Iterative Hard Thresholding and L_0 Regularisation

Thomas Blumensath, M. Yaghoobi and Michael E. Davies

ITERATIVE HARD THRESHOLDING AND L_0 REGULARISATION

T. Blumensath, M. Yaghoobi and M. E. Davies

IDCOM & Joint Research Institute for Signal and Image Processing
Edinburgh University
King's Buildings, Mayfield Road
Edinburgh, EH9 3JL, UK

ABSTRACT

Sparse signal approximations are approximations that use only a small number of elementary waveforms to describe a signal. In this paper we proof the convergence of an iterative hard thresholding algorithm and show, that the fixed points of that algorithm are local minima of the sparse approximation cost function, which measures both, the reconstruction error and the number of elements in the representation. Simulation results suggest that the algorithm is comparable in performance to a commonly used alternative method.

Index Terms— Sparse Approximations, Iterative Thresholding, L_0 Regularisation.

1. INTRODUCTION

Sparse signal approximations have over the last decade gained in popularity in the signal processing community and a wide range of applications such as source coding, denoising, source separation and pattern analysis have benefited from progress made in this area. The problem of sparse signal approximation is to the solution of the linear equation:

$$\mathbf{x} = \sum_i \phi_i y_i + \mathbf{e},$$

where $\mathbf{x}$ is the signal of interest and $\{\phi_i\}$ a set of elements, commonly called the dictionary. For convenience we use the linear operator defined as $\Phi Y = \sum_i \phi_i y_i$. We assume $Y \in \mathcal{H}$, where $\mathcal{H}$ is a Hilbert space. For a given $\mathbf{x}$ and Φ the sparse approximation problem is to find coefficient set $Y = \{y_i\}$ minimising the cost function:

$$C(Y) = \|\mathbf{x} - \Phi Y\|_2^2 + \lambda \|Y\|_0, \quad (1)$$

where $\|Y\|_0$ is defined as $|\Gamma_1(Y)|$, where $\Gamma_1(Y) = \{y_i : y_1 \neq 0\}$ is the set of non-zero coefficients and $|\Gamma_1(Y)|$ is the size of this set. $\|Y\|_0$ is therefore the number of non-zero coefficients.

The problem of minimising equation (1) for general $\mathbf{x}$ and Φ is known to be an NP-hard optimisation problem [1]. Therefore, two common themes have been adopted to approximately solve the problem, greedy optimisation strategies and relaxation of the cost function. Greedy strategies, such as Matching Pursuit (MP) type algorithms [2], are iterative procedures and possess no guarantee of optimising the above cost function, however, they are often relatively fast and have therefore been used extensively in practical applications. Relaxation methods, which replace the $\|Y\|_0$ constraint by a differentiable cost function, such as the FOCUSS algorithm [3] or the Basis Pursuit Denoising method [4] often offer better performance, but are computationally demanding.

Recently, iterative thresholding algorithms have been proposed for convex regularisations of the above problem as in [5] and [6], while in [7] an iterative hard thresholding algorithm was proposed. Unfortunately, the algorithm as proposed in [7] is in general unstable and no convergence analysis was presented. This is the task we set out to achieve in this paper. What is more, we also show, that the iterative hard thresholding algorithm does decrease the above cost function in each step and is guaranteed, under conditions on the singular values of Φ to converge to a local minimum of equation (1).

The paper is organised as follows. In the next section we propose a surrogate objective function replacing equation (1). The minimum of this surrogate function is shown to be achieved using hard thresholding of a Landweber iteration. Iterating this step is shown to decrease equation (1). Section 3 presents a convergence proof that guarantees that the algorithm converges. Finally, section 4 presents an experimental evaluation of the algorithm.

This research was supported by EPSRC grant D000246/1. MED acknowledges support of his position from the Scottish Funding Council and their support of the Joint Research Institute with the Heriot-Watt University as a component part of the Edinburgh Research Partnership.

