

# PARNES: A RAPIDLY CONVERGENT ALGORITHM FOR ACCURATE RECOVERY OF SPARSE AND APPROXIMATELY SPARSE SIGNALS

MING GU\*, LEK-HENG LIM\*, AND CINNA JULIE WU\*

**Abstract.** In this article we propose an algorithm, PARNES, for the basis pursuit denoise problem  $\text{BP}(\sigma)$  which approximately finds a minimum one-norm solution to an underdetermined least squares problem. PARNES, (1) combines what we think are the best features of currently available methods SPGL1 [35] and NESTA [3], and (2) incorporates a new improvement that exhibits linear convergence under the assumption of the *restricted isometry property* (RIP). As with SPGL1, our approach ‘probes the Pareto frontier’ and determines a solution to the BPDN problem  $\text{BP}(\sigma)$  by exploiting the relation between the LASSO problem  $\text{LS}(\tau)$  and  $\text{BP}(\sigma)$  given by their Pareto curve. As with NESTA we rely on the accelerated proximal gradient method proposed by Yu. Nesterov [27, 26] that takes a remarkable  $O(\sqrt{L/\epsilon})$  iterations to come within  $\epsilon > 0$  of the optimal value, where  $L$  is the Lipschitz constant of the gradient of the objective function. Furthermore we introduce an ‘outer loop’ that regularly updates the prox center. Nesterov’s accelerated proximal gradient method then becomes the ‘inner loop’. The restricted isometry property together with the Lipschitz differentiability of our objective function permits us to derive a condition for switching between the inner and outer loop in a provably optimal manner. A by-product of our approach is a new algorithm for the LASSO problem that also exhibits linear convergence under RIP.

**1. Introduction.** We would like to find a solution to the sparsest recovery problem

$$\min \|x\|_0 \quad \text{s.t.} \quad Ax = b \tag{1.1}$$

where  $\|x\|_0$  is the number of nonzero entries of  $x \in \mathbb{R}^n$ . In the presence of noise we want the residual  $r := b - Ax$  to be small in the least squares sense and so we look for

$$\min \|x\|_0 \quad \text{s.t.} \quad A^\top r = 0. \tag{1.2}$$

Typically,  $A$  is an  $m$ -by- $n$  matrix with  $m \ll n$ . Note that  $A^\top r = 0$  is the normal equation. Often, the noise level will be specified by some  $\sigma$ , and  $A^\top r = 0$  will be replaced with  $\|Ax - b\| \leq \sigma$ . Both problems are known to be ill-posed and NP-hard [20, 25]; in particular, (1.1) is the subset selection problem [33].

**1.1. Motivation.** In the field of digital technology one hopes to efficiently find *digital* representations of *analog* signals, such as images and audio. The classical two stage approach consists first of collecting measurements of the analog signal. The number of measurements can be quite large depending on the desired resolution. In the case of an image, the number of measurements can be in the millions. In the second step, called *source coding*, the measurements are quantized, and the resulting signal is sparsified and compressed using an appropriate basis - e.g., Fourier, discrete cosine transform. In some cases, even sparser representations may be found by using a *redundant dictionary*. These dictionaries consist of unions of different bases. Mathematically, the aim is to have

$$f \approx Bx \tag{1.3}$$

where  $f \in \mathbb{R}^k$  is the sampled signal and  $B \in \mathbb{R}^{k \times n}$ ,  $k \leq n$ , represents a *redundant dictionary* under which there is a *compressed signal*  $x$  which is sparse. Since such dictionaries do not have a unique

---

\*Department of Mathematics, University of California at Berkeley, Berkeley, CA 94720-3840, USA (mgu, lekheng, cinnawu@math.berkeley.edu)

representation of  $f$ , they would only be useful if computationally tractable methods were available to recover a sparse  $x$ .

The method described above can be somewhat inefficient since most of the collected data is disregarded at the compression stage. In fact, for settings such as magnetic resonance imaging (MRI), it can be inconvenient to take large amounts of measurements. Thus, in such cases where measurements are difficult, one would like a method which only samples what is necessary without affecting the quality of the resulting signal. The ideas of *compressed sensing* aim to sample and compress in one stage.

Compressed sensing techniques may be useful in applications where signals are sparse in a known basis and/or measurements are expensive but computations are cheap. Some useful settings include image processing [31], seismics [24, 23], astronomy [6], and model selection in regression [17]. An interesting application is the single-pixel camera developed at Rice University [15]. Refer to Appendix 1 for a description and how it applies the ideas of compressed sensing.

To understand the basics of compressed sensing, suppose  $f \in \mathbb{R}^k$  is a signal that admits a sparse representation under the redundant dictionary  $B \in \mathbb{R}^{k \times n}$ . The idea of compressed sensing is to choose a measurement matrix  $M \in \mathbb{R}^{m \times k}$  with  $m \leq k$  which will allow us to recover the sparse signal  $x$ . In turn, this will let us approximately recover  $f = Bx$ . We have

$$b \approx Mf \quad \text{with} \quad f = Bx$$

where  $b \in \mathbb{R}^m$  is the set of measurements. In practice one would like to choose  $M$  so that  $m \ll k$ .

Setting  $A = MB$ , compressed sensing hopes to recover  $x$  by solving the optimization problem

$$\min \|x\|_0 \quad \text{s.t.} \quad \|Ax - b\|_2 \leq \sigma, \quad (1.4)$$

where  $\sigma$  is a bound on the noise level of  $b$ . Since this problem is known to be NP-hard, the various  $l_1$ -relaxed formulations described below are often used.

**1.2. Various  $l_1$ -relaxed formulations.** An  $l_1$ -relaxation is naturally motivated by the fact that  $\|\cdot\|_1$  is the largest convex underestimator of  $\|\cdot\|_0$  on unit  $l_\infty$ -ball. There has been a large amount of work done to show that this gives an effective approximation of the solution to (1.4); see [11, 14, 33]. In particular, under certain conditions on the sparsity of the solution  $x$ ,  $x$  can be recovered exactly provided that the matrix  $A$  satisfies the *restricted isometry property*; see Appendix 3. We will make this and some other reasonable assumptions on  $A$  to prove the linear convergence of PARNES.

Relaxing the zero-norm in (1.1) gives the basis pursuit problem [13]

$$\text{BP} \quad \min \|x\|_1 \quad \text{s.t.} \quad Ax = b. \quad (1.5)$$

Similarly, relaxing the zero-norm in (1.4) gives the basis pursuit denoise (BPDN) problem

$$\text{BP}(\sigma) \quad \min \|x\|_1 \quad \text{s.t.} \quad \|Ax - b\|_2 \leq \sigma. \quad (1.6)$$

The special case of  $\sigma = 0$  gives the basis pursuit problem.

Alternatively, there are two other commonly used  $l_1$ -relaxations. The penalized least-squares problem

$$\text{QP}(\lambda) \quad \min \|Ax - b\|_2^2 + \lambda \|x\|_1 \quad (1.7)$$

is one possible relaxation proposed by Chen, Donoho, and Saunders [13]. Another formulation is the LASSO problem [32] with the one-norm constrained by a parameter  $\tau$

$$\text{LS}(\tau) \quad \min \|Ax - b\|_2 \quad \text{s.t.} \quad \|x\|_1 \leq \tau. \quad (1.8)$$

In Appendix 2, we list the common  $l_1$ -relaxations of equations (1.1), (1.4), and their duals [9, 13, 32]. Note that among all the formulations,  $\text{BP}(\sigma)$  is often the most natural one in applications. This is because one often has an idea of what the noise level  $\sigma$  is, while it is not so clear what  $\lambda$  and  $\tau$  should be. Also note that the name, basis pursuit denoise, is often applied to what we call the penalized least-squares problem, originally proposed by Chen, Donoho, and Saunders in [13]. We will follow the naming convention in the SPGL1 and NESTA papers [35, 3].

