Convergence and Stability of a Class of Iteratively Re-weighted Least Squares Algorithms for Sparse Signal Recovery in the Presence of Noise

Behtash Babadi[1,2], **Demba Ba**[1,2], **Patrick L. Purdon**[1,2], and **Emery N. Brown**[1,2,3]

Behtash Babadi: behtash@nmr.mgh.harvard.edu; Demba Ba: demba@mit.edu; Patrick L. Purdon: patrickp@nmr.mgh.harvard.edu; Emery N. Brown: enb@neurostat.mit.edu

[1]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139

[2]Department of Anesthesia, Critical Care, and Pain Medicine, Massachusetts General Hospital, Boston, MA 02114

[3]Harvard-MIT Division of Health, Sciences and Technology

Abstract

In this paper, we study the theoretical properties of a class of iteratively re-weighted least squares (IRLS) algorithms for sparse signal recovery in the presence of noise. We demonstrate a one-to-one correspondence between this class of algorithms and a class of Expectation-Maximization (EM) algorithms for constrained maximum likelihood estimation under a Gaussian scale mixture (GSM) distribution. The IRLS algorithms we consider are parametrized by $0 < \nu \leq 1$ and $\varepsilon > 0$. The EM formalism, as well as the connection to GSMs, allow us to establish that the IRLS(ν, ε) algorithms minimize ε-smooth versions of the ℓ_ν 'norms'. We leverage EM theory to show that, for each $0 < \nu \leq 1$, the limit points of the sequence of IRLS(ν, ε) iterates are stationary point of the ε-smooth ℓ_ν 'norm' minimization problem on the constraint set. Finally, we employ techniques from Compressive sampling (CS) theory to show that the class of IRLS(ν, ε) algorithms is stable for each $0 < \nu \leq 1$, if the limit point of the iterates coincides the global minimizer. For the case $\nu = 1$, we show that the algorithm converges exponentially fast to a neighborhood of the stationary point, and outline its generalization to super-exponential convergence for $\nu < 1$. We demonstrate our claims via simulation experiments. The simplicity of IRLS, along with the theoretical guarantees provided in this contribution, make a compelling case for its adoption as a standard tool for sparse signal recovery.

I. Introduction

Compressive sampling (CS) has been among the most active areas of research in signal processing in recent years [1], [2]. CS provides a framework for efficient sampling and re-construction of sparse signals, and has found applications in communication systems, medical imaging, geophysical data analysis, and computational biology.

The main approaches to CS can be categorized as optimization-based methods, greedy/ pursuit methods, coding-theoretic methods, and Bayesian methods (see [2] for detailed discussions and references). In particular, convex optimization-based methods such as ℓ_1-

minimization, the Dantzig selector, and the LASSO have proven successful for CS, with theoretical performance guarantees both in the absence and in the presence of observation noise. Although these programs can be solved using standard optimization tools, iteratively re-weighted least squares (IRLS) has been suggested as an attractive alternative in the literature. Indeed, a number of authors have demonstrated that IRLS is an efficient solution technique rivalling standard state-of-the-art algorithms based on convex optimization principles [3], [4], [5], [6], [7], [8]. Gorodnitsky and Rao [3] proposed an IRLS-type algorithm (FOCUSS) years prior to the advent of CS and demonstrated its utility in neuroimaging applications. Donoho et al. [4] have suggested the usage of IRLS for solving the basis pursuit de-noising (BPDN) problem in the Lagrangian form. Saab et al. [5] and Chartrand et al. [6] have employed IRLS for non-convex programs for CS. Carrillo and Barner [7] have applied IRLS to the minimization of a smoothed version of the ℓ_0 'norm' for CS. Wang et al. [8] have used IRLS for solving the ℓ_ν-minimization problem for sparse recovery, with $0 < \nu \leq 1$. Most of the above-mentioned papers lack a rigorous analysis of the convergence and stability of the IRLS in the presence of noise, and merely employ IRLS as a solution technique for other convex and non-convex optimization techniques. However, IRLS has also been studied in detail as a stand-alone optimization-based approach to sparse reconstruction in the absence of noise by Daubechies et al. [9].

In [10], Candès, Wakin and Boyd have called CS the "modern least-squares": the ease of implementation of IRLS algorithms, along with their inherent connection with ordinary least-squares provide a compelling argument in favor of its adoption as a standard algorithm for recovery of sparse signals [9].

In this work, we extend the utility of IRLS for compressive sampling in the presence of observation noise. For this purpose, we use the Expectation-Maximization (EM) theory for Normal/Independent (N/I) random variables and show that IRLS applied to noisy compressive sampling is an instance of the EM algorithm for constrained maximum likelihood estimation under a N/I assumption on the distribution of its components. This important connection has a two-fold advantage. First, the EM formalism allows to study the convergence of IRLS in the context of the EM theory. Second, one can evaluate the stability of the IRLS as a maximum likelihood problem in the context of noisy CS. More specifically, we show that the said class of IRLS algorithms, parametrized by $0 < \nu \leq 1$ and $\varepsilon > 0$, are iterative procedures to maximize ε-smooth approximations to the ℓ_ν 'norms'. We use EM theory to prove convergence of the algorithms to stationary points of the objective, for each $0 < \nu \leq 1$. We employ techniques from CS theory to show that the IRLS(ν, ε) algorithms are stable for each $0 < \nu \leq 1$, if the limit point of the iterates coincides with the global minimizer (which is trivially the case for $\nu = 1$, under mild conditions standard for CS). For the case $\nu = 1$, we show that the algorithm converges exponentially fast to a neighborhood of the stationary point, for small enough observation noise. We further outline the generalization of this result to super-exponential convergence for the case of $\nu < 1$. Finally, through numerical simulations we demonstrate the validity of our claims.

The rest of our treatment begins with Section II, where we introduce a fairly large class of EM algorithms for likelihood maximization within the context of N/I random variables. In the following section, we show a one-to-one correspondence between the said class of EM

algorithms and IRLS algorithms which have been proposed in the CS literature for sparse recovery. In Sections IV and V, we prove the convergence and stability of the IRLS algorithms identified previously in Section III. We derive rates of convergence in Section VI and demonstrate our theoretical predictions through numerical experiments in Section VII. Finally, we give concluding remarks in Section VIII.

II. Normal/Independent random variables and the Expectation-Maximization algorithm

A. N/I random variables

Consider a positive random variable U with probability distribution function $p_U(u)$, and an M-variate normal random vector Z with mean zero and non-singular covariance matrix Σ. For any constant M-dimensional vector μ, the random vector

$$Y = \mu + U^{-1/2} Z \quad (1)$$

is said to be a normal/independent (N/I) random vector [11]. N/I random vectors encompass large classes of multi-variate distributions such as the Generalized Laplacian and multi-variate t distributions. Many important properties of N/I random vectors can be found in [12] and [11]. In particular, the density of the random vector Y is given by

$$p_Y(y) = \frac{1}{(2\pi)^{M/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2} \kappa \left((y-\mu)^T {\textstyle\sum}^{-1} (y-\mu) \right) \right), \quad (2)$$

with

$$\kappa(x) := -2\ln \left(\int_0^\infty u^{M/2} e^{-ux/2} p_U(u) du \right) \quad (3)$$

for $x \geq 0$ [11].

N/I random vectors are also commonly referred to as Gaussian scale mixtures (GSMs). In the remainder of our treatment, we use the two terminologies interchangeably.