2. SPARSE APPROXIMATION AND OPTIMISATION TRANSFER

2.1. Optimisation Transfer

Instead of optimising equation (1), which is NP-hard, let us introduce a surrogate objective function as proposed in [9]:

$$C^S(Y, A) = \|\mathbf{x} - \Phi Y\|_2^2 + \lambda \|Y\|_0 - \|\Phi Y - \Phi A\|_2^2 + \|Y - A\|_2^2. \quad (2)$$

Note that $C(Y) = C^S(Y, Y)$. Equation (2) can be rewritten as:

$$C^S(Y, A) = \sum_i [y_i^2 - 2y_i(A_i + \phi_i^H \mathbf{x} - \phi_i^H \Phi A) + \lambda |y_i|^0] + \|\mathbf{x}\|_2^2 + \|A\|_2^2 - \|\Phi A\|_2^2,$$

where $|y_i|^0$ is one if $y_i \neq 0$ and zero otherwise. Now the y_i are decoupled. Therefore, the minimum of equation (2) can be calculated by minimising with respect to each y_i individually. To derive the minimum, we distinguish two cases, $y_i = 0$ and $y_i \neq 0$. In the first case, the element wise cost is (ignoring the constant terms) λ. In the second case the cost is (again ignoring the constant terms):

$$y_i^2 - 2y_i(A_i + \phi_i^H \mathbf{x} - \phi_i^H \Phi A),$$

the minimum of which is achieved at

$$y_i^\star = A_i + \phi_i^H \mathbf{x} - \phi_i^H \Phi A.$$

Comparing the cost for both cases (i.e $y_i = 0$ and $y_i = A_i + \phi_i^H \mathbf{x} - \phi_i^H \Phi A$) we see that the minimum of equation (2) is attained at:

$$Y = H_{\lambda^{0.5}}(A + \Phi^H(\mathbf{x} - \Phi A)),$$

where we use the element-wise hard thresholding operator:

$$H_{\lambda^{0.5}}(y_i) \begin{cases} 0 & \text{if } |y_i| \leq \lambda^{0.5} \\ y_i & \text{if } |y_i| > \lambda^{0.5}. \end{cases}$$

Note that the minimum need not be unique whenever $A_i + \phi_i^H \mathbf{x} - \phi_i^H \Phi A = \lambda^{0.5}$. However, using a strict inequality in the definition of the thresholding operator as done here guarantees a unique update.

The iterative hard thresholding algorithm is now defined as:

$$Y^{n+1} = H_{\lambda^{0.5}}(Y^n + \Phi^H(\mathbf{x} - \Phi Y^n)). \quad (3)$$

This is the same algorithm suggested in [7], however, we found that this algorithm is not stable in general. In this paper we show that a sufficient requirement for the above algorithm to converge is that the eigenvalues of the linear operator $\mathbf{I} - \Phi^H\Phi$ are $0 < \lambda(\mathbf{I} - \Phi^H\Phi) \leq 1$. Using the singular value decomposition of $\Phi = USV^H$ we can write $(\mathbf{I} - \Phi^H\Phi) = (\mathbf{I} - VS^HSV^H) = (V(I - S^HS)V^H)$, so that we can express the above requirement as a restriction on the singular values $\sigma(\Phi)$ of Φ, i.e. $\sigma(\Phi) < 1$.

2.2. Relationship Between Optimisation of the Surrogate Function and the Original Cost Function

In this subsection we prove the following lemma:

Lemma 2.1. *Let $Y^{n+1} = H_\theta(Y^n + \Phi^H(\mathbf{x} - \Phi Y^n))$. The sequences $(C(Y^n))_n$ and $(C^S(Y^{n+1}, Y^n))_n$ are non-increasing.*

Proof. Define the operator $L = \sqrt{\mathbf{I} - \Phi^H\Phi}$. Then:

$$\begin{aligned} C(Y^{n+1}) &\leq C(Y^{n+1}) + \|L(Y^{n+1} - Y^n)\|_2^2 \\ &= C^S(Y^{n+1}, Y^n) \\ &\leq C^S(Y^n, Y^n) \\ &= C(Y^n) \\ &\leq C(Y^n) + \|L(Y^n - Y^{n-1})\|_2^2 \\ &= C^S(Y^n, Y^{n-1}), \end{aligned}$$

where the first equality is the definition of C^S and the second inequality is due the fact the Y^{n+1} is the minimiser of $C^S(Y, Y^n)$. □

3. CONVERGENCE PROOF

We have shown in the previous section that the iterative hard thresholding algorithm is guaranteed not to increase the cost function in equation (1). In this section we prove an even more interesting property of the algorithm, namely, the algorithm converges to a local minimum of equation (1).