Such formulations are useful since they are all convex optimization problems. In particular,  $\text{BP}$  is a linear program while  $\text{QP}(\lambda)$ ,  $\text{LS}(\tau)$ , and  $\text{BP}(\sigma)$  are quadratic programming problems. In fact, for appropriate parameter choices of  $\sigma$ ,  $\lambda$ , and  $\tau$ , the solutions of  $\text{BP}(\sigma)$ ,  $\text{QP}(\lambda)$ , and  $\text{LS}(\tau)$  coincide [30]. The exact dependence is hard to compute unless  $A$  is orthogonal [35]. However, there are solution methods which exploit these relationships. In fact, PARNES is based on the algorithm SPGL1 [35] which solves  $\text{BP}(\sigma)$  by solving a sequence of  $\text{LS}(\tau)$  problems. In SPGL1, the  $\text{LS}(\tau)$  problems are solved using a spectral gradient-projection method whereas we use a method based on the ideas in NESTA [3].

There is now a wide variety of available algorithms which solve the  $\text{BP}(\sigma)$ ,  $\text{QP}(\lambda)$ , and  $\text{LS}(\tau)$  problems. The NESTA algorithm is part of a larger class of algorithms often described as proximal gradient methods. Algorithms such as HOMOTOPY [29, 28] and LARS [17] belong to an early class of methods using a homotopy method; the solution to  $\text{BP}(\sigma)$  is found by solving  $\text{QP}(\lambda)$  for various values of  $\lambda$ . Another class consists of using Bregman iterative procedures. An example of this type of algorithm is the fixed-point continuation method, FPC [21, 22], which is tested in our experiments.

We focus on solving the relaxation  $\text{BP}(\sigma)$  where  $\sigma$  is an estimate of the noise level in the data. We present an algorithm that solves  $\text{BP}(\sigma)$  for any value of  $\sigma \geq 0$  by combining certain features of the available solvers SPGL1 and NESTA.

As mentioned above, SPGL1 approaches the solution of  $\text{BP}(\sigma)$  by solving  $\text{LS}(\tau)$  for a sequence of  $\tau$ 's. We incorporate it into PARNES since we have observed that it is suitable for large-scale problems and faster than many other solvers. In SPGL1, the  $\text{BP}(\sigma)$  problem is interpreted as finding the root of a single-variable nonlinear equation. An approximate Newton's method is used, and each iteration involves solving the LASSO problem and its dual.

The algorithm NESTA solves the  $\text{BP}(\sigma)$  problem based on a method by Nesterov which has been shown to achieve an optimal convergence rate [27, 26]. Combined with continuation techniques [21, 22], the algorithm is experimentally shown to be accurate, robust and comparable in speed to SPGL1. In PARNES, we use the methods of Nesterov to approximately solve the LASSO problem in each iteration of our Newton's iteration.

**1.3. Organization of the paper.** In Sections 2 and 3, we present and describe the background of PARNES. The background includes the SPGL1 algorithm, Nesterov's method, and the one-norm projector which is also used in SPGL1. In Section 4 we make some reasonable assumptions about the minimization problem and then provide a proof of linear convergence. Lastly, we compare PARNES with current competitive solvers in a series of extensive numerical experiments.

Since PARNES is partly based on SPGL1, we follow the assumptions made in its paper [35]. For simplicity of discussion, the authors assume that  $b \in \text{range}(A)$  and  $b \neq 0$ . This allows  $\text{BP}(\sigma)$  to be feasible for all nonnegative  $\sigma$ . This is a reasonable assumption since in applications,  $A$  typically has full row rank.

**2. SPGL1.** Here we describe the root-finding algorithm SPGL1 [35] for solving  $\text{BP}(\sigma)$  problems. The approach is to view  $\text{BP}(\sigma)$  as finding a root to the single-variable nonlinear equation whose graph is called the Pareto curve. An inexact version of Newton's method is used to find the root, and at each iteration, an approximate solution to the LASSO problem,  $\text{LS}(\tau)$ , is found. Refer to [12] for more information on the inexact Newton method.

**2.1. Pareto Curve.** The points on the Pareto curve are given by  $(\|x_\tau\|_1, \varphi(\tau))$  where  $\varphi(\tau) = \|Ax_\tau - b\|_2$  and  $x_\tau$  solves  $\text{LS}(\tau)$  [7]. Thus, the Pareto curve gives the optimal trade off between the 2-norm of the residual and 1-norm of the solution to  $\text{LS}(\tau)$ . It can be shown that the Pareto curve also characterizes the optimal trade off between the 2-norm of the residual and 1-norm of the solution to  $\text{BP}(\sigma)$ .

The Pareto curve is restricted to the interval  $\tau \in [0, \tau_{\text{BP}}]$  where  $\tau_{\text{BP}}$  is the optimal objective value of  $\text{BP}$  since  $\varphi(0) = \|b\|_2 > 0$  and  $\varphi(\tau_{\text{BP}}) = 0$ . The following theorem, proved in [35], shows that the Pareto curve is convex, strictly decreasing over the interval  $\tau \in [0, \tau_{\text{BP}}]$ , and continuously differentiable for  $\tau \in (0, \tau_{\text{BP}})$ .

**THEOREM 2.1.** *The function  $\varphi$  is*

1. *convex and nonincreasing.*
2. *continuously differentiable for  $\tau \in (0, \tau_{\text{BP}})$  with  $\varphi'(\tau) = -\lambda_\tau$  where  $\lambda_\tau = \|A^T y_\tau\|_\infty$  is the optimal dual variable to  $\text{LS}(\tau)$  and  $y_\tau = r_\tau / \|r_\tau\|_2$  with  $r_\tau = Ax_\tau - b$ .*
3. *strictly decreasing and  $\|x_\tau\|_1 = \tau$  for  $\tau \in [0, \tau_{\text{BP}}]$ .*




FIG. 2.1. An example of a Pareto curve. The solid line is the Pareto curve; the dotted red lines give two iterations of Newton's method.

**2.2. Root Finding.** Since the Pareto curve characterizes the optimal trade-off for both  $\text{BP}(\sigma)$  and  $\text{LS}(\tau)$ , solving  $\text{BP}(\sigma)$  for a fixed  $\sigma$  can be interpreted as finding a root of the non-linear equation  $\varphi(\tau) = \sigma$ . The iterations consist of finding the solution to  $\text{LS}(\tau)$  for a sequence of parameters  $\tau_k \rightarrow \tau_\sigma$  where  $\tau_\sigma$  is the optimal objective value of  $\text{BP}(\sigma)$ .

Applying Newton's method to  $\varphi$  gives

$$\tau_{k+1} = \tau_k + (\sigma - \varphi(\tau_k))/\varphi'(\tau_k).$$

Since  $\varphi$  is convex, strictly decreasing and continuously differentiable,  $\tau_k \rightarrow \tau_\sigma$  superlinearly for all initial values  $\tau_0 \in (0, \tau_{\text{BP}})$  (see Proposition 1.4.1 in [4]).

Evaluating  $\varphi(\tau_k)$  involves solving a potentially large optimization problem. By Theorem 2.1,  $\varphi(\tau_k)$  is the optimal value to  $\text{LS}(\tau_k)$  and  $\varphi'(\tau_k)$  is the dual solution to  $\text{LS}(\tau_k)$ . Thus, an inexact Newton method is carried out with approximations of  $\varphi(\tau_k)$  and  $\varphi'(\tau_k)$ .

The authors of [35] have proved the following convergence result. Let  $\bar{y}_\tau$  and  $\bar{\lambda}_\tau$  be the approximations of  $y_\tau$  and  $\lambda_\tau$  as defined in Theorem 2.1. Then the duality gap at each iteration is given by

$$\eta_\tau = \|\bar{r}_\tau\|_2 - (b^T \bar{y}_\tau - \tau \bar{\lambda}_\tau).$$

**THEOREM 2.2.** *Suppose  $A$  has full rank,  $\sigma \in (0, \|b\|_2)$ , and the inexact Newton method generates a sequence  $\tau_k \rightarrow \tau_\sigma$ . If  $\eta_k := \eta_{\tau_k} \rightarrow 0$  and  $\tau_0$  is close enough to  $\tau_\sigma$ , we have*

$$|\tau_{k+1} - \tau_\sigma| = \gamma_1 \eta_k + \zeta_k |\tau_k - \tau_\sigma|,$$

where  $\zeta_k \rightarrow 0$  and  $\gamma_1$  is a positive constant.