Eq. (2) is a representation of the density of an elliptically-symmetric random vector Y [13]. Eq. (3) gives a canonical form of the function $\kappa(\cdot)$ that arises from a given N/I distribution. However, when substituted in Eq. (2), not all $\kappa(\cdot)$ lead to a distribution in the GSM family, i.e., to random vectors which exhibit a decomposition as in Eq. (1). This will be important in our treatment because we will show that IRLS algorithms which have been proposed for sparse signal recovery correspond to specific choices of $\kappa(x)$ which *do* lead to GSMs. In [14], Andrews et al. give sufficient and necessary conditions under which a symmetric density belongs to the family of GSMs. In [11], Lange et al. generalize these results by giving sufficient and necessary conditions under which a spherically-symmetric random vector is a GSM (note that any elliptically-symmetric density as in Eq. (2), with non-singular covariance matrix, can be linearly transformed into a spherically-symmetric

density). The following theorem gives necessary and sufficient conditions under which a given choice of $\kappa(x)$ leads to a density in the N/I family.

Proposition 1 (Conditions for a GSM)—A function $f(x) : \mathbb{R} \mapsto \mathbb{R}$ is called completely mono-tone iff it is infinitely differentiable and $(-1)^k f^{(k)}(x) \geq 0$ for all non-negative integers k and all $x \geq 0$. Suppose Y is an elliptically-symmetric random vector, with representation as in Eq. (2), then Y is a N/I random vector iff $\kappa'(x)$ is completely monotone.

We refer the reader to [11] for a proof of this result.

B. EM algorithm

Now, suppose that one is given a total of P samples from multi-variate N/I random vectors, y_i, with mean and covariances $\mu_i(\theta)$ and $\Sigma_i(\theta)$ respectively, for $i = 1, 2, \cdots, P$, all parametrized by an unknown parameter vector θ. Let

$$\delta_i^2(\theta) := (y_i - \mu_i(\theta))^T \sum_i^{-1}(\theta)(y_i - \mu_i(\theta)) \quad (4)$$

for $i = 1, 2, \cdots, P$. Then, the log-likelihood of the P samples, parametrized by θ, is given by:

$$L\left(\{y_i\}_{i=1}^P; \theta\right) = -\frac{1}{2} \sum_{i=1}^P \left\{ \kappa\left(\delta_i^2(\theta)\right) + \ln |\sum_i(\theta)| \right\}. \quad (5)$$

An Expectation-Maximization algorithm for maximizing the log-likelihood results if one linearizes the function $\kappa(x)$ at the current estimate of θ. This is due to the fact that

$$\kappa'(x)\big|_{\delta_i^2(\theta)} = \mathbb{E}\{U|y_i, \theta\} \quad (6)$$

and hence linearization of $\kappa(x)$ at the current estimate of θ gives the Q-function by taking the scale variable U as the unobserved data [15], [11]. If $\theta^{(\ell)}$ is the current estimate of θ, then the $(\ell + 1)^{\text{th}}$ iteration of an EM algorithm maximizes the following Q-function:

$$Q\left(\theta|\theta^{(\ell)}\right) = -\frac{1}{2} \sum_{i=1}^P \left\{ \kappa'\left(\delta_i^2(\theta^{(\ell)})\right) \delta_i^2(\theta) + \ln |\sum_i(\theta)| \right\}. \quad (7)$$

Maximization of the above Q-function is usually more tractable than maximizing the original likelihood function.

The EM algorithm is an instance of the more general class of Majorization-Minimization (MM) algorithms [16]. The EM algorithm above can also be derived in the MM formalism, that is, without recourse to missing data or other statistical constructs such as marginal and complete data likelihoods. In [11], Lange et al. take the MM approach (without missing data) and point out that the key ingredient in the MM algorithm is the $\kappa(\cdot)$ function, which is related to the missing data formulation of the algorithm through Eq. (6).

III. Iterative Re-weighted Least Squares

In this section, we define a class of IRLS algorithms and show that they correspond to a specific class of EM algorithms under GSM assumptions.

A. Definition

Let $x \in \mathbb{R}^M$ be such that $|\{x_i : x_i \neq 0\}| \leq s$, for some $s < M$. Then, x is said to be an s-sparse vector. Consider the following observation model

$$b = Ax + n \quad (8)$$

where $b \in \mathbb{R}^N$, with $N < M$ is the observation vector, $A \in \mathbb{R}^{N \times M}$ is the measurement matrix, and $n \in \mathbb{R}^N$ is the observation noise. The noisy compressive sampling problem is concerned with the estimation of x given b, A and a model for n. Suppose that the observation noise n is bounded such that $\|n\|_2 \leq \eta$, for some fixed $\eta > 0$. Let $\mathcal{C} := \{x : \|b - Ax\|_2 \leq \eta\}$. Let $w \in \mathbb{R}^M$ such that $w_i > 0$ for all $i = 1, 2, \cdots, M$. Then, for all $x, y \in \mathbb{R}^M$, the inner-product defined by

$$\langle x, y \rangle_w := \sum_{i=1}^{M} w_i x_i y_i \quad (9)$$

induces a norm $\|x\|_w^2 := \langle x, x \rangle_w$.

Definition 2—Let $\nu \in (0, 1]$ be a fixed constant. Given an initial guess $x^{(0)}$ of x (e.g. the least-squares solution), the class of IRLS(ν, ε) algorithms for estimating x generates a sequence $\{x^{(\ell)}\}_{\ell=1}^{\infty}$ of iterates/refined estimates of x as follows:

$$x^{(\ell+1)} := \arg \min_{x \in \mathcal{C}} \|x\|_{w^{(\ell)}}^2, \quad (10)$$

with

$$w_i^{(\ell)} := \frac{1}{\left(\varepsilon^2 + \left(x_i^{(\ell)}\right)^2\right)^{1-\nu/2}} \quad (11)$$

for $i = 1, 2, \cdots, M$ and some fixed $\varepsilon > 0$.

Each iteration of the IRLS algorithm corresponds to a weighted least-squares problem constrained to the closed quadratic convex set $\mathcal{C}$, and can be efficiently solved using the standard convex optimization methods. The Lagrangian formulation of the IRLS has a simple closed form expression which makes it very appealing for implementation purposes [4]. Moreover, if the output SNR is greater than 1, that is, $\|n\|_2 \leq \eta < \|b\|_2$, then 0 is not a feasible solution. Hence, the gradient of $\|x\|_w^2$ is non-vanishing over $\mathcal{C}$. Therefore, the solutions lie on the boundary of $\mathcal{C}$, given by $\|b - Ax\|_2 = \eta$. Such a problem has been extensively studied in the optimization literature in its dual form, for which several robust

and efficient solutions exist (See [17] and references therein). Finally, note that when $\eta = 0$, the above algorithm is similar to the one studied by Daubechies et al. [9]. Throughout the paper, we may drop the dependence of IRLS(ν, ε) on ν and ε, and simply denote it by IRLS, wherever there is no ambiguity.