More formally, we have the following theorem:

Theorem 3.1. *Assume $Y \in \mathcal{H}$, where $\mathcal{H}$ is a Hilbert space. If $C(Y^0) < \infty$ and if the eigenvalues of the operator L^2 obey $0 < \lambda(\mathbf{I} - \Phi^H\Phi) \leq 1$, then the sequence $(Y^n)_n$ defined by the iterative procedure in equation (3) converges to a local minimum of equation (1).*

Note that the condition $C(Y^0) < \infty$, which we use in lemma 3.2 is only really of importace in infinite dimensional spaces, where it implies that only a finite number of y_i^0 are non-zero.

To prove theorem 3.1 we need a few lemmas. To simplify notation we introduce the non-linear operator $TY = H_\theta(Y + \Phi^H(\mathbf{x} - \Phi Y))$. [1]

Lemma 3.2. $\forall\, \epsilon > 0$, $\exists\, N$ *such that* $\forall\, n > N$, $\|Y^{n+1} - Y^n\|_2^2 \leq \epsilon$.

Proof. We show that $\sum_n^N \|Y^{n+1} - Y^n\|_2^2$ converges, which implies the lemma [8, Theorem 3.23]. This is done by showing that $\sum_n^N \|Y^{n+1} - Y^n\|_2^2$ is monotonically increasing and bounded. We have monotonicity by:

$$\sum_{n=1}^{N-1} \|Y^{n+1} - Y^n\|_2^2 + \|Y^{N+1} - Y^N\|_2^2 \geq \sum_{n=1}^{N-1} \|Y^{n+1} - Y^n\|_2^2.$$

and boundedness follows from:

$$\begin{aligned} \sum_{n=0}^{N} \|Y^{n+1} - Y^n\|_2^2 &\leq \frac{1}{c}\sum_{n=0}^{N} \|L(Y^{n+1} - Y^n)\|_2^2 \quad (4) \\ &\leq \frac{1}{c}\sum_{n=0}^{N} [C(Y^n) - C(Y^{n+1})] \\ &= \frac{1}{c}(C(Y^0) - C(Y^{N+1})) \\ &\leq \frac{1}{c}C(Y^0), \end{aligned} \quad (5)$$

where c is a lower bound on the spectrum of the linear operator L^HL where we use $L = \sqrt{\mathbf{I} - \Phi^H\Phi}$, which by assumption is strictly greater than zero. $\|L(Y^{n+1} - Y^n)\|_2^2 \leq C(Y^n) - C(Y^{n+1})$ (see proof of Lemma 2.1) is here used to derive the second inequality. □

We also need the following fixed point condition:

Lemma 3.3. *Let ϕ_i^H be the i^{th} row of Φ^H and define the sets $\Gamma_0 = \{i : y_i^\star = 0\}$ and $\Gamma_1 = \{i : y_i^\star > \lambda^{0.5}\}$. Then at a fixed point of algorithm 3, i.e at points such that $Y^\star = T(Y^\star)$ we have*

$$|\phi_i^H(\mathbf{x} - \Phi Y^\star)| \begin{cases} = 0 & \text{if } i \in \Gamma_1 \\ \leq \lambda^{0.5} & \text{if } i \in \Gamma_0. \end{cases}$$

Proof. This is clear from looking at $Y^\star = T(Y^\star)$ element wise. Inserting the algorithm into the fixed point condition we have:

$$y_i^\star = H_{\lambda^{0.5}}(y_i^\star + \phi_i^H(\mathbf{x} - \Phi Y^\star)),$$

If $y_i^\star = 0$, then $|\phi_i^H(\mathbf{x} - \Phi Y^\star)| \leq \lambda^{0.5}$. Similarly for $i \in \Gamma_1$ we have :

$$y_i^\star = y_i^\star + \phi_i^H(\mathbf{x} - \Phi Y^\star),$$

where we have dropped the thresholding operator, as $y_i^\star \neq 0$. □

[1] A less formal proof could be derived based on Lemma 2.1. As $C(T(Y)) \leq C(Y)$, the cost function converges, however, $C(T(Y^\star)) = C(Y^\star)$ does not necessarily imply $T(Y^\star) = Y^\star$. By changing the algorithm to guarantee that if $C(T(Y^n)) = C(Y^n)$, then $Y^{n+1} = Y^n$, convergence can be guaranteed. However, the more formal proof presented here, shows that the algorithm converges even without such a modification.