**2.3. Solving the LASSO problem.** Approximating  $\varphi(\tau_k)$  and  $\varphi'(\tau_k)$  require approximately minimizing  $\text{LS}(\tau)$ . The algorithm SPGL1 uses a spectral projected-gradient (SPG) algorithm. The method follows the algorithm by Birgin, Martínez, and Raydan [5] and is shown to be globally convergent. The costs include evaluating  $Ax$ ,  $A^\top r$ , and a projection onto the one-norm ball  $\|x\|_1 \leq \tau$ . In PARNES, we replace this SPG algorithm with a version of Nesterov's method that we called NESTA-LASSO (cf. Algorithm 2).

**3. PARNES.** Our method builds on the following: (1) the fast  $l^1$ -projector of Duchi, Shalev-Shwartz, Singer, and Chandra [16]; (2) the accelerated proximal gradient algorithm of Nesterov [27, 26]; (3) the Pareto curve approach of van den Berg and Friedlander [35]. The pseudocode of PARNES is given in Algorithm 3; logically it depends on NESTA-LASSO in Algorithm 2.

**3.1. Nesterov's Algorithm.** Let  $Q \subseteq \mathbb{R}^n$  be a convex closed set. Let  $f : Q \rightarrow \mathbb{R}$  be convex and Lipschitz differentiable with  $L$  as the Lipschitz constant of its gradient, i.e.

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|, \quad \text{for all } x, y \in Q.$$

Nesterov's accelerated proximal gradient algorithm iteratively defines a sequence  $x_k$  as a judiciously chosen convex combination of two other sequences  $y_k$  and  $z_k$ , which are in turn solutions to two quadratic optimization problems on  $Q$ :

$$\begin{aligned} y_k &= \operatorname{argmin}_{y \in Q} \nabla f(x_k)^\top (y - x_k) + \frac{L}{2} \|y - x_k\|_2^2, \\ z_k &= \operatorname{argmin}_{z \in Q} \sum_{i=0}^k \frac{i+1}{2} [f(x_i) + \nabla f(x_i)^\top (z - x_i)] + \frac{L}{2} \|z - c\|_2^2, \\ x_k &= \frac{2}{k+3} z_k + \frac{k+1}{k+3} y_k. \end{aligned}$$

Note that Nesterov’s algorithm allows for the use of a more general *prox-function*  $d(x)$  as long as it satisfies

$$d(x) \geq \frac{\alpha}{2} \|x - c\|_2^2. \quad (3.1)$$

For simplicity we choose the right hand side of (3.1) with  $\alpha = 1$  as our prox-function throughout this paper. The  $c$  in the prox-function is called the *prox center*.

Nesterov showed that if  $x^*$  is the optimal solution to

$$\min_{x \in Q} f(x),$$

then the iterates defined above satisfy

$$f(y_k) - f(x^*) \leq \frac{2L}{(k+1)(k+2)} \|x^* - c\|_2^2 = O\left(\frac{L}{k^2}\right).$$

An implication is that the algorithm requires  $O(\sqrt{L/\varepsilon})$  iterations to bring  $f(y_k)$  to within  $\varepsilon > 0$  of the optimal value.

---

**Algorithm 1** Accelerated proximal gradient method for convex minimization

---

**Input:** function  $f$ , gradient  $\nabla f$ , Lipschitz constant  $L$ , prox center  $c$ .

**Output:**  $x^* = \operatorname{argmin}_{x \in Q} f(x)$

- 1:  $x_0$ ;
  - 2: **for**  $k = 0, 1, 2, \dots$ , **do**
  - 3:   compute  $f(x_k)$  and  $\nabla f(x_k)$ ;
  - 4:    $y_k = \operatorname{argmin}_{y \in Q} \nabla f(x_k)^\top (y - x_k) + \frac{L}{2} \|y - x_k\|_2^2$ ;
  - 5:    $z_k = \operatorname{argmin}_{z \in Q} \sum_{i=0}^k \frac{i+1}{2} [f(x_i) + \nabla f(x_i)^\top (z - x_i)] + \frac{L}{2} \|z - c\|_2^2$ ;
  - 6:    $x_k = \frac{2}{k+3} z_k + \frac{k+1}{k+3} y_k$ ;
  - 7: **end for**
- 

**3.2. NESTA-LASSO: An accelerated proximal gradient algorithm for LASSO.** While we are ultimately interested in solving the BPDN problem in (1.6), a by-product of our approach is a new algorithm for the LASSO problem in (1.8).

We apply Nesterov’s accelerated proximal gradient method in Algorithm 1 to the LASSO problem  $\text{LS}(\tau)$ , and in each iteration, we use the fast  $l^1$ -projector  $\text{proj}_1$  described in the next section. The pseudocode for this is given in Algorithm 2. We make one slight improvement to Algorithm 1, namely, we update our prox centers from time to time. In fact, we will see in Section 4 that the prox centers may be updated in an optimal fashion and that this leads to linear convergence under a suitable application of RIP (see Corollary 4.2 for details).

In our case,  $f = \frac{1}{2} \|b - Ax\|_2^2$ ,  $\nabla f = A^\top (Ax - b)$ , and  $Q$  is the one-norm ball  $\|x\|_1 \leq \tau$ . The

initial point  $x_0$  is used as the prox center  $c$ . To compute the iterate  $y_k$ , we have

$$\begin{aligned}
y_k &= \operatorname{argmin}_{\|y\|_1 \leq \tau} \nabla f(x_k)^\top (y - x_k) + \frac{L}{2} \|y - x_k\|_2^2 \\
&= \operatorname{argmin}_{\|y\|_1 \leq \tau} y^\top y - 2(x_k - \nabla f(x_k)/L)^\top y \\
&= \operatorname{argmin}_{\|y\|_1 \leq \tau} \|y - (x_k - \nabla f(x_k)/L)\|_2 \\
&= \operatorname{proj}_1(x_k - \nabla f(x_k)/L, \tau)
\end{aligned}$$

where  $\operatorname{proj}_1(c, \tau)$  returns the projection of the vector  $c$  onto the one-norm ball of radius  $\tau$ . By similar reasoning, computing  $z_k$  can be shown to be equivalent to computing

$$z_k = \operatorname{proj}_1\left(x_k - \frac{1}{L} \sum_{i=0}^k \frac{i+1}{2} \nabla f(x_i), \tau\right).$$

In lines 9 to 12 of Algorithm 2, the duality gap for  $\operatorname{LS}(\tau)$  is computed, and the stopping criterion is met when the gap is less than the tolerance  $\eta$ .

---

**Algorithm 2** NESTA-LASSO algorithm with prox center updates

---

**Input:** initial point  $x_0$ , LASSO parameter  $\tau$ , tolerance  $\eta$ .

**Output:**  $x_\tau = \operatorname{argmin}\{\|b - Ax\|_2 \mid \|x\|_1 \leq \tau\}$ .

```

1:  $h_0 = 0, r_0 = b - Ax_0, g_0 = -A^\top r_0;$ 
2: for  $j = 0, \dots, j_{\max}$ , do
3:   for  $k = 0, \dots, k_{\max}$ , do
4:      $y_k = \operatorname{proj}_1(x_k - g_k/L, \tau);$ 
5:      $h_k = h_k + \frac{k+1}{2} g_k;$ 
6:      $z_k = \operatorname{proj}_1(x_k - h_k/L, \tau);$ 
7:      $x_k = \frac{2}{k+3} z_k + \frac{k+1}{k+3} y_k;$ 
8:      $r_k = b - Ax_k;$ 
9:      $g_k = -A^\top r_k;$ 
10:     $\eta_k = \|r_k\|_2 - (b^\top r_k - \tau \|g_k\|_\infty) / \|r_k\|_2;$ 
11:    if  $\eta_k \leq e^{-2} \eta_0$  then
12:      return
13:    end if
14:  end for
15:   $x_0 = x_k;$ 
16:  if  $\eta_k \leq \eta$  then
17:    return  $x_\tau = x_0;$ 
18:  end if
19: end for

```

---

**3.3. One Projector.** The one-projector,  $\operatorname{proj}_1$ , is used twice in each iteration of Algorithm 2. We briefly describe the algorithm of Duchi, Shalev-Schwartz, Singer, and Chandra [16] for fast projection to an  $l^1$ -ball in high-dimension. A similar algorithm is presented in [35]. The algorithm costs  $O(n \log n)$  in the worst case, but it has been shown to cost much less experimentally [35].