B. IRLS as an EM algorithm

Consider an M-dimensional random vector $y \in \mathbb{R}^M$ with independent elements distributed according to

$$p_{Y_i}(y_i) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\kappa\left((y_i - x_i)^2\right)\right) \quad (12)$$

for some function $\kappa(x)$ with completely monotone derivative. Note that y is parametrized by $\theta := (x_1, x_2, \cdots, x_M)^T \in \mathcal{C}$. The Q-function of the form (7) given the observation $y = 0 \in \mathbb{R}^M$ is given by:

$$Q\left(x | x^{(\ell)}\right) := -\frac{1}{2}\sum_{i=1}^{M}\kappa'\left(\left(x_i^{(\ell)}\right)^2\right)x_i^2 \quad (13)$$

Identifying $\kappa'\left(\left(x_i^{(\ell)}\right)^2\right)$ with $w_i^{(\ell)}$ in Eq. (10), we have

$$\kappa'(x) = \frac{1}{(\varepsilon^2 + x)^{1-\nu/2}} \quad (14)$$

for $x \geq 0$. It is not hard to show that $\kappa'(x)$ is completely monotone [11] and hence, according to Proposition 1, $\kappa(x)$ given by

$$\kappa(x) = \kappa(0) + \int_0^x \frac{1}{(\varepsilon^2 + t)^{1-\nu/2}}dt = \kappa(0) + \left(\varepsilon^2 + x\right)^{\nu/2} \quad (15)$$

defines an N/I univariate random variable with density given by Eq. (12). The log-likelihood corresponding to the zero observation is then given by

$$L(x) := -\frac{1}{2}\sum_{i=1}^{M}\left(\varepsilon^2 + x_i^2\right)^{\nu/2} \quad (16)$$

Therefore, the IRLS algorithm can be viewed as an iterative solution, which is an EM algorithm [11], for the following program:

$$\min_{x \in \mathscr{C}} \quad f_\nu(x) := \sum_{i=1}^{M}\left(\varepsilon^2 + x_i^2\right)^{\nu/2}. \quad (17)$$

Note that the above program corresponds to minimizing a ε-smoothed version of the ℓ_ν 'norm' of x (Figure 1) subject to the constraint $\|b - Ax\|_2 \quad \eta$, that is

$$\min_{x \in \mathscr{C}} \sum_{i=1}^{M} |x_i|^\nu \quad (18)$$

The function $f_\nu(x)$ has also been considered in [9] in the analysis of the IRLS algorithm for noiseless CS. However, the above parallel to EM theory can be generalized to any other weighting scheme with a completely monotone derivative. For instance, consider the IRLS algorithm with the weighting:

$$w_i^{(\ell)} = \frac{1}{\varepsilon + \left| x_i^{(\ell)} \right|} \quad (19)$$

for some $\varepsilon > 0$. Using the connection to EM theory [11], it can be shown that this IRLS is an iterative solution to

$$\min_{x \in \mathscr{C}} \quad \left\{ \|x\|_1 - \varepsilon \sum_{i=1}^{M} \ln\ (\varepsilon + |x_i|) \right\} \quad (20)$$

which is a perturbed version of ℓ_1 minimization subject to $\|b - Ax\|_2 \quad \eta$.

IV. Convergence

The convergence of the IRLS iterates in the absence of noise have been studied in [9], where the proofs rely on the null space property of the constraint set. The connection to EM theory allows to derive convergence results in the presence of noise using the rich convergence theory of EM algorithms.

A. Convergence of IRLS as an EM Algorithm

It is not hard to show that the EM algorithm provides a sequence of iterates $\{x^{(\ell)}\}_{\ell=0}^{\infty}$ so that sequence of log-likelihoods $\{L(x^{(\ell)})\}_{\ell=1}^{\infty}$ converges. However, one needs to be more prudent when making statements about the convergence of the iterates $\{x^{(\ell)}\}_{\ell=0}^{\infty}$. Let $\mathcal{C}$ denote a, non-empty, closed, strictly convex subset of $\mathbb{R}^N$. Let $\mathcal{M}: \mathbb{R}^M \mapsto \mathbb{R}^M$ be the map

$$\mathcal{M}(z) := \arg\min_{x \in \mathscr{C}} \|x\|_{w(z)}^2 \quad (21)$$

for all $z \in \mathbb{R}^M$, where

$$w_i(z) := \frac{1}{\left(\varepsilon^2 + z_i^2 \right)^{1-\nu/2}}, \quad \text{for } i = 1, 2, \cdots, M. \quad (22)$$

Results from convex analysis [18] imply the following sufficient and necessary optimality condition for $x^* \in \mathcal{C}$, the unique minimizer of $\|x\|^2_{w(z)}$ over $\mathcal{C}$:

$$\langle x^*, x - x^* \rangle_{w(z)} \geq 0, \text{ for all } x \in \mathcal{C}. \quad (23)$$

Moreover, continuity of $\|x\|^2_{w(z)}$ in x and z implies that $\mathcal{M}$ is a continuous map [19]. We prove this latter fact formally in Appendix A. The proof of convergence of the EM iterates to a stationary point of the likelihood function can be deduced from variations on the global convergence theorem of Zangwill [20] (See [19] and [21]). For completeness, we present a convergence theorem tailored for the problem at hand:

Theorem 3 (Convergence of the sequence of IRLS iterates)—Let $x^{(0)} \in \mathcal{C}$ and $\{x^{(\ell)}\}_{i=0}^{\infty} \in \mathcal{C}$ be a sequence defined as $x^{(\ell+1)} = \mathcal{M}(\mathrm{x}^{(\ell)})$ for all ℓ. Then, (i) $x^{(\ell)}$ is bounded and $\|x^{(\ell)} - x^{(\ell+1)}\|_2 \to 0$, (ii) every limit point of $\{x^{(\ell)}\}_{\ell=1}^{\infty}$ is a fixed point of $\mathcal{M}$, (iii) every limit point of $\{x^{(\ell)}\}_{\ell=0}^{\infty}$ is a stationary point of the function $f_\nu(x):=\sum_{i=1}^{M}\left(\varepsilon^2 + x_i^2\right)^{\nu/2}$ over $\mathcal{C}$, and (iv) $f_\nu(x^{(\ell)})$ converges monotonically to $f_\nu(x^*)$, for some stationary point x^*.

Proof: (i) is a simple extension of Lemmas 4.4 and 5.1 in [9], where one substitutes $\langle x^{(\ell+1)}, x^{(\ell)} - x^{(\ell+1)} \rangle_{w(x^{(\ell)})}$ 0 for the optimality conditions at each iterate ℓ.

(ii) From (*i*), $\{x^{(\ell)}\}_{\ell=0}^{\infty}$ is a bounded sequence. The Bolzano-Weierstrass theorem establishes that $\{x^{(\ell)}\}_{\ell=0}^{\infty}$ has at least one convergent subsequence. Let $x^{(\ell_k)} \to \bar{x}$ be one such convergent sub-sequence:

$$\lim_{k\to\infty} x^{(\ell_k+1)} = \lim_{k\to\infty} x^{(\ell_k)} = \bar{x}, \quad (24)$$

Since $x^{(\ell_k+1)} = \mathcal{M}(x^{(\ell_k)})$, the continuity of the map $\mathcal{M}$ implies that

$$\bar{x} = \lim_{k\to\infty} x^{(\ell_k+1)} = \mathcal{M}\left(\lim_{k\to\infty} x^{(\ell_k)}\right) = \mathcal{M}(\bar{x}). \quad (25)$$

Therefore, $\bar{x}$ is a fixed point of the mapping $\mathcal{M}$.

(iii) To establish (iii), we will show that the limit point of *any* convergent subsequence $\{x^{(\ell_k)}\}_k$ of $\{x^{(\ell)}\}_{\ell=0}^{\infty}$ satisfies the necessary conditions of the stationary points of the minimization of $f_\nu(x):=\sum_{i=1}^{M}\left(\varepsilon^2 + x_i^2\right)^{\nu/2}$ over $\mathcal{C}$. Note that $\bar{x} = \mathcal{M}(\bar{x})$ if and only if $\langle \bar{x}, x - \bar{x} \rangle_{w(\bar{x})}$ 0 for all $x \in \mathcal{C}$. Moreover,

$$\langle \overline{x}, x - \overline{x} \rangle_{w(\overline{x})} = \sum_{i=1}^{M} \frac{1}{\left(\varepsilon^2 + \overline{x}_i^2\right)^{1-\nu/2}} \overline{x}_i \left(x_i - \overline{x}_i\right) \quad = \frac{2}{\nu} \sum_{i=2}^{M} \underbrace{\frac{\nu \overline{x}_i}{2\left(\varepsilon^2 + \overline{x}_i^2\right)^{1-\nu/2}}}_{(\nabla f_\nu(\overline{x}))_i} \left(x_i - \overline{x}_i\right)$$

$$= \frac{2}{\nu} \langle \nabla f_\nu(\overline{x}), x - \overline{x} \rangle. \tag{26}$$

Note that $\langle \nabla f_\nu(x), x - \overline{x} \rangle \geq 0$ is the necessary condition for a stationary point $\overline{x}$ of $f_\nu(x)$ over the strictly convex set $\mathcal{C}$ [18]. Finally, (iv) follows from the continuity of $\|x\|_{w(z)}^2$ in x and z, and convexity of $\mathcal{C}$ (See Theorem 2 of [21]). This concludes the proof of the theorem.