Lemma 3.4. *A fixed point $Y^\star = TY^\star$ is a local minimum of equation (1).*

Proof. Given a fixed point $Y^\star = TY^\star$ and any small perturbation $|\partial h_i| < \epsilon$, for some $\epsilon > 0$. We show that $C(Y^\star + \partial h) > C(Y^\star)$. However, we first show $\exists\ \epsilon > 0 : \forall\ \|\partial h\| < \epsilon$ the following inequality holds:

$$C^S(Y^\star + \partial h, Y^\star) \geq C^S(Y^\star, Y^\star) + \|\partial h\|_2^2.$$

$$\begin{aligned}&C^S(Y^\star + \partial h, Y^\star) - C^S(Y^\star, Y^\star) = \\ &\sum_i (y_i + \partial h_i)^2 - 2(y_i + \partial h_i)y_i - 2(y_i + \partial h_i)(\Phi^H \mathbf{x} - \Phi^H \Phi Y^\star)_i \\ &- y_i^2 + 2y_i^2 + 2y_i(\Phi^H \mathbf{x} - \Phi^H \Phi Y^\star)_i - \lambda|y_i|^0 + \lambda|y_i + \partial h_i|^0.\end{aligned}$$

After simplification of the above equation, we split the summation into two parts, one for $\Gamma_0 = \{i : y_i = 0\}$ and one for $\Gamma_1 = \{i : y_i \neq 0\}$. We get:

$$\begin{aligned}&C^S(Y^\star + \partial h, Y^\star) - C^S(Y^\star, Y^\star) = \\ &\quad \|\partial h\|_2^2 + \sum_{\Gamma_0} \lambda|\partial h_i|^0 - 2\partial h_i(\Phi^H \mathbf{x} - \Phi^H \Phi Y^\star) \\ &\quad\quad + \sum_{\Gamma_1} -2\partial h_i(\Phi^H \mathbf{x} - \Phi^H \Phi Y^\star)\end{aligned}$$

For a fixed point $Y^\star$ the last line is zero as stated in lemma 3.3. For the summation over Γ_0 we have to consider two cases, if $\partial h_i = 0$, then this term is zero. If $\partial h_i \neq 0$, then choosing $|\partial h_i| \leq |\frac{\lambda}{2(\phi_i^H(\mathbf{x} - \Phi Y))}|$ guarantees the non-negativity of this term. Note that we also need the condition that $|\partial h_i| \leq \mathbf{y}_i$ for all $i \in \Gamma_1$ such that $y_i - h_i \neq 0$. This condition is required when splitting the cost function $|y_i - \partial h_i|^0$. Therefore $\exists\ \epsilon : \forall\ \partial h, |\partial h_i| \leq \epsilon, C^S(Y^\star + \partial h, Y^\star) \geq C^S(Y^\star, Y^\star) + \|\partial h\|_2^2$. Using this we get:

$$\begin{aligned}C(Y^\star + \partial h) &= C^S(Y^\star + \partial h, Y^\star) - \|L\partial h\|_2^2 \\ &\geq C^S(Y^\star + \partial h, Y^\star) - \|\partial h\|_2^2 \geq C^S(Y^\star, Y^\star) = C(Y^\star)\end{aligned}$$

□

Proof of theorem 3.1. In lemma 3.2 take $\epsilon < \lambda$. If $y_i^n > \lambda^{0.5}$ and $y_i^{n+1} = 0$, then $\|Y^{N+1} - Y^N\|_2^2 \geq \lambda$, which by lemma 3.2 is impossible for $n > N$ for some N. Therefore, for large N, the set of zero and non-zero coefficients will not change and $|y_i^n| > \lambda^{0.5}, \forall i \in \Gamma_1, n > N$. For $y_i^n, i \in \Gamma_1$ the algorithm then reduces to the standard Landweber algorithm with guaranteed convergence [10]. Note that the largest (smallest) eigenvalue of $(I - \Phi^H \Phi)$ will not increase (decrease) if we delete columns from Φ ensuring that the eigenvalue constraint required for the Landweber convergence is satisfied.

Also, by lemma 3.4 the fixed point is a local minimum of equation (1). □

4. EXPERIMENTAL EVALUATION

In [7], where the iterative hard thresholding approach was first suggested, results are presented that study the performance of the method in the case where the signal is known to have an exact sparse representation and the measure used is the probability of finding that representation.