---

**Algorithm 3** PARNES algorithm

---

**Input:** initial point  $x_0$ , BPDN parameter  $\sigma$ , tolerance  $\eta$ .

**Output:**  $x_\sigma = \operatorname{argmin}\{\|x\|_1 \mid \|Ax - b\|_2 \leq \sigma\}$

```
1:  $\tau_0 = 0, \varphi_0 = \|b\|_2, \varphi'_0 = \|A^\top b\|_\infty;$ 
2: for  $k = 0, \dots, k_{\max}$ , do
3:    $\tau_{k+1} = \tau_k + (\sigma - \varphi_k)/\varphi'_k;$ 
4:    $x_{k+1} = \text{NESTA-LASSO}(x_k, \tau_{k+1}, \eta);$ 
5:    $r_{k+1} = b - Ax_{k+1};$ 
6:    $\varphi_{k+1} = \|r_{k+1}\|_2;$ 
7:    $\varphi'_{k+1} = -\|A^\top r_{k+1}\|_\infty/\|r_{k+1}\|_2;$ 
8:   if  $\|r_{k+1}\|_2 - \sigma \leq \eta \cdot \max\{1, \|r_{k+1}\|_2\}$  then
9:     return  $x_\sigma = x_{k+1};$ 
10:  end if
11: end for
```

---

Note that the two calls to one-projector in each iteration can be reduced to one call with results in Tseng's paper [34]. However, we make no change to our algorithm due to the low cost of the one-projector.

Consider the projection of an  $n$ -vector  $c$  onto the one-norm ball  $\|x\|_1 \leq \tau$ . This is given by the minimization problem

$$\operatorname{proj}_1(c, \tau) := \operatorname{argmin}_x \|c - x\|_2 \quad \text{s.t.} \quad \|x\|_1 \leq \tau. \quad (3.2)$$

Without loss of generality, the symmetry of the one-norm ball allows us to assume that  $c$  has all nonnegative entries. Assuming the coefficients of the vector  $c$  are ordered from largest to smallest, the solution  $x^* = \operatorname{proj}_1(c, \tau)$  is given by

$$x_i^* = \max\{0, c_i - \eta\} \quad \text{with} \quad \eta = \frac{\tau - (c_1 + \dots + c_k)}{k} \quad (3.3)$$

where  $k$  is the largest index such that  $\eta \leq c_k$ . Please refer to [16, 35] for more details on the implementation.

**4. Optimality.** We describe here how we update the prox center,  $c$ , in the prox-function (3.1) to speed up our algorithm. We prove that the centers can be updated in an optimal way. Nesterov has shown [27, 26] that when his accelerated proximal gradient algorithm is applied to minimize the objective function  $f$ , the  $k$ th iterate  $x_k$  and the minimizer  $x_*$  satisfy

$$f(x_k) - f(x^*) \leq \frac{L}{k(k+1)} d(x^*) \quad (4.1)$$

where  $L$  is the Lipschitz constant for  $\nabla f$  and  $d(x)$  is the prox-function, typically chosen to be  $\frac{1}{2}\|x - c\|_2^2$  where  $c$  is the prox center. For simplicity, we have absorbed a factor of 2 into the  $L$ .

In our case,  $f(x) = \|Ax - b\|_2^2$  where the matrix  $A$  is understood to represent a dictionary of atoms, usually sampled via some measurement, i.e.  $A = MB$  where  $M$  is the measurement operator and  $B$  is the dictionary. We will assume that  $A$  satisfies the *restricted isometry property* of order  $2s$  as described in [10, 9], namely, there exists a constant  $\delta_{2s} \in (0, 1)$  such that

$$(1 - \delta_{2s})\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta_{2s})\|x\|_2^2 \quad (4.2)$$

whenever  $\|x\|_0 \leq 2s$ . This is reasonable since it has been proven in [10] that this property helps ensure that the solution to (1.4) is closely approximated by the solution to (1.6). Under this assumption and the assumption that  $x^*$  is  $s$ -sparse, it can also be shown that solution to the LASSO problem is unique. Since the one-norm ball is compact, this implies that  $x_k$  converges to  $x^*$ . Refer to Appendix 3 for more details on the RIP condition and a proof of the uniqueness of the LASSO solution.

Additionally, we make the assumption that the solution  $x^*$  is  $s$ -sparse, and  $x_k$  is  $s$ -sparse when it is sufficiently close to  $x^*$ . This is reasonable since each iteration involves projecting the current iteration onto a 1-norm ball. Due to the geometry of the projection, there is a high likelihood that our assumption will hold.

Let  $\delta = 1 - \delta_{2s}$ . It follows from (4.2) that

$$\|A(x^* - x_k)\|_2^2 + 2(x_k - x^*)A^\top(Ax^* - b) = f(x_k) - f(x^*) \geq \|A(x^* - x_k)\|_2^2 \geq \delta\|x_k - x^*\|_2^2$$

and then from (4.1) that

$$\delta\|x_k - x^*\|_2^2 \leq \frac{L}{k(k+1)}\|x^* - c\|_2^2.$$

We get

$$\|x_k - x^*\|_2 \leq \sqrt{\frac{L}{k(k+1)\delta}}\|x^* - c\|_2 \leq \frac{1}{k}\sqrt{\frac{L}{\delta}}\|c - x^*\|_2. \quad (4.3)$$

The above relation and (4.1) suggest that we can speed up the algorithm by updating the prox center,  $c$ , with the current iterate  $x_k$  every  $k$  steps. With our assumptions, we prove in the following theorem that there is an optimal number of such steps. Allow the iterates to be represented by  $x_{jk}$  where  $j$  is the number of times the prox center has been changed (the outer iteration) and  $k$  is number of iterations after the last prox center change (the inner iteration).

**THEOREM 4.1.** *Suppose  $A$  satisfies the restricted isometry property of order  $2s$ , the solution  $x^*$  is  $s$ -sparse, and the iterates  $x_{jk}$  are eventually  $s$ -sparse. Then for each  $j$ ,*

$$\|x_{jk_{\text{opt}}} - x^*\| \leq 1/e\|x_{j1} - x^*\|$$

where

$$k_{\text{opt}} = e\sqrt{\frac{L}{\delta}} \quad (4.4)$$

and  $e$  is the base of the natural logarithm. Moreover, the total number of iterations,  $j_{\text{tot}}k_{\text{opt}}$ , to get  $\|x_{jk} - x^*\|_2 \leq \varepsilon$  is minimized with this choice of  $k_{\text{opt}}$ .

*Proof.* First observe that (4.3) implies

$$\|x_{jk} - x^*\|_2 \leq \frac{1}{k}\sqrt{\frac{L}{\delta}}\|x_{j1} - x^*\|_2 \leq \left(\frac{1}{k}\sqrt{\frac{L}{\delta}}\right)^j \|x_{11} - x^*\|_2 \leq \varepsilon\|x_{11} - x^*\|_2$$

when

$$j \log \left( \frac{1}{k}\sqrt{\frac{L}{\delta}} \right) = \log \varepsilon.$$

This relation allows us to choose  $k$  to minimize the product  $jk$ . Since

$$jk = \frac{k \log \varepsilon}{\log \sqrt{L/\delta} - \log k},$$

taking derivative of the expression on the right shows that  $jk$  is minimized when

$$k_{\text{opt}} = e\sqrt{\frac{L}{\delta}}$$

where  $e$  is the base of the natural logarithm. Thus,

$$j_{\text{tot}}k_{\text{opt}} = -e\sqrt{\frac{L}{\delta}} \log \varepsilon.$$

□

The assumption that the iterates  $x_{jk}$  would eventually be  $s$ -sparse is reasonable precisely because we expect the LASSO relaxation (1.8) to be a good proxy for (1.2). In other words, we expect our solutions to (1.8) to be sparse. Of course this argument is merely meant to be a heuristic; a more rigorous justification along the lines of [14] may be possible and is in our future plans.

