UNCONSTRAINED L1 OPTIMIZATION WITH APPLICATIONS TO SIGNAL AND IMAGE PROCESSING

CARLOS ANDRES RAMIREZ VILLAMARIN

Computational Science Program



APPROVED:

Miguel Argáez, Ph.D., Chair

Sergio Cabrera, Ph.D.

Osvaldo Mendez, Ph.D.

Benjamin C. Flores, Ph.D.
Dean of the Graduate School

*A mi madre y*

*en memoria de mi padre,*

*con amor.*

UNCONSTRAINED L1 OPTIMIZATION WITH APPLICATIONS TO SIGNAL AND IMAGE PROCESSING

by

CARLOS ANDRES RAMIREZ VILLAMARIN, M.S.



DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at El Paso

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

Computational Science Program

THE UNIVERSITY OF TEXAS AT EL PASO

May 2013

# Acknowledgements

I would like to express my deep gratitude to my advisor Doctor Miguel Argáez, for his guidance, orientation, and most importantly, for giving me the opportunity to start pursuing my dream of becoming a scientist. He is a person of opportunities, he believed in me when I had first failed the TOEFL exam, when the only thing he knew about me was just that I loved mathematics. I am also profoundly indebted to Doctor Leticia Velázquez who has always been concerned with my growth as a researcher and as a person. She is also responsible for such a wonderful life-opportunity of studying in the U.S.

I am grateful with my committee members Doctor Osvaldo Mendez, who with admirable enthusiasm helped me to clarify some mathematical ideas, and specially with Doctor Sergio Cabrera who has always had an open door for me.

With all my heart, I thank all those people who were beside me during my graduate studies and offered me their friendship. I specially acknowledge my lovely girl Aline and the most important people in my life: My mom Clara Ines and my sister Liana. This degree is dedicated to them with love.

Last but not least, I gratefully acknowledge the financial support provided by REPSOL U.S.A, the College of Science at UTEP and the ARL Grant No. W911NF-07-2-0027 which have made possible my graduate studies.

# Abstract

In recent years, the applied mathematical community has witnessed a revolution that is changing the paradigm of classical signal and image processing. Novel and efficient numerical algorithms have emerged for solving new challenges in large scale signal retrieval, where both constrained and unconstrained $\ell_1$ minimization methods play a fundamental role.

In this work, we present a new methodology for solving unconstrained $\ell_1$ minimization problems in the context of image and signal processing. Our approach consists in solving a sequence of relaxed unconstrained minimization problems depending on a positive regularization parameter $\mu$ that converges to zero. The optimality conditions of each subproblem are characterized through a fixed point equation, where a preconditioned conjugate gradient algorithm is applied to solve a sequence of resulting linear systems. The preconditioner used in the conjugate gradient algorithm is designed in such a way that it exploits the structure of the induced matrix at each subproblem. Moreover, we show that the distribution of the eigenvalues in the preconditioned system is bounded regardless of the value of the regularization parameter $\mu$.

We prove global convergence for the iterative scheme derived from our method and conduct several numerical experiments showing that the proposed algorithm performs comparable with some state of the art solvers. In particular, we present numerical evidence that our algorithm is very competitive when recovering high dynamic range signals. Finally, we solve a set of real world applications including image processing problems, data and signal processing in micrometeorology, texture segmentation in seismic data and total variation.

# Table of Contents

# Chapter 1

# Introduction

## 1.1 Overview

Providing a precise and meaningful description of a phenomenon or dataset is of central interest in the scientific research community. Inspired by this task, multi-resolution analysis led to the popularization of wavelets as a tool for representing a broad range of signals in the early 90's. This effort was followed by a new area known as *sparse representation* where certain classes of signals are expressed as a linear combination of only a few elementary signals. We say that a given signal-vector $u$ has a sparse representation $x$ if it can be expressed as

$$u = \Psi x,$$

where the number of nonzero elements of the vector $x$ is by far less than its dimension, and $\Psi$ is referred to as the sparsifying matrix. We thus say that $u$ is compressible under $\Psi$, and that the vector $x$ is sparse.