B. Discussion

Note that Theorem 3 implies that if the minimizer of $f_\nu(x)$ over $\mathcal{C}$ is unique, then the IRLS iterates will converge to this unique minimizer. Moreover, by Theorem 5 of [21], the limit points of IRLS lie in a compact and connected subset of the set $\{x : f_\nu(x) = f_\nu(x^*)\}$. In particular, if the set of stationary points of $f_\nu(x)$ is finite, the IRLS sequence of iterates will converge to a unique stationary point. However, in general the IRLS is not guaranteed to converge (i.e. the set of limit points of the sequence of iterates is not necessarily a singleton).

There are various ways to choose ε adaptively or in a static fashion. Daubechies et al. [9] suggest a scheme where ε is possibly decreased in each step. This way $f_\nu(x)$ provides a better approximation to the ℓ_ν norm. Saab et al. [5] cascade a series of IRLS with fixed but decreasing ε, so that the output of each is used as the initialization of the next.

The result of Theorem 3 can be generalized to incorporate iteration-dependent changes of ε. Let $\{\varepsilon^{(\ell)} > 0\}_{\ell=1}^{\infty}$ be a sequence such that $\lim_{\ell \to \infty} \varepsilon^{(\ell)} = \overline{\varepsilon} \geq 0$. It is not hard to show that Theorem 3 holds for such a choice of $\{\varepsilon^{(\ell)}\}$, by defining $f_\nu(x) := \sum_{i=1}^{M} \left(\overline{\varepsilon}^2 + x_i^2\right)^{\nu/2}$, if $\overline{\varepsilon} > 0$. Let $\overline{x}$ be a limit point of the IRLS iterates. Then, if $\overline{\varepsilon} = 0$ and $T_{\overline{x}} := \text{supp}(\overline{x}) \subseteq \{1, 2, \cdots, M\}$, then the results of parts (iii) and (iv) of Theorem 3 holds for minimization of $f_\nu(x) := \sum_{i \in T_{\overline{x}}} |x_i|^\nu = \|x_{T_{\overline{x}}}\|_\nu^\nu$ over $\mathcal{C} \cap \{x : x_{T_{\overline{x}}^c} = 0\}$. This general result encompasses both the approach of Daubechies et al. [9] (in the absence of noise), and Saab et al. [5] as special cases. The technicality under such iteration-dependent choices of ε arises in showing that $\overline{x}$ is the fixed point of the limit mapping $\mathcal{M}(z)$, which can be established by invoking the uniform convergence of $\arg\min_{x \in \mathcal{C}} \|x\|_{w^{(\ell)}(z^{(\ell)})}^2$ to $\mathcal{M}(x)$ for any given subsequence $\{z^{(\ell)}\}$ converging to $\overline{x}$. A formal proof is given in Appendix B. For simplicity and clarity of presentation, the remaining results of this paper are presented with the assumption that $\varepsilon > 0$ is fixed.

V. Stability of IRLS for noisy CS

Recall that $f_\nu(x)$ is a smoothed version of $\|x\|_\nu^\nu$. Hence, for $0 < \nu \leq 1$, the global minimizer of $f_\nu(x)$ over $\mathcal{C}$ is expected to be close to the s-sparse x, given sufficient regularity conditions on the matrix A [8], [5]. Bounding the distance of this minimizer to the s-sparse x provides the

desired stability bounds. For $v = 1$, $f_1(x)$ is strictly convex. Therefore the solution of the minimization of $f_1(x)$ over the convex set $\mathcal{C}$ is unique [18]. Hence, the IRLS iterates will converge to the unique minimizer in this case. However, for $v < 1$, the IRLS iterates do not necessarily converge to a global minimizer of $f_v(x)$ over $\mathcal{C}$. In practice, the IRLS is applied with randomly chosen initial values, and the limit point with the highest log-likelihood is chosen [7].

Recall that the matrix $A \in \mathbb{R}^{N \times M}$ is said to have Restricted Isometry Property (RIP) [22] of order $s < M$ with constant $\delta_s \in (0, 1)$, if for all $x \in \mathbb{R}^M$ supported on any index set $T \subset \{1, 2, \cdots, M\}$ satisfying $|T| \quad s$, we have

$$(1-\delta_s)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1+\delta_s)\|x\|_2^2. \quad (27)$$

The following theorem establishes the stability of the minimization of $f_v(x)$ over $\mathcal{C}$ in the noisy setting:

Theorem 4

Let $b = Ax + n$ be given such that $x \in \mathbb{R}^M$ is s-sparse. Let m be a fixed integer and suppose that $A \in \mathbb{R}^{N \times M}$ satisfies

$$\delta_m + \left(\frac{s}{m}\right)^{1-2/v} \delta_{m+s} < \left(\frac{s}{m}\right)^{1-2/v} - 1. \quad (28)$$

Suppose that $\|n\|_2 \quad \eta$ and let $\mathcal{C} := \{x : \|b - Ax\|_2 \quad \eta\}$. Let $\varepsilon > 0$ be a fixed constant. Then, the solution to the following program

$$\overline{x} = \arg\min_{x \in \mathcal{C}} \sum_{i=1}^{M} \left(\varepsilon^2 + x_i^2\right)^{v/2} \quad (29)$$

satisfies

$$\|\overline{x} - x\|_2^v \leq C_1 \eta^v + C_2 s^{v/2} \varepsilon^v \quad (30)$$

where C_1 and C_2 are constants depending only on v, s/m, δ_m and δ_{m+s}.

Proof—The proof is a modification of the proof of Theorem 4 in [5], which is based on the proof of the main result of [23]. Let $T_0 := \{i : x_i \quad 0\}$. Let $S \subseteq \{1, 2, \cdots, M\}$. We define

$$f_v(x_S) := \sum_{i \in S} \left(\varepsilon^2 + x_i^2\right)^{v/2}. \quad (31)$$

Let $\overline{x}$ be a global minimizer of $f_v(x)$ over $\mathcal{C}$ and let $h := \overline{x} - x$. It is not hard to verify the following fact:

$$f_\nu(h_{T^c}) \le f_\nu(h_T). \quad (32)$$

The above inequality is the equivalent of the cone constraint in [23]. Moreover, it can be shown that

$$f_\nu(x_S) \le s^{1-\nu/2}\left(\|x_S\|_2^\nu + s^{\nu/2}\varepsilon^\nu\right) \quad (33)$$

for any $x \in \mathbb{R}^M$ and $S \subseteq \{1, 2, \cdots, M\}$ such that $|S| \quad s$. By dividing the set T_0^c into the sets T_1, T_2, $\cdots$ of size m, sorted according to decreasing magnitudes of the elements of $h_{T_0^c}$, it can be shown that

$$\|h_{T_{01}^c}\|_2^\nu \le \frac{f_\nu(h_{T_0})}{(2/\nu-1)^{\nu/2}m^{1-\nu/2}} \le \frac{1}{(2/\nu-1)^{\nu/2}}\left(\frac{s}{m}\right)^{1-\nu/2}\left(\|h_{T_{01}}\|_2^\nu + s^{\nu/2}\varepsilon^\nu\right), \quad (34)$$

where $T_{01} := T_0 \cup T_1$. Also, by the construction of $\mathcal{C}$ and the hypothesis of the theorem about A, one can show that [5], [23]:

$$(2\eta)^\nu \ge \|Ah\|_2^\nu \ge \left\{(1-\delta_{m+s})^{\nu/2} - (1+\delta_m)^{\nu/2}\left(\frac{s}{m}\right)^{1-\nu/2}\right\}\|h_{T_{01}}\|_2^\nu \quad (35)$$