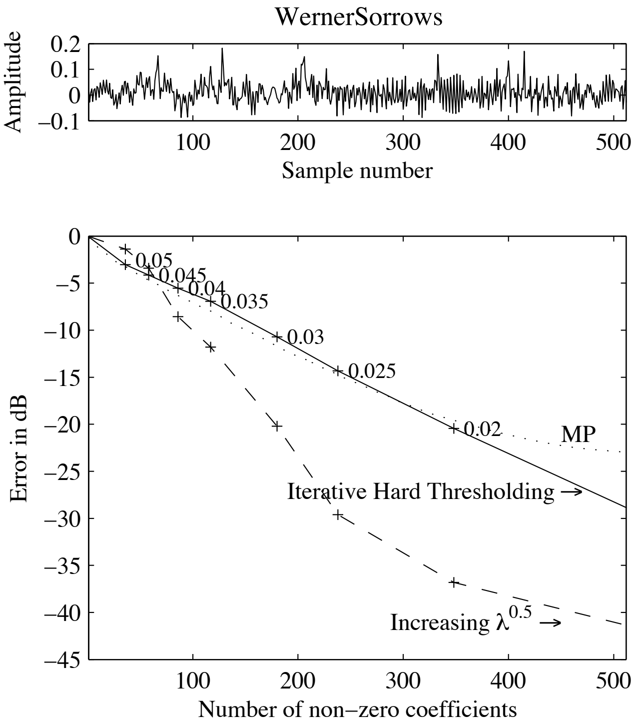

Fig. 1. The test signal (top) and the comparison of reconstruction error in dB vs. the number of non-zero coefficients used (bottom). Also shown are the different thresholds $\lambda^{0.5}$ used to calculate each of the points for the thresholding algorithm.Slowly increasing the threshold from zero to $\lambda^{0.5}$ during the first 500 iterations produced the results shown with the dashed line. The dotted line shows the results for achieved with the Matching Pursuit (MP) algorithm.

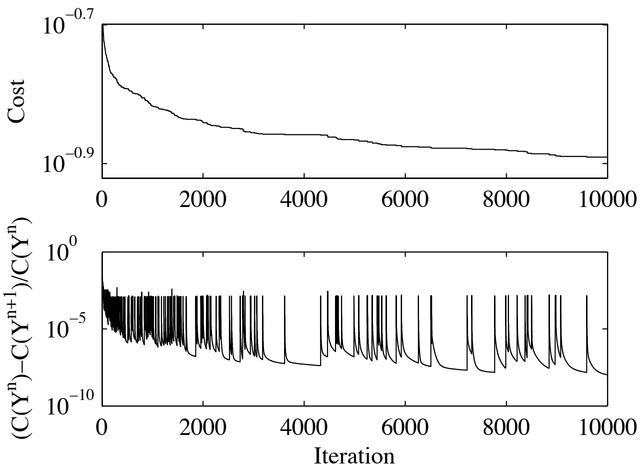

Fig. 2. Plot of the convergence of the algorithm for the threshold values of 0.015. Top plot is the cost function whilst the lower plot shows the normalised diference $\ln((C(Y^n) - C(Y^{n+1}))/C(Y^n))$.

Here we present additional and complementary results for a more general signal, for which it cannot be assumed that an exact sparse representation is available. In our experiments we used the test sig-

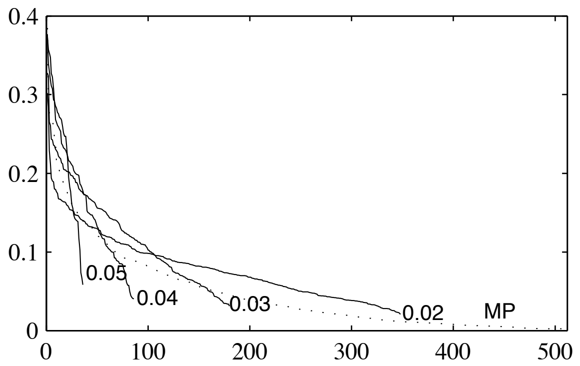

Fig. 3. Ordered magnitude of the coefficients for different thresholds. Also shown are the MP coefficients (dotted).

nals provided with the atomizer toolbox [11], as these have often been used for similar algorithms and therefore can be seen as a 'standard' benchmark. We generated test signals of length 512 and normalised the signal to unit L_2 length. We used the same toolbox to generate a six times overcomplete Wavelet packed dictionary containing Daubechies wavelets with 16 vanishing moments. The dictionary was then multiplied by a scalar, such that $\|\Phi\|_2 < 1$.