Let  $k_{\text{opt}}$  be as in (4.4) and for each  $j$  we perform  $k_{\text{opt}}$  inner iterations and let  $y_j = x_{jk_{\text{opt}}}$  be the output to the  $j$ th outer iteration. An immediate consequence of Theorem 4.1 is that the relative decrease after  $k_{\text{opt}}$  steps of the inner iteration in Algorithm 2 is  $e^{-2}$ , i.e.

$$\|y_j - x^*\|_2^2 \leq e^{-2} \|y_{j-1} - x^*\|_2^2$$

and in general, this is the best possible. A consequence is the following linear convergence under appropriate assumptions.

**COROLLARY 4.2.** *If  $A$  satisfies the restricted isometry property of order  $2s$ , the solution  $x^*$  is  $s$ -sparse, and the iterates  $y_j$  are eventually  $s$ -sparse, then Algorithm 2 is linearly convergent.*

In our experiments, there are some cases where updating the prox center will eventually cause the duality gap to jump to a higher value than the previous iteration. This can cause the algorithm to run for more iterations than necessary. A check is added to prevent the prox center from being updated if it no longer helps.

**5. Other solution techniques and tools.** Here we describe the solvers and tools we test PARNES with. Many of these algorithms are currently some of the most competitive algorithms for sparse reconstruction problems. In addition, they are all available online. Since the experiments in the NESTA paper [3] are extensive and test hard, realistic problems, we repeat some of their tests and follow their experimental standards in order to obtain a good comparison. The code used to run the experiments is available at <http://www.acm.caltech.edu/~nesta>. To be consistent, the parameters that we describe below are chosen to be the same as theirs. Refer to the NESTA paper for a more complete description of the experimental details. Note that the following solvers either solve BP( $\sigma$ ) or QP( $\lambda$ ).

**5.1. NESTA.** In the NESTA paper, the accuracy of NESTA is confirmed in Section 4. For each of the following test problems, NESTA is used to solve the problem BP( $\sigma$ ). Its code is available at <http://www.acm.caltech.edu/~nesta>. The parameters for NESTA are set to be

$$x_0 = A^*b, \quad \mu = 0.02,$$

where  $x_0$  is the initial guess and  $\mu$  is the smoothing parameter for the one-norm function in  $\text{BP}(\sigma)$ .

Continuation techniques are also used to speed up NESTA in [3]. Such techniques are useful when it is observed that a problem involving some parameter  $\lambda$  is faster for large  $\lambda$ , [21, 29]. Thus, the idea of continuation is to solve a sequence of problems for decreasing values of  $\lambda$ . In the case of NESTA, it can be observed that the larger the  $\mu$  value, the faster the convergence. When continuation is used in the experiments, there are four continuation steps with  $\mu_0 = \|x_0\|_\infty$  and  $\mu_t = (\mu/\mu_0)^{t/4}\mu_0$  for  $t = 1, 2, 3, 4$ .

**5.2. GPSR: Gradient Projection for Sparse Reconstruction [18].** GPSR is used to solve the penalized least-squares problem  $\text{QP}(\lambda)$ . The code is available at <http://www.lx.it.pt/~mtf/GPSR>. The problem is first recasted as a bound-constrained quadratic program (BCQP) by using a standard change of variables on  $x$ . Here,  $x = u_1 - u_2$ , and the variables are now given by  $[u_1, u_2]$  where the entries are positive. The new problem is then solved using a gradient projection (GP) algorithm. The parameters are set to the default values in the following experiments.

A version of GPSR with continuation is also tested. The number of continuation steps is 40, the variable TOLERANCEA is set to  $10^{-3}$  and the variable MINITERA is set to 1. All other parameters are set to their default values. Note that in the code used to run the experiments, the stopping criterion for the continuation steps are different from the original GPSR code.

**5.3. SpARSA: Sparse reconstruction by separable approximation [19].** SPARSA is used to minimize functions of the form  $\phi(x) = f(x) + \lambda c(x)$  where  $f$  is smooth and  $c$  is non-smooth and non-convex. The  $\text{QP}(\lambda)$  problem is a special case of functions of this form. The code for SPARSA is available at <http://www.lx.it.pt/~mtf/SpARSA>.

In a sense, SPARSA is an iterative shrinkage/thresholding algorithm like FPC. Utilizing continuation and a Brazilai-Borwein heuristic, see [1], to find step sizes, the speed of the algorithm can be increased. The number of continuation steps is set to 40 and the variable MINITERA is set to 1. All remaining variables are set to their default values.

**5.4. SPGL1: Spectral projected gradient and SPARCO [35, 36].** The code to SPGL1 is available at <http://www.cs.ubc.ca/labs/sc1/spgl1>. The parameters for our numerical experiments are set to their default values.

Due to the vast number of available and upcoming algorithms for sparse reconstruction, the authors of SPGL1 and others have created SPARCO. In SPARCO, they provide a much needed testing framework for benchmarking algorithms. It consists of a large collection of imaging, compressed sensing, and geophysics problems. Moreover, it includes a library of standard operators which can be used to create new test problems. SPARCO is implemented in MATLAB and was originally created to test the algorithm SPGL1. The toolbox is available at <http://www.cs.ubc.ca/labs/sc1/sparco>.

**5.5. FISTA: Fast iterative soft-thresholding algorithm [2].** FISTA solves  $\text{QP}(\lambda)$ . It can be thought of as a simplified version of the Nesterov algorithm described in Section 3 since it only involves two sequences of iterates instead of three. In the experiments in [3], FISTA is used to compute reference solutions due to its efficiency and accuracy. These solutions will be used in our paper to determine the accuracy of the other computed solutions. The code for FISTA is included in the code used to run the following experiments.

**5.6. FPC: Fixed point continuation [21, 22].** FPC solves the general problem  $\min_x \|x\|_1 + \lambda f(x)$  where  $f(x)$  is differentiable and convex. Note that the special case with  $f(x) = \frac{1}{2} \|Ax - b\|_2^2$  is the  $\text{QP}(\lambda)$  problem. The algorithm is available at <http://www.caam.rice.edu/~optimization/L1/fpc>.

The approach is based on the observation that the solution solves a fixed-point equation  $x = F(x)$  where the operator  $F$  is a composition of a gradient descent-like operator and a shrinkage operator. It can be shown that the algorithm has  $q$ -linear convergence and also, finite-convergence for some components of the solution. Since the parameter  $\lambda$  affects the speed of convergence, continuation techniques are used to slowly decrease  $\lambda$  for faster convergence. A more recent version of FPC, FPC-BB, uses Brazilai-Borwein steps to speed up convergence. Both versions of FPC are tested with their default parameters.

**5.7. FPC-AS: Fixed-point continuation and active set [37].** FPC-AS is an extension of FPC into a two-stage algorithm which solves  $\text{QP}(\lambda)$ . The code can be found at <http://www.caam.rice.edu/~optimization/L1/fpc>. It has been shown in [21] that applying the shrinkage operator a finite number of times yields the support and signs of the optimal solution. Thus, the first stage of FPC-AS involves applying the shrinkage operator until an active set is determined. In the second stage, the objective function is restricted to the active set and  $\|x\|_1$  is replaced by  $c^T x$  where  $c$  is the vector of signs of the active set. The constraint  $c_i \cdot x_i > 0$  is also added. Since the objective function is now smooth, many available methods can now be used to solve the problem. In the following tests, the solvers L-BFGS and conjugate gradients, CG (referred to as FPC-AS (CG)), are used. Continuation methods are used to decrease  $\lambda$  to increase speed. For experiments involving approximately sparse signals, the parameter controlling the estimated number of nonzeros is set to  $n$ , and the maximum number of subspace iterations is set to 10. The other parameters are set to their default values. All other experiments were tested with the default parameters.

**5.8. Bregman [38].** The Bregman Iterative algorithm consists of solving a sequence of  $\text{QP}(\lambda)$  problems for a fixed  $\lambda$  and updated observation vectors  $b$ . Each  $\text{QP}(\lambda)$  is solved using the Brazilai-Borwein version of FPC. Typically, very few (around four) outer iterations are needed. Code for the Bregman algorithm can be found at <http://www.caam.rice.edu/~optimization/L1/2006/10/bregman-iterative-algorithms-for.html>. All parameters are set to their default values.