Combining Eqs. (34) and (35) with the fact that $\|h\|_2^\nu \le \|h_{T_{01}}\|_2^\nu + \|h_{T_{01}^c}\|_2^\nu$, yields.

$$\|h\|_2^\nu \le C_1\eta^\nu + C_2 s^{\nu/2}\varepsilon^\nu \quad (36)$$

where

$$C_1 = 2^\nu \frac{1 + \frac{1}{(2/\nu-1)^{\nu/2}}\left(\frac{s}{m}\right)^{1-\nu/2}}{(1-\delta_{m+s})^{\nu/2} - (1+\delta_m)^{\nu/2}\left(\frac{s}{m}\right)^{1-\nu/2}} \quad (37)$$

and

$$C_2 = \frac{1}{(2/\nu-1)^{\nu/2}}\left(\frac{s}{m}\right)^{1-\nu/2}. \quad (38)$$

Remark

Note that the result of Theorem 4 can be extended to compressible signals in a straightforward fashion [5]. Moreover, as it will be shown in the next section, the hypothesis

of Eq. (28) can be relaxed to the sparse approximation property developed in [24], with a similar characterization of the global minimizer under study.

VI. Convergence Rate of IRLs for noisy CS

In presenting our results on the convergence rate of IRLS in presence of noise, it is more convenient to employ a slightly weaker notion of near isometry of the matrix A developed in [24]. This is due to the structure of the IRLS algorithm, which makes it more convenient to analyze the convergence rate in the ℓ_1 sense, and as it becomes clear shortly, the sparse approximation property is the more appropriate choice of regularity condition on the matrix A.

A. Sparse approximation property and its consequences

We say that a matrix A has the sparse approximation property of order s if

$$\|x_S\|_2 \leq D\|Ax\|_2 + \frac{\beta}{\sqrt{s}}\|x_{S^c}\|_1 \quad (39)$$

for all $x \in \mathbb{R}^M$, where S is an index set such that $|S| \leq s$, and D and β are positive constants. Note that RIP of order $2s$ implies sparse approximation property [24], but the converse is not necessarily true. The error bounds obtained in Theorem 4 can be expressed in terms of D and β in a straightforward fashion [24]. A useful consequence of the sparse approximation property is the reverse triangle inequality in presence of noise:

Proposition 5—Let A satisfy the sparse approximation property of order s with constants $\beta < 1$ and D. Let $x_1, x_2 \in \mathcal{C} := \{x : \|b - Ax\|_2 \leq \eta\}$ and suppose that x_1 is s-sparse. Then, we have:

$$\|x_2 - x_1\|_1 \leq \frac{4D}{1-\beta}\sqrt{s}\eta + \frac{1+\beta}{1-\beta}(\|x_2\|_1 - \|x_1\|_1).$$

Proof: Let T be the support of x_1. Then, we have:

$$\begin{aligned}
\|(x_2 - x_1)_{T^c}\|_1 &\leq \|x_{2T^c}\|_1 + \|x_{1T^c}\|_1 \\
&= \|x_2\|_1 - \|x_{2T}\|_1 \\
&\leq \|x_2\|_1 - \|x_1\|_1 + \|(x_2 - x_1)_T\|_1
\end{aligned} \quad (40)$$

The sparse approximation property implies that

$$\|(x_2 - x_1)_T\|_1 \leq \sqrt{s}D\|A(x_2 - x_1)\|_2 + \beta\|(x_2 - x_1)_{T^c}\|_1 \quad (41)$$

Moreover, $\|A(x_2 - x_1)\|_2 \leq 2\eta$, by the construction of $\mathcal{C}$. Hence, combining Eqs. (40) and (41) yield:

$$\left\| (x_2 - x_1)_T \right\|_1 \leq \frac{2D}{1-\beta} \sqrt{s}\eta + \frac{\beta}{1-\beta} (\|x_2\|_1 - \|x_1\|_1), \quad (42)$$

which together with Eq. (40) gives the statement of the proposition.

The above reverse triangle inequality allows to characterize the stability of the IRLS in the ℓ_1 sense. This is indeed the method used by Daubechies et. al [9] in the absence of noise. Let $\bar{x}$ be the minimizer of $f_1(x)$ over $\mathcal{C}$. Then, it is straightforward to show that

$$\|\bar{x}\|_1 \leq f_1(\bar{x}) \leq f_1(x) \leq \|x\|_1 + M\varepsilon \quad (43)$$

Combining the above inequality with the statement of Proposition 5 yields:

$$\|\bar{x} - x\|_1 \leq \frac{4D}{1-\beta} \sqrt{s}\eta + \frac{1+\beta}{1-\beta} M\varepsilon \quad (44)$$

Note that the above bound is optimal in terms of η up to a constant, since $\|\eta\|_1 \leq \sqrt{N}\|n\|_2 = \sqrt{N}\eta$, and in fact $\|n\|_1$ may achieve the value of $\sqrt{N}\eta$.

B. Convergence rate of the IRLS

Let $\{x^{(\ell)}\}_{\ell=0}^{\infty}$ be a sequence of IRLS iterates that converge to a stationary point $\bar{x}$, for $\nu = 1$. We have the following theorem regarding the convergence rate of IRLS:

Theorem 6—Suppose that the matrix A satisfies the sparse approximation property of order s with constants D and β. Suppose that for some $\rho < 1$ we have

$$\mu := \frac{\beta(1+\beta)}{1-\rho} < 1. \quad (45)$$

and let R_0 be the right hand side of Eq. (44), so that $\|\bar{x} - x\|_1 \quad R_0$. Assume that

$$\min_{i \in T} |x_i| > R_0. \quad (46)$$

Let

$$e^{(\ell)} := x^{(\ell)} - \bar{x}. \quad (47)$$

Then, there exists a finite ℓ_0 such that for all $\ell > \ell_0$ we have:

$$\|e^{(\ell+1)}\|_1 \leq \mu \|e^{(\ell)}\|_1 + R_1 \quad (48)$$

for some R_1 comparable to R_0, which is explicitly given in this paper.