We run the algorithm for 10 different threshold values $\lambda^{0.5} \in \{0.1, 0.05, 0.045, 0.04, 0.035, 0.03, 0.025, 0.02, 0.015, 0.01\}$ and a total of 10 000 iterations. We started the algorithm each time with a zero coefficient vector. For comparison we also used the Matching Pursuit algorithm [2]. We also tried a slightly modified version of the iterative hard thresholding algorithm, in which we slowly increased the threshold from zero to $\lambda^{0.5}$ during the first 500 iterations. In all our experiments, this modification improved the results.

The results obtained for the 'WernerSorrows' signal are shown in figure 1, where we plot the L_2 reconstruction error in dB vs. the number of non-zero coefficients. The convergence is shown in figure 2 for $\lambda^{0.5} = 0.015$, where we also plot the normalised gradient $\ln \frac{(C(Y^n) - C(Y^{n+1}))}{C(Y^n)}$. The jumps in the normalised gradient are associated with changes in the set of non-zero coefficients. In figure 3 we further show the ordered magnitude of the non-zero coefficients for a range of $\lambda^{0.5}$ values.

The main conclusions to be drawn from these experiments are:

- When initialising the algorithm with a zero vector the results are comparable to Matching Pursuit.
- Convergence can be slow.
- The convergence rate depends on the threshold.
- The computational complexity of each iteration of the algorithm is comparable to the complexity of each iteration of Matching Pursuit.
- Increasing the threshold also increases the largest non-zero coefficients in the results and we found an affine relationship between the variance of the non-zero coefficients and $\lambda^{0.5}$.
- Using an adaptive threshold during the first iterations can significantly improve algorithm performance.

5. CONCLUSION

Solving the L_0 penalised optimisation problem studied in this paper is NP hard. Previous approaches either relax the L_0 penalty function (for example they use L_1 instead) or use greedy heuristics. The iterative greedy algorithm studied here is to our knowledge the first greedy algorithm that directly operates by iteratively reducing the original cost function and we have shown that this algorithm is guaranteed to find a local minimum of this cost function. Obviously, finding a local minimum is not guaranteed to find a good solution and it is also not clear, whether the algorithm is able to find the global minimum at all. For example, should the global minimum contain a coefficient that is smaller than the threshold, this solution is not reachable by the algorithm. With the current implementation it was found that the convergence of the method can be slow. This might be overcome by including an overrelaxation parameter into the update, this would however require a detailed convergence analysis. Other acceleration methods from the literature on iterative optimisation procedures might also be investigated. We found experimentally that when starting the algorithm with a zero vector the performance of the simple iterative procedure was comparable to the Matching Pursuit algorithm. However, the experiments with an increasing threshold show that additional performance increases are possible. This approach and different initialisations of the method remain to be investigated in more detail in the future.

6. REFERENCES

[1] G. Davis, *Adaptive Nonlinear Approximations*, Ph.D. thesis, New York University, 1994.

[2] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.

[3] J. F. Murray and K. Kreutz-Delgado, "An improved FOCUSS-based learning algorithm for solving sparse linear inverse problems," in *Conf. Record of the Thirty-Fifth Asilomar Conf. on Signals, Systems and Computers*, 2001, pp. 347–351.

[4] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Journal of Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1998.

[5] I. Daubechies, M. Defries, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Pure and applied Mathematics*, vol. 57, pp. 1413–1457, 2004.

[6] M. Elad, "Why simple shrinkage is still relevant for redundant representation," *to appear in IEEE Transactions on Information Theory*.

[7] K. K. Herrity, A. C. Gilbert, and J. A. Tropp, "Sparse approximation via iterative thresholding.," in *Proceedings of the Int. Conf. on Acoustics, Speech and Signal Processing*, 2006, pp. 624-627.

[8] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed. McGraw-Hill, 1976.

[9] K. Lange, D. R. Hunter, and I. Yang, "Optimization transfer using surrogate objective functions," *Journal of Computational and Graphical Statistics*, vol. 9, no. 1, pp. 1–20, Mar. 2006.

[10] L. Landweber, "An iterative formula for Fredholm integrals of the first kind," *American Journal of Mathematics*, vol. 73, no. 3, pp. 615–624, Jul. 1951.

[11] S. S. Chen, D. L. Donoho, M. A. Saunders, I. Johnston, and J. Scargle, "About atomizer," Tech. Rep., Stanford University, 1995.