**6. Numerical results.** As mentioned above, some of the algorithms we test solve  $\text{QP}(\lambda)$  and some solve  $\text{BP}(\sigma)$ . Comparing the algorithms thus require a way of finding a  $(\sigma, \lambda)$  pair which for which the solutions coincide. The tests in [3] use a two-step procedure. From the noise level,  $\sigma$  is chosen, and then SPGL1 is used to solve  $\text{BP}(\sigma)$ . The SPGL1 dual solution gives an estimate of the corresponding  $\lambda$  and then FISTA is used to compute a second  $\sigma$  corresponding to this  $\lambda$  with high accuracy. FISTA's solution will be used to determine the accuracy of the other solutions.

As in [3], we terminate PARNES at iteration  $k$  when the solution  $\hat{x}_k$  satisfies

$$\|\hat{x}_k\|_{\ell_1} \leq \|x_{\text{NES}}\|_{\ell_1} \quad \text{and} \quad \|b - A\hat{x}_k\|_{\ell_2} \leq 1.05 \|b - Ax_{\text{NES}}\|_{\ell_2}, \quad (6.1)$$

or

$$\lambda \|\hat{x}_k\|_{\ell_1} + \frac{1}{2} \|A\hat{x}_k - b\|_{\ell_2}^2 \leq \lambda \|x_{\text{NES}}\|_{\ell_1} + \frac{1}{2} \|Ax_{\text{NES}} - b\|_{\ell_2}^2, \quad (6.2)$$

where  $x_{\text{NES}}$  is NESTA's solution (using continuation). The rationale for these stopping criteria, as explained in [3], is to reduce any potential bias arising from the fact that some algorithms solve  $\text{QP}(\lambda)$ , for which (6.2) is the most natural, while others solve  $\text{BP}(\sigma)$ , for which (6.1) is the most natural. It is evident from the tables below that there is not a significant difference whether (6.1) or (6.2) is used. Note that the algorithms are recorded to not have converged (DNC) if the number of calls to  $A$  or  $A^*$  exceeds 20,000.

TABLE 6.1

Comparison of accuracy using experiments from Table 6.2. Dynamic range 100 dB,  $\sigma = 0.100$ ,  $\mu = 0.020$ , sparsity level  $s = m/5$ . Stopping rule is (6.1).

Methods	$N_A$	$\ x\ _1$	$\ Ax - b\ _2$	$\frac{\ x - x^*\ _1}{\ x^*\ _1}$	$\ x - x^*\ _\infty$	$\ x - x^*\ _2$
PARNES	632	942197.606	2.692	0.000693	8.312	46.623
NESTA	15227	942402.960	2.661	0.004124	45.753	255.778
NESTA + CT	787	942211.581	2.661	0.000812	9.317	52.729
GPSR	DNC	DNC	DNC	DNC	DNC	DNC
GPSR + CT	11737	942211.377	2.725	0.001420	15.646	90.493
SPARSA	693	942197.785	2.728	0.000783	9.094	51.839
SPGL1	504	942211.520	2.628	0.001326	14.806	84.560
FISTA	12462	942211.540	2.654	0.000363	4.358	26.014
FPC-AS	287	942210.925	2.498	0.000672	9.374	45.071
FPC-AS (CG)	361	942210.512	2.508	0.000671	9.361	45.010
FPC	9614	942211.540	2.719	0.001422	15.752	90.665
FPC-BB	1082	942209.854	2.726	0.001378	15.271	87.963
BREGMAN-BB	1408	942286.656	1.326	0.000891	9.303	52.449

In Tables 6.2 and 6.3, we repeat the experiments done in Tables 5.1 and 5.2 of [3] with PARNES. These experiments involve recovering an unknown signal that is exactly  $s$ -sparse with  $n = 262144$ ,  $m = n/8$ , and  $s = m/5$ . The experiments are performed with increasing values of the dynamic range  $d$  where  $d = 20, 40, 60, 80, 100$  dB. For each run, the measurement operator is a randomly subsampled discrete cosine transform, and the noise level is set to 0.1.

Note that the dynamic range,  $d$ , is a measure of the ratio between the largest and smallest magnitudes of the non-zero coefficients of the unknown signal. Problems with a high dynamic range occur often in applications. In these cases, high accuracy becomes important since one must be able to detect and recover lower-power signals with small amplitudes which may be obscured by high-power signals with large amplitudes.

The last two tables, Tables 6.4 and 6.5, replicate Tables 5.3 and 5.4 of [3]. There are five runs of the experiment. Each run involves an approximately sparse signals obtained from a permutation of the Haar wavelet coefficients of a  $512 \times 512$  image. The measurement vector  $b$  consists of  $m = n/8 = 512^2/8 = 32,768$  random discrete cosine measurements, and the noise level is set to 0.1. For more specific details, refer to [3]. In applications, the signal to be recovered is often approximately sparse rather than exactly sparse. Again, high accuracy is important when solving these problems.

It can be seen that NESTA + CT, SPARSA, SPGL1, PARNES, and both versions of FPC-AS perform well in the case of exactly sparse signals for all values of the dynamic range. However, in the case of approximately sparse signals, all versions of FPC and SPARSA no longer converge in 20,000 function calls. PARNES still performs well, converging in under 2000 iterations for all runs. The accuracy of the various algorithms is compared in Table 6.1.

**6.1. Choice of parameters.** As Tseng observed, accelerated proximal gradient algorithms will converge so long as the condition given as equation (45) in [34] is satisfied. In our case this

TABLE 6.2

Number of function calls where the sparsity level is  $s = m/5$  and the stopping rule is (6.1).

Method	20 dB	40 dB	60 dB	80 dB	100 dB
PARNES	122	172	214	470	632
NESTA	383	809	1639	4341	15227
NESTA + CT	483	513	583	685	787
GPSR	64	622	5030	DNC	DNC
GPSR + CT	271	219	357	1219	11737
SPARSA	323	387	465	541	693
SPGL1	58	102	191	374	504
FISTA	69	267	1020	3465	12462
FPC-AS	209	231	299	371	287
FPC-AS (CG)	253	289	375	481	361
FPC	474	386	478	1068	9614
FPC-BB	164	168	206	278	1082
BREGMAN-BB	211	223	309	455	1408

TABLE 6.3

Number of function calls where the sparsity level is  $s = m/5$  and the stopping rule is (6.2).

Method	20 dB	40 dB	60 dB	80 dB	100 dB
PARNES	74	116	166	364	562
NESTA	383	809	1639	4341	15227
NESTA + CT	483	513	583	685	787
GPSR	62	618	5026	DNC	DNC
GPSR + CT	271	219	369	1237	11775
SPARSA	323	387	463	541	689
SPGL1	43	99	185	365	488
FISTA	72	261	1002	3477	12462
FPC-AS	115	167	159	371	281
FPC-AS (CG)	142	210	198	481	355
FPC	472	386	466	1144	9734
FPC-BB	164	164	202	276	1092
BREGMAN-BB	211	223	309	455	1408

translates into

$$\min_{x \in \mathbb{R}^n} \left\{ \nabla f(y_k)^\top x + \frac{L}{2} \|x - x_k\|_2^2 + P(x) \right\} \geq \nabla f(y_k)^\top y_k + P(y_k), \quad (6.3)$$

upon setting  $\gamma_k = 1$  and

$$P(x) = \begin{cases} 0 & \text{if } \|x\|_1 \leq \tau, \\ \infty & \text{otherwise,} \end{cases}$$

in (45) in [34]. In other words, the value of  $L$  need not necessarily be fixed at the Lipschitz constant of  $\nabla f$  but may be decreased and that decreasing  $L$  has the same effect as increasing the

TABLE 6.4

Recovery results of an approximately sparse signal (with Gaussian noise of variance 1 added) and with (6.2) as a stopping rule.