Proof: The proof of the theorem is mainly based on the proof of Theorem 6.4 of [9]. The convergence of the IRLS iterates implies that $e^{(\ell)} \to 0$. Let T be the support of the s-sparse vector x. Therefore, there exists ℓ_0 such that

$$\|e^{(\ell_0)}\|_1 \leq \rho \min_{i \in T} |\overline{x}_i|. \quad (49)$$

Clearly the right hand side of the above inequality is positive, since

$$\min_{i \in T} |\overline{x}_i| \geq \min_{i \in T} |x_i| - R_0 > 0 \quad (50)$$

by hypothesis. Following the proof method of [9], we want to show (by induction) that for all $\ell > \ell_0$, we have

$$\|e^{(\ell+1)}\|_1 \leq \mu \|e^{(\ell)}\|_1 + R_1 \quad (51)$$

for some R_1 that we will specify later. Consider $e^{(\ell+1)} = x^{(\ell+1)} - x$. The first order necessary conditions on $x^{(\ell+1)}$ give:

$$\sum_{i=1}^{M} \frac{x_i^{(\ell+1)}}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} (\overline{x}_i - x_i^{(\ell+1)}) \geq 0. \quad (52)$$

Substituting $x^{(\ell+1)}$ by $x \mp e^{(\ell+1)}$ yields

$$\sum_{i=1}^{M} \frac{|e_i^{(\ell+1)}|^2}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \leq -\sum_{i=1}^{M} \frac{e_i^{(\ell+1)} \overline{x}_i}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \quad (53)$$

We intend to bound the term on the right hand side. First, note that

$$\left| \sum_{i \in T} \frac{e_i^{(\ell+1)} \overline{x}_i}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right| \leq \frac{1}{1-\rho} \|e_T^{(\ell+1)}\|_1, \quad (54)$$

since

$$\left| \frac{\overline{x}_i}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right| \leq \frac{|\overline{x}_i|}{|x_i^{(\ell)}|} \leq \frac{|\overline{x}_i|}{|\overline{x}_i + e_i^{(\ell+1)}|} \leq \frac{1}{1-\rho}$$

by hypothesis. Moreover, the sparse approximation property implies that

$$\|e_T^{(\ell+1)}\|_1 \leq \sqrt{s}D\|Ae^{(\ell+1)}\|_2 + \beta\|e_{T^c}^{(\ell+1)}\|_1 \quad \leq \sqrt{s}D\|A(x^{(\ell+1)} - \overline{x})\|_2 + \beta\|e_{T^c}^{(\ell+1)}\|_1$$
$$\leq 2D\eta\sqrt{s} + \beta\|e_{T^c}^{(\ell+1)}\|_1, \tag{55}$$

thanks to the tube constraint $\|b - Ax\|_2 \quad \eta$. Hence, the left hand side of Eq. (54) can be bounded as:

$$\left| \sum_{i \in T} \frac{e_i^{(\ell+1)} \overline{x}_i}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right| \leq \frac{1}{1-\rho} \left(2D\sqrt{s}\eta + \beta\|e_{T^c}^{(\ell+1)}\|_1 \right). \tag{56}$$

Also, we have:

$$\left| \sum_{i \in T^c} \frac{e_i^{(\ell+1)} \overline{x}_i}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right| \leq \max_{i \in T^c} \left\{ \frac{e_i^{(\ell+1)}}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right\} \|\overline{x}_{T^c}\|_1 \leq \frac{\|e^{(\ell+1)}\|_\infty}{\varepsilon} \|\overline{x}_{T^c}\|_1 := \gamma^{(\ell)} \|\overline{x}_{T^c}\|_1 \tag{57}$$

Note that $\gamma^{(\ell)} \to 0$, since $e^{(\ell)} \to 0$. Therefore, we have

$$\left| \sum_{i=1}^M \frac{|e_i^{(\ell+1)}|^2}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right| \leq \frac{1}{1-\rho} \left(2D\sqrt{s}\eta + \beta\|e_{T^c}^{(\ell+1)}\|_1 \right) + \gamma^{(\ell)} \|\overline{x}_{T^c}\|_1 \tag{58}$$

An application of the Cauchy-Schwarz inequality yields:

$$\|e_{T^c}^{(\ell+1)}\|_1^2 \leq \left(\sum_{i \in T^c} \frac{|e_i^{(\ell+1)}|^2}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right) \left(\sum_{i \in T^c} \sqrt{x_i^{(\ell)^2} + \varepsilon^2} \right)$$
$$\leq \left(\sum_{i \in T^c} \frac{|e_i^{(\ell+1)}|^2}{\sqrt{x_i^{(\ell)^2} + \varepsilon^2}} \right) \left(\sum_{i \in T^c} (|e_i^{(\ell)}| + |\overline{x}_i| + \varepsilon) \right) \tag{59}$$
$$\leq \left(\frac{1}{1-\rho}(2D\sqrt{s}\eta + \beta\|e_{T^c}^{(\ell+1)}\|_1) + \gamma^{(\ell)}\|\overline{x}_{T^c}\|_1 \right) \left(\|e^{(\ell)}\|_1 + \|\overline{x}_{T^c}\|_1 + M\varepsilon \right).$$

Hence,

$$\|e_{T^c}^{(\ell+1)}\|_1 \leq \frac{\beta}{1-\rho}\|e^{(\ell)}\|_1 + \left(\frac{\beta}{1-\rho} + \gamma^{(\ell)}\frac{1-\rho}{\beta} \right) \|\overline{x}_{T^c}\|_1 + \frac{\beta}{1-\rho}M\varepsilon + \frac{2D}{\beta}\sqrt{s}\eta. \tag{60}$$

First, note that $\gamma^{(\ell)}$ is a bounded sequence for $\ell > \ell_0$. Let γ_0 be an upper bound on $\gamma^{(\ell)}$ for all $\ell > \ell_0$. We also have:

$$\|\overline{x}_{T^c}\|_1 = \|(\overline{x} - x)_{T^c}\|_1 \leq \|\overline{x} - x\|_1 \leq R_0 \tag{61}$$

Now, we have:

$$\|e^{(\ell+1)}\|_1 \leq \|e_T^{(\ell+1)}\|_1 + \|e_{T^c}^{(\ell+1)}\|_1 \quad \leq (1+\beta)\|e_{T^c}^{(\ell+1)}\|_1 + 2D\sqrt{s}\eta$$
$$\leq \frac{\beta(1+\beta)}{1-\rho}\|e^{(\ell)}\|_1 + R_1 \qquad (62)$$

where

$$R_1 := (1+\beta)\left(\frac{\beta}{1-\rho} + \gamma_0 \frac{1-\rho}{\beta}\right) R_0 + \left(\frac{(4\beta+2)}{\beta}\right) D\sqrt{s}\eta + \left(\frac{\beta}{1-\rho}M\right)\varepsilon \quad (63)$$

This concludes the proof of the theorem.

C. Discussion

Eq. (62) implies that $\limsup_{\ell\to\infty}\|e^{(\ell)}\|_1 \leq \left(1-\frac{\beta(1+\beta)}{1-\rho}\right)^{-1} R_1$. Therefore, the IRLS iterates approach a neighborhood of radius $\left(1-\frac{\beta(1+\beta)}{1-\rho}\right)^{-1} R_1$ (in the ℓ_1 sense) around the stationary point $\bar{x}$ exponentially fast. Note that the radius of this neighborhood is comparable to the upper bound on the distance of $\bar{x}$ to the s-sparse vector x (in the ℓ_1 sense) given by R_0. Hence, it is expected that with relatively few iterations of IRLS, one gets a reasonable estimate of x (Indeed, numerical studies in Section VII-C confirm this observation). Although the bound of the theorem holds for all $\ell > \ell_0$, it is most useful when $(1-\mu)^{-1} R_1$ is less than $\rho \min_{i\in T}|x_i|$. A sufficient condition to guarantee this is that

$\min_{i\in T}|x_i| \geq R_0 + \frac{1}{\rho(1-\mu)} R_1$, which gives an upper bound on the noise level η and ε.

It is straightforward to extend the above theorem to the case of $\nu < 1$. As shown in [9], the local convergence of $\|e^{(\ell)}\|_\nu^\nu$ in the case of $\nu < 1$ and in the absence of noise is super-linear, with exponent $2 - \nu$. It is not hard to show that in the presence of noise, one can recover the super-linear local convergence with exponent $2 - \nu$. We refer the reader to Theorem 7.9 of [9], which can be extended to the noisy case with the respective modifications to the proof of Theorem 6.