In the last ten years, the sparse representation paradigm has led to the development of new approaches for solving signal and image processing problems exploiting compressibility as a prior information (see for example Donoho et al., and Figueiredo et al. [27, 34]), not to mention the improvement of current image compression-encoding standards such as JPEG and JPEG-2000. Moreover, important advances towards understanding the role of sparsity in image processing came along with the new sampling theory of compressed sensing introduced by Candès, Donoho and Tao [12, 23, 14], and with statistical image processing where non-smooth regularizers promoting sparsity are utilized under a maximum penalized likelihood estimator and/or maximum a posteriori schemes (see for example Chaux et al.

and Figueiredo et al. [28, 34]).

In practical compressed sensing problems, it is required to solve an optimization problem of the form

$$\min_x \|x\|_0 \text{ subject to } Ax = b, \tag{1.1}$$

for decoding a sparse signal $x \in \mathbb{R}^n$ that has been significantly sub-sampled by a sampling matrix $A \in \mathbb{R}^{m\times n}$ with $m < n$. Here, $\|x\|_0$ counts the number of non-zero entries of the vector $x$, and is commonly referred to as the $\ell_0$ norm.

Solving (1.1) amounts to finding the sparsest vector $x$ that simultaneously satisfies the linear system $Ax = b$. Nevertheless, finding such a vector is by nature combinatorial and generally an NP-hard problem [54]. The most popular strategy to overcome this difficulty is to replace the $\ell_0$ norm with the $\ell_1$ norm. This technique is known as *convex relaxation* since the $\ell_0$ norm is relaxed with its closest convex function. In this view, Problem (1.1) becomes

$$\min_x \|x\|_1 \text{ subject to } Ax = b. \tag{1.2}$$

Problem (1.2) possesses several desired properties. Firstly, it is a convex problem, thus any local minimizer is a global minimum. Secondly, formulation (1.2) can be posed as a linear programming problem which can be solved in polynomial time with interior point methods. Lastly, Problem (1.2) has proven to be successful at approximately solving (1.1) [14, 25, 19, 24]. In fact, under some mild conditions, the optimal set of (1.2) is equal to the optimal set of (1.1) (see Baraniuk, Candès, Candès et al., Donoho et al. and Donoho [5, 14, 12, 25, 23]).

In the sparse representation setting, Problem (1.2) is viewed as a non-smooth regularization of the linear system $Ax = b$. When $A$ has more columns than rows (assuming full-row rank), the linear system has infinitely many solutions and further conditions are needed in order to obtain a meaningful solution. When $A$ has more rows than columns, existence of the solution fails (unless $b$ is in the column space of $A$ which is not the usual case), and additional conditions are needed to stabilize the least squares solution. When $A$ is square but not invertible, an acceptable solution is obtained by relaxing the equality constraint in

(1.2) and imposing additional conditions at the solution. Finally, when $A$ is non-singular but ill-conditioned, stability fails and additional conditions are needed to stabilize the solution. In all these cases, the additional condition is imposed by minimizing a functional that depends on the solution $x$. In particular, if $x$ is known to be sparse, a suitable functional is $\|x\|_1$ as indicated in (1.2).

In practical applications, however, the observed signal $b$ may be contaminated with additive noise of finite energy. Consequently, Problem (1.2) becomes

$$\min_x \; \|x\|_1 \;\text{ subject to }\; \|Ax - b\|_2 \leq \epsilon, \tag{1.3}$$

where $\epsilon$ is a measure of corruption of the observed signal assuming normality on the noise distribution.

Although several approaches such as SPGL1 [72] and NESTA [7] have been proposed to solve (1.3), a well known and equivalent formulation is

$$\min_x \; f(x) = \frac{1}{2}\|Ax - b\|_2^2 + \lambda\|x\|_1, \tag{1.4}$$

where $\lambda \in (0, \|A^T b\|_\infty)$ is a penalization parameter that balances the solution between fidelity with the linear model and sparsity. The equivalence between (1.3) and (1.4) is conceived in the sense that for any $\epsilon$, there exists $\lambda$ such that the solution of (1.3) also solves (1.4) and vice versa. A formal discussion of this equivalence can be found in Rockafellar [65] or alternatively in Ramirez [60].

