise ratio of each unquantized measurement is about 3.16.

We consider a number of possible measurement scenarios. We vary b , the number of quantization bins used from, 2 to 22 and m , the number of measurements, from 50 to 500. We choose the bin thresholds so as to make each bin have approximately equal probability, assuming a Gaussian distribution.

For each estimation scenario we use each of the two penalty functions described in §3. For the case of ℓ_1 -regularized least squares we set the approximation \bar{y}_i to be the approximate centroid of $[l_i, u_i)$, assuming z_i has $\mathcal{N}(0, 0.1)$ distribution (which it does not).

We generate 100 random instances of x and y , while keeping A fixed, and we record the average percentage of true positive and false positive in the sparsity pattern of the resulting estimate of x . Our results are summarized in figure 2, which shows the true positive rate, and figure 3, which shows the false positive rate, as a function of b and m for both methods.

What these figures show is that there is a large region in the (b, m) space in which we get very good estimation performance, *i.e.*, a high likelihood of getting the correct sparsity pattern in x . For more than 150 measurements and around 10 quantization bins (corresponding to a bit more than 3 bits per measurement), we can estimate the sparsity pattern of x quite accurately. Both methods perform well in this region. This agrees with the theoretical results on compressed sensing which state that each nonzero entry in x requires about 4–5 samples for accurate estimation (which translates to an m of around 100–125 for our experiments).

From the contour lines we can also see that ℓ_1 -regularized maximum likelihood outperforms the more naive ℓ_1 -regularized least squares, both in terms of true positive rate and false positive rate, especially when we have a large number of coarsely quantized measurements (*i.e.*, small b and large m) or a small number of finely quantized measurements (*i.e.*, large b and small m). This is an accordance with the results in [JHF09] where the authors show that a more sophisticated method outperforms the simple basis pursuit denoising method when m is large compared to the sparsity of x .

6 Conclusions

We have presented two methods for carrying out compressed sensing with quantized measurements: ℓ_1 -regularized maximum likelihood, and a more naive method based on ℓ_1 -regularized least squares. Numerical simulations show that both methods work relatively well, with the first method outperforming the second one for coarsely quantized measurements. Other authors (*e.g.*, [JHF09, DPM09, BB08]) have already observed that compressed sensing can be carried out with very coarsely quantized measurements, in cases when the cases in which the quantization effects dominate the noise; our conclusion is that the combined effects of noise and coarse quantization can be simultaneously handled.




Figure 2: True positive rate as a function of m and b for ℓ_1 -ML (left) and ℓ_1 -LS (right).




Figure 3: False positive rate as a function of m and b for ℓ_1 -ML (left) and ℓ_1 -LS (right).

References

- [BB08] P. T. Boufounos and R. G. Baraniuk. 1-bit compressive sensing. In *42nd Annual Conference on Information Sciences and Systems (CISS)*, pages 16–21, 2008.
- [BBC09] S. Becker, J. Bobin, and E. Candès. NESTA: A fast and accurate first-order method for sparse recovery, 2009. Submitted for publication.
- [BC98] P. Blomgren and T. Chan. Color TV: Total variation methods for restoration of vector-valued images. *IEEE Transactions on Image Processing*, 7(3):304–309, 1998.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [CM73] J. Claerbout and F. Muir. Robust modeling with erratic data. *Geophysics*, 38:826, 1973.
- [CRT05] E. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2005.
- [CRT06] E. Candès, 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, 2006.
- [Don06] D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [DPM09] W. Dai, H. V. Pham, and O. Milenkovic. Distortion-rate functions for quantized compressive sensing, 2009. Arxiv preprint: <http://arxiv.org/abs/0901.0749>.
- [EHJT04] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of statistics*, pages 407–451, 2004.
- [FNW07] M. Figueiredo, R. Nowak, and S. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE Journal of Selected Topics in Signal Processing*, 1(4):586, 2007.
- [HRTZ04] T. Hastie, S. Rosset, R. Tibshirani, and J. Zhu. The entire regularization path for the support vector machine. *The Journal of Machine Learning Research*, 5:1391–1415, 2004.
- [HYZ08] E. Hale, W. Yin, and Y. Zhang. Fixed-point continuation for ℓ_1 -minimization: Methodology and convergence. *SIAM Journal on Optimization*, 19:1107–1130, 2008.

- [JHF09] L. Jacques, D. K. Hammond, and M. J. Fadili. Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine, 2009. Arxiv preprint: <http://arxiv.org/abs/0902.2367>.
- [KKBG09] S.-J. Kim, K. Koh, S. Boyd, and D. Gorinevsky. ℓ_1 trend filtering. *SIAM Review*, 51(2), 2009.
- [KKL⁺07] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky. An interior-point method for large-scale ℓ_1 -regularized least squares. *IEEE Journal of Selected Topics in Signal Processing*, 1(4):606–617, 2007.
- [LFB07] M. Lobo, M. Fazel, and S. Boyd. Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research*, 152(1):341–365, 2007.
- [OBG⁺06] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An iterative regularization method for total variation-based image restoration. *Multiscale Modeling and Simulation*, 4(2):460–489, 2006.
- [ROF92] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1-4):259–268, 1992.
- [TBM79] H. Taylor, S. Banks, and J. McCoy. Deconvolution with the ℓ_1 norm. *Geophysics*, 44(1):39–52, 1979.
- [Tib96] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society*, 58(1):267–288, 1996.
- [Tro06] J. Tropp. Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE Transactions on Information Theory*, 52(3):1030–1051, 2006.
- [Tse91] P. Tseng. Applications of a splitting algorithm to decomposition in convex programming and variational inequalities. *SIAM Journal on Control and Optimization*, 29:119–138, 1991.
- [Tse00] P. Tseng. A modified forward-backward splitting method for maximal monotone mappings. *SIAM Journal on Control and Optimization*, 38(2):431–446, 2000.
- [VBG97] L. Vandenberghe, S. Boyd, and A. El Gamal. Optimal wire and transistor sizing for circuits with non-tree topology. In *IEEE/ACM international conference on Computer-aided design*, pages 252–259, 1997.
- [VBG98] L. Vandenberghe, S. Boyd, and A. El Gamal. Optimizing dominant time constant in RC circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17(2):110–125, 1998.
- [Wid56] B. Widrow. A study of rough amplitude quantization by means of Nyquist sampling theory. *IRE Transactions on Circuit Theory*, 3(4):266–276, 1956.

- [YOGD08] W. Yin, S. Osher, D. Goldfarb, and J. Darbon. Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing. *SIAM Journal of Imaging Science*, 1(1):143–168, 2008.
- [ZBG09] A. Zymnis, S. Boyd, and D. Gorinevsky. Relaxed maximum a posteriori fault identification. *Signal Processing*, 89(6):989–999, 2009.