VII. Numerical Experiments

In this section, we use numerical simulations to explore and validate the stability and convergence rate analyses of the previous sections. In particular, we compare ℓ_1-minimization to $f_\nu(\cdot)$-minimization, in both cases in the presence of noise, for different values of ν, ε, and signal-to-noise ratio (SNR).

A. Experimental set-up

For fixed ν, ε and η,

1. Select N and M so that A is an $N \times M$ matrix; sample A with independent Gaussian entries.

2. Select $1 \leq S < M/2$.

3. Select T_0 of size S uniformly at random and set $x_j = 1$ for all $j \in T_0$, and 0 otherwise.

4. Make $b = Ax + n$, where each entry of n is drawn uniformly in $(-a, a)$, for some a that depends on η; find the solution $\bar{x}$ to the program of Eq. (18) by IRLS.

5. Compare $\bar{x}$ to x.

6. Repeat 50 times.

For each ν, ε and η, we compare the program solved in Step 4 to solving the program of Eq. (19) for $\nu = 1 (\ell_1$-minimization). We solve each IRLS iteration, as well as the ℓ_1-minimization problem, using CVX, a package for specifying and solving convex programs [25], [26].

Remark—Modulo some constants, both η and ε appear in the same proportion in the stability bound derived in Theorem 3. Intuitively, this means that, the higher the SNR (small η), the smaller the value of ε one should pick to solve the program. In our experiment, we start with a fixed ε for the smallest SNR value (5 dB), and scale this value linearly for each subsequent value of the SNR. In particular, we use $\varepsilon(\text{SNR}) = \varepsilon(5) \frac{\eta(\text{SNR})}{\eta(5)}$, where we use the loose notation $\eta(\text{SNR})$ to reflect the fact that each choice of SNR corresponds to a choice of η, and vice versa. In summary, our experimental set-up remains the same, except for the fact that we choose values of ε which depend on η.

B. Analysis of Stability

Figures 2 and 3 demonstrate the stability of IRLS for $\varepsilon = 10^{-4}$, respectively for choices of $\nu = 1$ and $\nu = 1/2$. Figure 2 shows (as expected) that the stability of IRLS is comparable to that of ℓ_1-minimization for $\nu = 1$ and small ε. Moreover, only a few number of IRLS iterations are required to reach a satisfactory value of the MSE. These observations also apply to Figure 3, which further highlights the sparsifying properties of $f_\nu(\cdot)$-minimization for $\nu < 1$. Indeed, the MSE achieved for $\nu = 1/2$ are smaller than those achieved for $\nu = 1$. Figure 4 shows that the approximation to the ℓ_1-norm improves as one decreases the value of ε. In all three figures, we can clearly identify the log-linear dependence of the MSE as a function of η, which is predicted by the bound we derived in Theorem 4.

C. Convergence rate analysis

Figure 5 shows that the IRLS algorithm for $f_1(\cdot)$-minimization converges exponentially fast to a neighborhood of the fixed-point of the algorithm. Moreover, the larger the SNR, the faster the convergence. These two observations are as predicted by the bound of Theorem 6. Figure 6 shows an alternate depiction of these observations in the log scale. Figure 7 shows that the IRLS algorithm for $f_{1/2}(\cdot)$-minimization converges super-exponentially fast to a neighborhood of the fixed-point of the algorithm. As observed in Figure 5, the larger the SNR, the faster the convergence. Figure 8 shows an alternate depiction of these observations in the log scale.

VIII. Discussion

In this paper, we provided a rigorous theoretical analysis of various iteratively re-weighted least-squares algorithms which have been proposed in the literature for recovery of sparse signals in the presence of noise [3], [4], [5], [6], [7], [8], [9]. We framed the recovery problem as one of constrained likelihood maximization using EM under Gaussian scale mixture assumptions. On the one hand, we were able to leverage the power of the EM theory to prove convergence of the said IRLS algorithms, and on the other hand, we were able to employ tools from CS theory to prove the stability of these IRLS algorithms and to derive explicit rates of convergence. We supplemented our theoretical analysis with numerical experiments which confirmed our predictions.

The EM interpretation of the IRLS algorithms, along with the derivation of the objective functions maximized by these IRLS algorithms, are novel. The proof of convergence is novel and uses ideas from Zangwill [20] which, in a sense, are more general than the proof presented by Daubechies [9] in the noiseless case. We have not presented the proof in the most general setting. However, we believe that the key ideas in the proof could be useful in various other settings involving iterative procedures to solve optimization problems. The proof of stability of the algorithms is novel; it relies on various properties of the function $f_\lambda(\cdot)$, along with techniques developed by Candés et al. [23]. The analysis of the rates of convergence is novel and makes interesting use of the sparse approximation property [24], along with some of the techniques introduced in [9].

Although we have opted for a fairly theoretical treatment, we would like to emphasize that the beauty of IRLS lies in its simplicity, not in its theoretical properties. Indeed, the simplicity of IRLS alone makes it appealing, especially for those who do not possess formal training in numerical optimization: no doubt, it is easier to implement least-squares, constrained or otherwise, than it is to implement a solver based on barrier or interior-point methods. Our hope is that a firm theoretical understanding of the IRLS algorithms considered here will increase their adoption as a standard framework for sparse approximation.

References

1. Donoho DL. Compressed sensing. IEEE Transactions on Information Theory. Apr.2006 52:1289–1306.

2. Bruckstein A, Donoho D, Elad M. From sparse solutions of systems of equations to sparse modeling of signals and images. SIAM Review. 2009; 51(1):34–81.

3. Gorodnitsky I, Rao BD. Sparse signal reconstruction from limited data using focuss: a recursive weighted norm minimization algorithm. IEEE Transactions on Signal Processing. 1997; 45(3):600–616.

4. Donoho DL, Elad M, Temlyakov VN. Stable recovery of sparse overcomplete representations in the presence of noise. IEEE Transactions on Information Theory. Jan.2006 52:6–18.

5. Saab, R.; Chartrand, R.; Yilmaz, O. Stable sparse approximations via nonconvex optimization. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2008); 2008. p. 3885-3888.

6. Chartrand, R.; Yin, W. Iteratively reweighted algorithms for compressive sensing. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2008); 2008. p. 3869-3872.

7. Carrillo, RE.; Barner, K. Iteratively re-weighted least squares for sparse signal reconstruction from noisy measurements," in. 43rd Annual Conference on Information Sciences and Systems (CISS 2009); March 2009; p. 448-453.

8. Wang W, Xu W, Tang A. On the performance of sparse recovery via ℓ_p-minimization. IEEE Transactions on Information Theory. Nov.2011 57:7255–7278.

9. Daubechies I, DeVore R, Fornasier M, Gntrk CS. Iteratively reweighted least squares minimization for sparse recovery. Comm Pure Appl Math. 2010; 63(1):1–38.

10. Candès EJ, Wakin M, Boyd S. Enhancing sparsity by reweighted ℓ_1 minimization. J Fourier Anal Appl. Dec.2008 14:877–905.

11. Lange K, Sinsheimer JS. Normal/independent distributions and their applications in robust regression. Journal of Computational and Graphical Statistics. 1993; 2(2):175–198.

12. Dempster, AP.; Laird, NM.; Rubin, DB. Iteratively reweighted least squares for linear regression when errors are normal/independent distributed. In: Krishnaiah, PR., editor. Multivariate Analysis V. Elsevier Science Publishers; 1980. p. 35-57.