Method	Run 1	Run 2	Run 3	Run 4	Run 5
PARNES	838	810	1038	1098	654
NESTA	8817	10867	9887	9093	11211
NESTA + CT	3807	3045	3047	3225	2735
GPSR	DNC	DNC	DNC	DNC	DNC
GPSR + CT	DNC	DNC	DNC	DNC	DNC
SPARSA	2143	2353	1977	1613	DNC
SPGL1	916	892	1115	1437	938
FISTA	3375	2940	2748	2538	3855
FPC-AS	DNC	DNC	DNC	DNC	DNC
FPC-AS (CG)	DNC	DNC	DNC	DNC	DNC
FPC	DNC	DNC	DNC	DNC	DNC
FPC-BB	5614	7906	5986	4652	6906
BREGMAN-BB	3288	1281	1507	2892	3104

TABLE 6.5

Recovery results of an approximately sparse signal (with Gaussian noise of variance 0.1 added) and with (6.2) as a stopping rule.

Method	Run 1	Run 2	Run 3	Run 4	Run 5
PARNES	1420	1772	1246	1008	978
NESTA	11573	10457	10705	8807	13795
NESTA + CT	7543	13655	11515	3123	2777
GPSR	DNC	DNC	DNC	DNC	DNC
GPSR + CT	DNC	DNC	DNC	DNC	DNC
SPARSA	12509	DNC	DNC	3117	DNC
SPGL1	1652	1955	2151	1311	2365
FISTA	10845	12165	10050	7647	11997
FPC-AS	DNC	DNC	DNC	DNC	DNC
FPC-AS (CG)	DNC	DNC	DNC	DNC	DNC
FPC	DNC	DNC	DNC	DNC	DNC
FPC-BB	DNC	DNC	DNC	DNC	DNC
BREGMAN-BB	3900	3684	2045	3292	3486

stepsize. Tseng suggested to decrease  $L$  adaptively by a constant factor until (45) is violated and then backtrack and repeat the iteration (cf. Note 6 in [34]). For simplicity, and very likely at the expense of speed, we do not change our  $L$  adaptively in PARNES and NESTA-LASSO. Instead, we simply choose a small fixed  $L$  by trying a few different values so that (6.3) is satisfied for all  $k$ , and likewise for the tolerance  $\eta$  in Algorithm 2. However, even with this crude way of selecting  $L$  and  $\eta$ , the results obtained are still rather encouraging.

**7. Conclusions.** As seen in the numerical results, the algorithms SPGL1 and NESTA are among some of the top performing algorithms available for basis pursuit denoise problems. We have therefore combined their best features, (1) use of the Pareto curve, and (2) Nesterov's accelerated

proximal gradient method into our algorithm PARNES. The name PARNES thus comes from a combination of the names Pareto and Nesterov. The placement of *Par* as the first part of PARNES represents that fact that our outer-most iterations consist of finding roots on the Pareto curve. Since our inner iterations consist of applying Nesterov’s methods, *Nes* takes up the second spot in PARNES. Through our experiments, we have shown that our algorithm is comparable to currently available state-of-the-art methods.

**Acknowledgments.** We would like to give special thanks to Emmanuel Candès for helpful discussions and ideas. The numerical experiments in this paper rely on the shell scripts and MATLAB codes<sup>1</sup> of Stephen Becker and Jérôme Bobin. We have also benefited from Michael Friedlander and Ewout van den Berg’s MATLAB codes<sup>2</sup> for SPGL1. We are grateful to them for generously making their codes available on the web. Finally, thanks to Rice University for the use of their single-pixel camera image.




FIG. 7.1. Illustration of the single-pixel camera. Lens 2 is the lens directly under the photon detector (PD). Courtesy of Rice University.

**Appendix 1: Single-pixel camera.** An interesting application of compressed sensing is the single-pixel camera developed at Rice University [15]. The camera uses a digital micromirror device (DMD) consisting of an array of  $k$  micromirrors to take measurements of the image  $f$ . Each micromirror corresponds to a pixel in  $f$  and can be oriented toward or away from Lens 2. Each orientation of the DMD gives a measurement of the image by reflecting the light onto Lens 2 which focuses the light on a single photon detector (the single pixel). The single pixel then sends the information to an A/D converter.

By the theory of compressed sensing, the user hopes to take  $m \ll k$  random measurements of the image using the DMD and later recover the image. The measurements satisfy the relationship

$$b \approx Mf = MBx \tag{7.1}$$

where  $b$  is the vector of measurements,  $M \in \mathbb{R}^{m \times k}$  is the measurement matrix with  $m \leq k$  and rows determined by the DMD, and  $x$  is a sparse representation of  $f$  under dictionary  $B$ . Setting  $A = MB$ , the problem of recovering  $x$  simply reduces to solving the sparse recovery problem (1.4).

<sup>1</sup>[http://www.acm.caltech.edu/~nesta/NESTA\\_ExperimentPackage.zip](http://www.acm.caltech.edu/~nesta/NESTA_ExperimentPackage.zip)

<sup>2</sup><http://www.cs.ubc.ca/labs/scl/spgl1>

**Appendix 2: Various  $l_1$ -relaxed formulations.** We give here a few common  $l_1$ -relaxations of (1.4) and their duals [9, 13, 32]. For simplicity, we write  $r = b - Ax$ .

$$\begin{array}{lll}
\text{BP} & \min \|x\|_1 & \text{s.t. } r = 0 \\
\text{QP}(\lambda) & \min \|r\|_2^2 + \lambda \|x\|_1 & \\
\text{LS}(\tau) & \min \|r\|_2 & \text{s.t. } \|x\|_1 \leq \tau \\
\text{BP}(\sigma) & \min \|x\|_1 & \text{s.t. } \|r\|_2 \leq \sigma \\
\text{DS}(\mu) & \min \|x\|_1 & \text{s.t. } \|A^\top r\|_\infty \leq \mu \\
\text{BP}^* & \min -b^\top y & \text{s.t. } \|A^\top y\|_\infty \leq 1 \\
\text{QP}^*(\lambda) & \min \|r\|_2^2 - b^\top r & \text{s.t. } \|A^\top r\|_\infty \leq \lambda \\
\text{LS}^*(\tau) & \min \tau \|A^\top r\|_\infty - b^\top r & \text{s.t. } \|x\|_1 \leq \tau \\
\text{BP}^*(\sigma) & \min \|r\|_2 - b^\top r & \text{s.t. } \|A^\top r\|_\infty \leq \sigma \\
\text{DS}^*(\mu) & \min \mu \|z\|_1 - b^\top r & \text{s.t. } \|A^\top r\|_\infty \leq \mu, \quad Az = r
\end{array} \tag{7.2}$$

The names BP( $\sigma$ ), QP( $\lambda$ ), LS( $\tau$ ), and DS( $\mu$ ) are for *basis pursuit denoise*, *penalized least squares*, *least absolute shrinkage and selection operator* (LASSO), and *Dantzig selector* respectively. The asterisked names denote the dual of the respective formulations. These are all convex optimization problems. In particular, BP, LS\*( $\tau$ ), DS( $\mu$ ), DS\*( $\mu$ ) are linear programming problems while QP( $\lambda$ ), LS( $\tau$ ), BP( $\sigma$ ), BP\*( $\sigma$ ) are quadratic programming problems.

**Appendix 3: The restricted isometry property.** Suppose we would like to recover the solution  $\tilde{x}$  to

$$\min \|x\|_0 \quad \text{s.t. } Ax = b$$

where  $\|x\|_0$  is the number of nonzero entries of  $x$  and  $A$  is underdetermined. It has been shown that the solution  $x^*$  to

$$\min \|x\|_1 \quad \text{s.t. } Ax = b$$

recovers  $\tilde{x}$  exactly if  $x$  is sparse enough and  $A$  satisfies the *restricted isometry property* [8].

**DEFINITION 7.1.** *A matrix  $A$  satisfies the restricted isometry property of order  $s$  with constant  $\delta_s \in (0, 1)$  if*

$$(1 - \delta_s)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta_s)\|x\|_2^2$$

whenever  $\|x\|_0 \leq s$ .

In [10], Candès restricts the definition above with the assumption that  $\delta_s \in (0, \sqrt{2} - 1)$ . He proves the following theorem. Let  $x_s$  be the vector  $x$  with all but the largest  $s$ -entries set to zero.