A considerable number of approaches for solving (1.4) have proliferated in the recent years including interior-point methods [13, 45], gradient projection [33], iterative shrinkage/soft thresholding [6, 21, 34, 9], forward backward splitting [41, 20] and the recently developed split Bregman [37]. These methodologies, with exception of interior-point and gradient projection methods, are related to iterative-thresholding ideas and have proven to be faster than standard second order methods. Nonetheless, as the target signal becomes less sparse and the dynamic range increases, standard first order methods do not provide sufficient accuracy [7]. In this work, we concentrate on developing numerical algorithms for solving problems that involve both high dynamic range and nearly sparse signals.

Before concluding this section, we want to point out that this work concentrates in extending the ideas presented by Argáez at al. in [2], in which the absolute value function is approximated by a sequence of continuously differentiable and strictly convex functions in the spirit of an homotopic principle. We believe that we have succeeded not only from a numerical and experimental viewpoint, but also in theoretical aspects.

## 1.2 Outline

This work is divided into two parts. The first part covers Chapters 2 through 4, and focuses in developing a new strategy for solving Problem (1.4) where the target signal has a relative high dynamic range. In Chapter 2 we formally state the main problem and provide a detailed description of it. In Chapter 3 we present the proposed methodology and discuss several computational and theoretical aspects of the method. In Chapter 4 we present a set of numerical experiments that support the proposed methodology compared with some state-of-the-art algorithms.

The second part of this work is presented in Chapter 5 and covers a wide range of real applications including standard image processing problems, data and signal processing in micrometeorology, total variation and texture segmentation in seismic data. Finally, in Chapter 6 we present some concluding remarks.



## 1.3 Notation and Definitions

Throughout this work, we model discrete signals as vectors. A signal $x$ of length $n$ is represented as a vector in $\mathbb{R}^n$, and its $i$-th component is denoted as $x_i$, for $i = 1, \ldots, n$. Unless stated differently, we always assume $x$ to be of length $n$. A discrete 2D image of size $\sqrt{n} \times \sqrt{n}$ is represented as a vector of size $n$ produced by stacking the columns of the image consecutively. With no loss of generality, we only consider square images whose sides length

are a power of two. The support of a vector $x$ is defined as

$$\operatorname{supp}(x) = \{i \mid x_i \neq 0\}.$$

We say that a signal $x$ is ***sparse*** if it has only a few nonzero components. That is, $|\operatorname{supp}(x)| \ll n$, where $|\cdot|$ is the cardinal operator [1]. The energy of a vector $x$ is measured with the $\ell_2$ norm, that is, $\|x\|_2$. A signal $x$ is said to be ***nearly sparse*** if most of its energy is concentrated in a small subset of indexes $I \subseteq \{1, \cdots, n\}$. A way to visualize this is by rearranging the entries of $x$ in absolute value, from the largest to the smallest, and checking whether they decay rapidly.

### 1.3.1 Measure of sparsity

Several measures of sparsity have been proposed in the mathematical analysis community. A summary of the most important of such measures is presented below.

**The $\ell_p$ norm**

The $\ell_p$ norm is defined as

$$\|x\|_p^p = \sum_{i=1}^{n} |x_i|^p, \quad p \geq 1.$$

For $p < 1$, the $\ell_p$ norm is no longer a formal norm since it does not satisfy the triangular inequality. In the sparse representation literature, however, it is common to use the term "norm" to refer this family of quasi-norms. We shall also use this terminology keeping in mind this reservation. When $p \in (0, 1]$, the $\ell_p$ norm tends to produce lower values for sparse vectors, and higher values for dense vectors. For this reason, the $\ell_p$ norm for $p \in (0, 1]$ is widely used for promoting and measuring sparsity. As an illustration, Figure 1.1 shows the behavior of the scalar function $|x|^p$ for $p = 0.1, 0.5, 1.0$, and $2.0$. Consider the case $p = 0.1$ which is illustrated in red. Here the shape of $|x|^p$ penalizes values of $x$ that are different from

[1] The symbol $|\cdot|$ also refers to the absolute value operator. The use of this notation will be clear from the context.

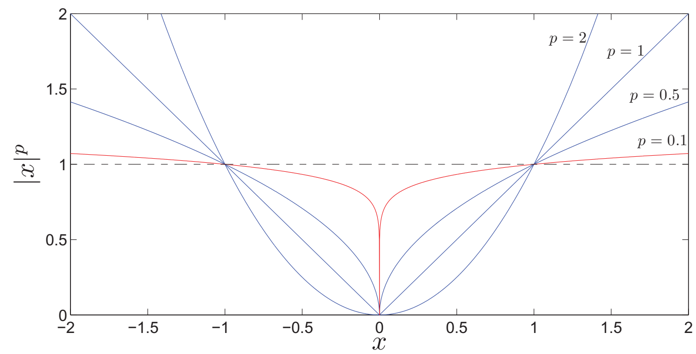


Figure 1.1: Behavior of the scalar function $|x|^p$ for different values of $p$ approaching to zero.

zero even if they are close to zero, and "forgives" values of $x$ only when they are exactly zero. This is not the case, for example, when $p = 2.0$ where values of $x$ close to zero are not severely penalized.

### The $\ell_0$ norm

Figure 1.1 gives us an intuition about the behavior of $\|x\|_p^p$ for $p \in (0, 1]$. It is natural to ask for the limit of the $\ell_p$ norm as $p$ goes to zero. This question gives rise to the definition of the $\ell_0$ norm:

$$\|x\|_0 = \lim_{p \to 0} \|x\|_p^p = \lim_{p \to 0} \sum_{i=1}^{n} |x_i|^p = |\mathrm{supp}(x)|.$$

Because $|\mathrm{supp}(x)|$ is equivalent to counting the number of nonzero components in the vector $x$, the $\ell_0$ norm is the most suitable measure of sparsity. Unfortunately, the function $\|x\|_0$ is not a convex function so that any local minimum is not necessarily a global minimum. Moreover, the $\ell_0$ norm is not continuous preventing standard numerical optimization techniques to be applied directly.

### 1.3.2 Sparsity in images

Previous work in the area of Computational Harmonic Analysis have shown that nearly-sparsity in the discrete cosine and wavelet domain can be used as a plausible image model [47, 48, 53]. In fact, the JPEG and JPEG-2000 encoding standards are based on the capability of the Discrete Cosine Transform (DCT) and Wavelet Transform (WT) to sparsely represent natural images, respectively [26]. It turns out that this is not an isolated case since natural images have redundant information due to the high local correlation derived at each pixel.

As an example, Figure 1.2 shows a typical natural image and a randomly generated image with the corresponding wavelet coefficients. The rearranged wavelet coefficients in absolute value are also presented. It can be seen that for the natural image, only a few wavelet coefficients capture almost all the energy of the image, whereas the energy of the randomly generated image is spread out along almost all the coefficients. This particular feature makes of special interest the problem of finding sparse signals among many others that are not.