**THEOREM 7.2.** *If  $A$  satisfies the RIP of order  $2s$  for some  $\delta_s \in (0, \sqrt{2} - 1)$ , then  $x^*$  satisfies*

$$\|x^* - x\|_2 \leq C_0 s^{-1/2} \|x - x_s\|_1 \tag{7.3}$$

where  $C_0$  is a constant. If  $x$  is  $s$ -sparse, we have exact recovery.

Candès also shows that the solutions to (1.4) and (1.6) are equivalent in sense of the following theorem.

**THEOREM 7.3.** *If  $A$  satisfies the RIP of order  $2s$  for some  $\delta_s \in (0, \sqrt{2} - 1)$ , then the solution  $x^*$  to  $\text{BP}(\sigma)$  satisfies*

$$\|x^* - x\|_2 \leq C_0 s^{-1/2} \|x - x_s\|_1 + C_1 \sigma \quad (7.4)$$

where  $C_0$  and  $C_1$  are small constants.

In particular, if  $A$  satisfies the RIP of order  $2s$  for some  $\delta_s \in (0, \sqrt{2} - 1)$ , the definition implies that there exists no nonzero  $x$  such that  $x$  is  $s$ -sparse and  $Ax = 0$ . Theorem 3.2 in [28] gives the following uniqueness result about LASSO problem.

**THEOREM 7.4.** *Suppose the solution to the LASSO problem (1.8) is  $s$ -sparse. Then the solution is unique if there exists no non-zero  $s$ -sparse null vector of  $A$ .*

This implies that if  $A$  satisfies the RIP condition of order  $2s$  and the solution is  $s$ -sparse, the LASSO problem,  $\text{LS}(\tau)$ , has a unique solution.

#### REFERENCES

- [1] J. Barzilai and J. Borwein, “Two point step size gradient method,” *IMA J. Numer. Anal.*, **8** (1988), no. 1, pp. 141–148.
- [2] A. Beck and M. Teboulle, “Fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM J. Imaging Sci.*, **2** (2009), no. 1, pp. 183–202.
- [3] S. Becker, J. Bobin, and E. J. Candès, “NESTA: a fast and accurate first-order method for sparse recovery,” *preprint*, (2009).
- [4] D. P. Bertsekas, *Nonlinear Programming*, 2nd Ed., Athena Scientific, Belmont, MA, 1999.
- [5] E. Birgin, J. Martínez, and M. Raydan, “Nonmonotone spectral projected gradient methods on convex sets” *SIAM J. Optim.*, **10** (2000), no. 4, pp. 1196–1211.
- [6] J. Bobin, J.-L. Stark, and R. Ottensamer, “Compressed sensing in astronomy,” *IEEE J. Selected Top. Signal Process.*, **2** (2008), no. 5, pp. 718–726.
- [7] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2004.
- [8] E. J. Candès and T. Tao “Decoding by linear programming,” *IEEE Trans. Inform. Theory*, **51** (2005), no. 12, pp. 4203–4215.
- [9] E. J. Candès and T. Tao, “The Dantzig selector: statistical estimation when  $p$  is much larger than  $n$ ,” *Ann. Statist.*, **35** (2007), no. 6, pp. 2313–2351.
- [10] E. J. Candès, “The restricted isometry property and its implications for compressed sensing,” *C. R. Math. Acad. Sci. Paris*, **346** (2008), no. 9–10, pp. 589–592.
- [11] E. J. Candès, J. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Comm. Pure Appl. Math.*, **59** (2006), no. 8, pp. 1207–1223.
- [12] R. S. Dembo, S. C. Eisenstat, and T. Steihaug, “Inexact newton methods,” *SIAM J. Numer. Anal.*, **19** (1982), no. 2, pp. 400–408.
- [13] S. Chen, D. L. Donoho and M. Saunders, “Atomic decomposition by basis pursuit,” *SIAM J. Sci. Comput.*, **20** (1998), no. 1, pp. 33–61.
- [14] D. L. Donoho, “For most large underdetermined systems of linear equations the  $\ell_1$ -norm solution is also the sparsest solution,” *Comm. Pure Appl. Math.*, **59** (2006), no. 6, pp. 797–829.
- [15] M. Duarte, M. Davenport, D. Takhar, J. Laska, T. Sun, K. Kelly, and R. Baraniuk, “Single-pixel imaging via compressive sampling,” *IEEE Signal Process. Mag.*, **25** (2008), no. 2, pp. 83–91.
- [16] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, “Efficient projections onto the  $\ell_1$ -ball for learning,” *Proc. Int. Conf. Mach. Learn. (ICML ’08)*, **25** (2008), no. 307, pp. 272–279.
- [17] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression,” *Ann. Statist.*, **32** (2004), no. 2, pp. 407–499.
- [18] M. Figueiredo, R. Nowak, and S. Wright, “Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems,” *IEEE J. Selected Top. Signal Process.*, **1** (2007), no. 4, pp. 586–597.
- [19] M. Figueiredo, R. Nowak, and S. Wright, “Sparse reconstruction by separable approximation,” *IEEE Trans. Signal Process.*, (to appear).

- [20] M. R. Garey, and D. S. Johnson, *Computers and Intractability. A guide to the theory of NP-completeness*, W. H. Freeman, New York, NY, 1979.
- [21] E. T. Hale, W. Yin, and Y. Zhang, "A fixed-point continuation method for  $\ell_1$ -regularized minimization with applications to compressed sensing," *Rice University Technical Report*, (2007).
- [22] E. T. Hale, W. Yin, and Y. Zhang, "Fixed-point continuation for  $\ell_1$ -minimization: Methodology and convergence," *SIAM J. Optimization*, **19** (2008), no. 3, pp. 1107–1130.
- [23] G. Hennenfent and F. J. Herrmann, "Simply denoise: wavefield reconstruction via jittered undersampling," *Geophysics*, **73** (2008), no. 3, pp. V19–V28.
- [24] G. Hennenfent and F. J. Herrmann, "Sparseness-constrained data continuation with frames: Applications to missing traces and aliased signals in 2/3-D," *SEG Tech. Program Expanded Abstracts* **24**, (2005), no. 1, pp. 2162–2165.
- [25] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, **24** (1995), no. 2, pp. 227–234.
- [26] Yu. Nesterov, "A method for solving the convex programming problem with convergence rate  $O(1/k^2)$ ," *Dokl. Akad. Nauk SSSR*, **269** (1983), no. 3, pp. 543–547.
- [27] Yu. Nesterov, "Smooth minimization of non-smooth functions," *Math. Program.*, **103** (2005), no. 1, pp. 127–152.
- [28] M. R. Osborne, B. Presnell, and B. A. Turlach, "On the LASSO and its dual," *J. Comput. Graph. Statist.*, **9** (2000), no. 2, pp. 319–337.
- [29] M. R. Osborne, B. Presnell, and B. A. Turlach, "A new approach to variable selection in least squares problems," *IMA J. Numer. Anal.*, **20** (2000), no. 3, pp. 389–403.
- [30] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.
- [31] J. Romberg, "Imaging via compressive sensing," *IEEE Trans. Signal Process.*, **25** (2008), no. 2, pp. 14–20.
- [32] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *J. Roy. Statist. Soc. Ser. B*, **58** (1996), no. 1, pp. 267–288.
- [33] J. A. Tropp, "Just relax: Convex programming methods for identifying sparse signals in noise," *IEEE Trans. Inform. Theory*, **52** (2006), no. 3, pp. 1030–1051.
- [34] P. Tseng, "On accelerated proximal gradient methods for convex-concave optimization," *preprint*, (2008).
- [35] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Sci. Comput.*, **31** (2008/09), no. 2, pp. 890–912.
- [36] E. van den Berg, M. P. Friedlander, G. Hennenfent, F. J. Herrmann, R. Saab, and Ö. Yilmaz, "Algorithm 890: SPARCO: a testing framework for sparse reconstruction," *ACM Trans. Math. Software*, **35** (2009), no. 4, Art. 29, pp. 16.
- [37] Z. Wen, W. Yin, D. Goldfarb, and Y. Zhang, "A fast algorithm for sparse reconstruction based on shrinkage, subspace optimization and continuation," *preprint*, (2009).
- [38] W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman iterative algorithms for  $l_1$  minimization with applications to compressed sensing," *SIAM J. Imaging Sci.*, **1** (2008), no. 1, pp. 143–168.