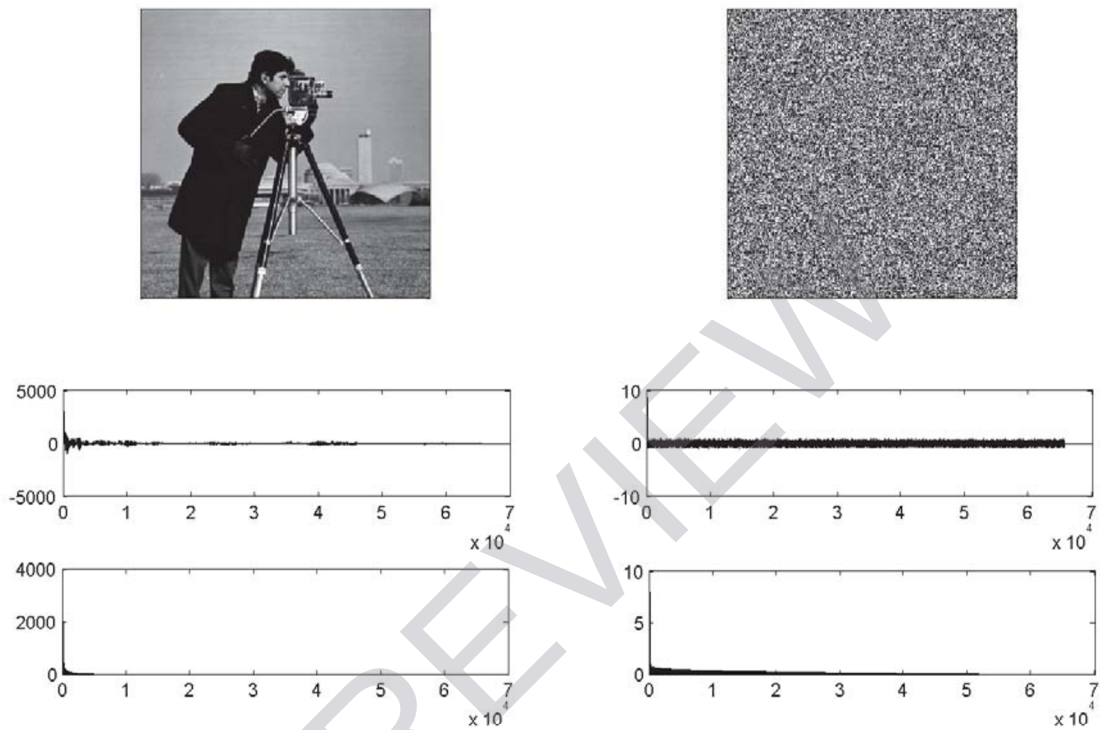


Figure 1.2: (First row) Left: a natural $256 \times 256$ image. Right: a randomly generated $256 \times 256$ image. (Second row) Left: wavelet coefficients of the natural image. Right: wavelet coefficients of the randomly generated image. (Third row) Left: rearranged wavelet coefficients of the natural image in absolute value. Right: rearranged wavelet coefficients of the randomly generated image in absolute value.

# Chapter 2

# Problem Formulation

In this chapter, we formally formulate our central problem and discuss its more relevant properties.

### Central Problem

We are interested in solving the problem

$$\min_x f(x) = \frac{1}{2}\|Ax-b\|_2^2+\lambda\|x\|_1, \tag{$P_1^\lambda$}$$

where $A \in \mathbb{R}^{m\times n}$, $b \in \mathbb{R}^m$, $\lambda \in (0, \|A^T b\|_\infty)$ and $m \leq n$. For the case $m < n$, we assume the matrix $A$ is full-rank, and for the case $m = n$, we assume singularity or severe ill-conditioning on $A$. This problem is previously stated as problem (1.4), and is rewritten here as $(P_1^\lambda)$ for convenience.

The parameter $\lambda$ is regarded as a penalization parameter that balances the optimal solution between sparsity and fidelity with respect to the linear model. If $\lambda \to 0^+$, the solution of $(P_1^\lambda)$ is the vector $x$ with minimum $\ell_1-$ norm that satisfies $A^T(Ax-b) = 0$ [45]. On the other hand, if $\lambda$ is too large, in particular if $\lambda \geq \|A^T b\|_\infty$, the least squares term in $(P_1^\lambda)$ is insignificant and therefore the solution $x$ must be zero.

### Properties

There are three fundamental properties of Problem $(P_1^\lambda)$ which we list below.

**Property 1** *The function $f$ is a convex function.*

**Proof** It follows from the fact that $f$ is a positive linear combination of two convex functions $g_1(x) = \|Ax - b\|_2^2$ and $g_2(x) = \|x\|_1$.

In turn, $g_1$ is convex since its Hessian $\nabla^2 g_1(x) = A^T A$ is positive semidefinite. Additionally, $g_2$ is convex since it is a norm.

**Property 2** *The function $f$ is proper.*

**Proof** For $f$ to be proper means that $f(x) < \infty$ for at least one $x \in \mathbb{R}^n$ and that $f(x) > -\infty$ for all $x \in \mathbb{R}^n$. The first condition is satisfied with e.g. $x = 0$, since $f(0) = \frac{1}{2}\|b\|_2^2 < \infty$. The second condition is trivially satisfied since $f(x) \geq 0$ for all $x$.

**Property 3** *The function $f$ is coercive.*

**Proof** For $f$ to be coercive means that $\lim_{\|x\| \to \infty} f(x) = +\infty$, which is easily verified from the definition of $f$.

The most important implication of these properties is that Problem $(P_1^\lambda)$ always has a solution, nonetheless, it is not necessarily unique [9, 45].

The major challenge in solving problem $(P_1^\lambda)$ is to deal with the non differentiability of the $\ell_1$ norm term. Several approaches have been proposed to circumvent this difficulty, including interior-point methods [13, 45], gradient projection [33], iterative shrinkage/soft thresholding [6, 21, 34, 9], forward backward splitting [41, 20] and the recently developed split Bregman [37]. We overcome the non-differentiability of the $\ell_1$ norm by approximating the absolute value with a continuously differentiable and strictly convex function that depends on a regularization parameter $\mu$ in the spirit of an homotopic principle [2]. The next chapter presents in detail this strategy and surveys some state-of-the-art algorithms.

# Chapter 3

# Proposed Methodology

In this chapter, we present a survey of some methodologies that address problem ($P_1^\lambda$) and focus on the Fixed Point Least Squares Preconditioned Conjugate Gradient (FPLS_PCG) algorithm introduced in this dissertation. We discuss several numerical properties of our FPLS_PCG algorithm, and provide a global convergence analysis giving theoretical guarantees of our method.

## 3.1 State-of-the-art Strategies

After the introduction of Compressed Sensing [12, 23] in the early 2000's, the applied mathematical and signal processing communities have paid a lot of attention to Problem ($P_1^\lambda$). As a result, several numerical algorithms for addressing this problem have been developed. In this section, we survey some of these algorithms referred as top solvers in recent literature.

### 3.1.1 l1_ls Algorithm

The l1_ls ($\ell_1$ Least Squares) algorithm, introduced in 2007 by Kim et al. [45], rewrites formulation ($P_1^\lambda$) as

$$\begin{aligned} \min_x \quad & \frac{1}{2}\|Ax-b\|_2^2 + \lambda \sum_{i=1}^{n} u_i \\ \text{subject to} \quad & -u_i \le x_i \le u_i, \end{aligned}$$

and solves it using logarithm barrier interior-point methods with inexact Newton's directions. The resulting Newton's system has the form

$$\begin{bmatrix} tA^TA + D_1 & D_2 \\ D_2 & D_1 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = - \begin{bmatrix} tA^T(Ax-b)+h_1 \\ t\lambda\mathbf{1} - h_2 \end{bmatrix},$$

where $D_1,\, D_2 \in \mathbb{R}^{n\times n}$ are diagonal matrices that depend on $x$ and $u$, and $h_1,\, h_2 \in \mathbb{R}^n$ are vectors also depending on $x$ and $u$. This linear system is always positive definite and is solved using a preconditioned conjugate gradient PCG method.

The l1_ls algorithm is implemented in MATLAB and can be freely downloaded from `http://www.stanford.edu/ boyd/l1_ls/`.

### 3.1.2 NESTA Algorithm

The NESTA (Nesterov's Algorithm) algorithm, introduced in 2010 by Becker et. al [7], considers the equivalent formulation

$$\min_x \|x\|_1 \;\; \text{subject to} \;\; \|Ax-b\|_2 \le \epsilon, \tag{$P_1^\epsilon$}$$

and applies the Nesterov's method for non-smooth convex optimization [55], where the $\ell_1$ norm is approximated by

$$\|x\|_1 \approx \max_{u\in Q_d} u^Tx - \mu\frac{1}{2}\|u\|_2^2$$

as $\mu \to 0$, and $Q_d = \{u : \; \|u\|_\infty \le 1\}$.

The NESTA algorithm is implemented in MATLAB and can be freely downloaded from `http://www-stat.stanford.edu/~candes/nesta/`.

### 3.1.3 SpaRSA Algorithm

The SpaRSA (Sparse Reconstruction by Separable Approximation) algorithm, introduced in 2008 by Wright et al. [79], generates a sequence of approximate solutions $\{x_k\}_{k=1,2,\cdots}$ where

$$x_{k+1} \in \arg\min_z \frac{1}{2}\|z-u_k\|_2^2 + \frac{\tau}{\alpha_k}\|z\|_1$$

and

$$u_k = x_k - \frac{1}{\alpha_k} A^T (A x_k - b).$$

This scheme follows the framework of iterative shrinkage/soft thresholding (IST) methods in which a gradient descent direction is computed followed by a shrinkage operation. The difference between SpaRSA and standard IST methods is that here the step size $\alpha_k$ is chosen according to Barzilai-Borwein formula, accelerating convergence.

The SpaRSA algorithm is implemented in MATLAB and can be freely downloaded from `http://www.lx.it.pt/~mtf/SpaRSA/`.

### 3.1.4 FISTA Algorithm

The FISTA (Fast Iterative Shrinkage-Thresholding Algorithm) algorithm, introduced in 2009 by Beck et al., is also an accelerated version of the standard IST scheme. That is, it performs a gradient descent direction followed by a shrinkage/soft thresholding operation. The fundamental difference lies in that FISTA incorporates information from the previous two iterations by constructing a linear combination of the form

$$y_{k+1} = x_k + \left( \frac{t_k - 1}{t_{k+1}} \right) (x_k - x_{k-1})$$

where

$$t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2} \quad t_1 = 1,$$

and then proceeding to the gradient descent and shrinkage/soft thresholding steps on $y_{k+1}$. This subtle modification, first introduced by Nesterov [56] in smooth convex optimization, results in a remarkable acceleration of the IST iterations.

The FISTA algorithm is implemented in MATLAB and can be freely downloaded from `http://www.eecs.berkeley.edu/~yang/software/l1benchmark/`.

### 3.1.5 PFSR Algorithm

The PFSR (Path Following Signal Recovery) algorithm, introduced in 2010 by Argáez et al. [2], considers the equivalent formulation

$$\min_x \|x\|_1 \text{ subject to } Ax+\nu = b, \tag{$P_1^\nu$}$$

where $\nu$ is a noise vector with $\|\nu\|_2 \leq \epsilon$, and solves a sequence of relaxed subproblems of the form

$$\min_x \sum_{i=1}^{n} \sqrt{x_i^2 + \mu} \text{ subject to } Ax + \nu = b, \tag{3.1}$$

as $\mu \to 0$.

Problem (3.1) is strictly convex and is guaranteed to have a unique global solution for a fixed $\mu$. Additionally, unlike the original problem ($P_1^\nu$), this new objective function is continuously differentiable, so classical constrained optimization techniques can be applied. Moreover, if $\mu$ is regarded as a continuous parameter, the set containing the solution of all subproblems (3.1) defines a smooth curve called the central path that converges to an optimal solution of problem ($P_1^\nu$) [2] (see Figure 3.1).

There are three fundamental aspects in the implementation of the PFSR algorithm: first, the characterization of the optimality set for each subproblem of the form (3.1), second, the equivalence between (3.1) as $\mu \to 0$ and ($P_1^\nu$), and third, the way in which the regularization parameter $\mu$ is updated.



#### Characterization of the optimality set

The Lagrangian function associated with (3.1) is

$$l(x, y) = \sum_{i=1}^{n} \sqrt{x_i^2 + \mu} + (Ax + \nu - b)^T y,$$

where $y \in \mathbb{R}^m$ is the Lagrange multiplier for the equality constraint. Therefore, the KKT conditions are given by the following square system of nonlinear equations:

$$F_\mu(x, y) = \begin{bmatrix} D_\mu(x)^{-1/2} x + A^T y \\ Ax + \nu - b \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tag{3.2}$$