13. Huber, P.; Ronchetti, E. MyiLibrary. Robust statistics. Vol. 1. Wiley Online Library; 1981.

14. Andrews D, Mallows C. Scale mixtures of normal distributions. Journal of the Royal Statistical Society. Series B (Methodological). 1974:99–102.

15. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society Series B (Methodological). 1977; 39(1):1–38.

16. Lange, K. Optimization. Springer; 2004.

17. Golub GH, von Matt U. Quadratically constrained least squares and quadratic problems. Numerische Mathematik. 1992; 59(1):561–580.

18. Bertsekas, DP. Convex Optimization Theory. 1. Athena Scientific; 2009.

19. Wu CFJ. On the convergence properties of the em algorithm. Ann Statist. 1983; 11(1):95–103.

20. Zangwill, WI. Nonlinear Programming: a unified approach. Prentice-Hall; 1969.

21. Nettleton D. Convergence properties of the em algorithm in constrained parameter spaces. The Canadian Journal of Statistics/La Revue Canadienne de Statistique. 1999; 27(3):639–648.

22. Candès EJ, Tao T. Decoding by linear programming. IEEE Trans on Information Theory. Dec. 2005 51:4203–4215.

23. Candès EJ, Romberg J, Tao T. Stable signal recovery for incomplete and inaccurate measurements. Commun Pure Appl Math. Aug.2006 59:1207–1223.

24. Sun Q. Sparse approximation property and stable recovery of sparse signals from noisy measurements. IEEE Transactions on Signal Processing. 2011; 59(10):5086–5090.

25. Grant, M.; Boyd, S. CVX: Matlab software for disciplined convex programming, version 1.21. Apr. 2011 http://www.stanford.edu/~boyd/software.html

26. Grant, M.; Boyd, S. Graph implementations for nonsmooth convex programs. In: Blondel, V.; Boyd, S.; Kimura, H., editors. Recent Advances in Learning and Control. Springer-Verlag Limited; 2008. p. 95-110.Lecture Notes in Control and Information Scienceshttp://stanford.edu/~boyd/graphdcp.html

27. Rockafellar, RT.; Wets, RJB. Variational Analysis. Springer-Verlag; 1997.

Appendix A. Continuity of (z)

The proof of Theorem 3 relies on the continuity of $\mathcal{M}(\cdot)$ as a map from $\mathbb{R}^M$ into $\mathcal{C}$. We establish continuity by showing that for all $z_n \to z$ as $n \to \infty$, $\mathcal{M}(z_n) \to \mathcal{M}(z)$. We will show that every convergent subsequence of $\mathcal{M}(z_n)$ converges to $\mathcal{M}(z)$. Since $\mathcal{C}$ is non-empty, there exists $\bar{x} \in \mathcal{C}$ such that

$$\left\| \mathscr{M}(z_n) \right\|_{w(z_n)}^2 \leq \left\| \overline{x} \right\|_{w(z_n)}^2 \leq \frac{\left\| \overline{x} \right\|_2^2}{\varepsilon^\nu}. \quad (64)$$

On the other hand, z_n is bounded because it is convergent. This implies that there exists B such that $\max_i |z_{ni}| \leq B$, so that $\|\mathcal{M}(z_n)\|_{w(z_n)}^2 \geq \frac{1}{(\varepsilon^2 + B^2)^{1-\nu/2}} \|\mathcal{M}(z_n)\|_2^2$. Therefore,

$$\|\mathcal{M}(z_n)\|_2^2 \leq \frac{(\varepsilon^2 + B^2)^{1-\nu/2}}{\varepsilon^\nu} \|\bar{x}\|_2^2, \quad (65)$$

so that $\mathcal{M}(z_n)$ is uniformly bounded. Therefore, there exists a convergent subsequence of $\mathcal{M}(z_n)$. Now, let $\mathcal{M}(z_{n_k})$ be *any* convergent subsequence, and let $\mathcal{M}_0$ be its limit. By definition of $\mathcal{M}(z_{n_k})$ and results from convex optimization [18], for each n_k, $\mathcal{M}(z_{n_k})$ is the unique element of $\mathcal{C}$ satisfying

$$\left\langle \frac{\mathcal{M}(z_{n_k})}{(\varepsilon^2 + z_{n_k}^2)^{(1-\nu/2)}}, x - \mathcal{M}(z_n) \right\rangle \geq 0, \text{ for all } x \in \mathcal{C}. \quad (66)$$

Taking limits and invoking continuity of the inner-product, we obtain

$$\left\langle \frac{\mathcal{M}_0}{(\varepsilon^2 + z^2)^{(1-\nu/2)}}, x - \mathcal{M}_0 \right\rangle \geq 0, \text{ for all } x \in \mathcal{C}. \quad (67)$$

Continuity of $\mathcal{M}(\cdot)$ follows from the fact that $\mathcal{M}(z)$ is the *unique* element of $\mathcal{C}$ which satisifies

$$\left\langle \frac{\mathcal{M}(z)}{(\varepsilon^2 + z^2)^{(1-\nu/2)}}, x - \mathcal{M}(z) \right\rangle \geq 0, \text{ for all } x \in \mathcal{C}. \quad (68)$$

Therefore, $\mathcal{M}_0 = \mathcal{M}(z)$, which establishes the continuity of $\mathcal{M}(\cdot)$.

Appendix B. Iteration-dependent choices of ε

Let $\{\varepsilon^{(\ell)} > 0\}_{\ell=1}^\infty$ be a non-increasing sequence such that $\lim_{\ell \to \infty} \varepsilon^{(\ell)} = \bar{\varepsilon} \geq 0$. In this case, the mapping $\mathcal{M}$ must be substituted by the mapping $\mathcal{M}^{(\ell)} \mathbb{R}^M \mapsto \mathbb{R}^M$:

$$M^{(\ell)}(z) := \arg\min_{x \in C} \|x\|_{w^{(\ell)}(z)}^2 \quad (69)$$

for all $z \in \mathbb{R}^M$, where

$$w_i^{(\ell)}(z) := \frac{1}{\left(\varepsilon^{(\ell)^2} + z_i^2\right)^{1-\nu/2}}, \quad \text{for } i = 1, 2, \cdots, M. \quad (70)$$

The main difference in the proof is in part (ii), where it is shown that $\bar{x}$ is a fixed point of the mapping $\mathcal{M}$. We consider two cases: 1) $\bar{\varepsilon} > 0$ and, 2) $\bar{\varepsilon} = 0$.

Case 1

Suppose that $\bar{\varepsilon} > 0$ and that $\{x^{(\ell_k)}\}_k$ is a converging subsequence of the IRLS iterates with the limit $\bar{x}$. Since the sequence $\{x^{(\ell)}\}$ is bounded, there exists an L, such that for all $\ell > L$, all the iterates $x^{(\ell)}$ lie in a bounded and closed ball $B_0 \subset \mathbb{R}^M$. Moreover, let

$$\mathcal{M}(\overline{x}) := \arg\min_{x \in \mathcal{C}} \sum_{i=1}^{M} \frac{x_i^2}{\left(\bar{\varepsilon}^2 + \overline{x}_i^2\right)^{1-\nu/2}}. \quad (71)$$

Clearly, $\mathcal{M}(\bar{x})$ is bounded (since the true vector $x \in \mathcal{C}$ is bounded). Let $B \subset \mathbb{R}^M$ be a closed ball in $\mathbb{R}^M$ such that $B_0 \subseteq B$ and $\mathcal{M}(\bar{x}) \in B$. Then, we have

$$
\begin{aligned}
\overline{x} = \limsup_{k \to \infty} x^{(\ell_k+1)} \quad &= \limsup_{k \to \infty} \mathcal{M}^{(\ell_k)}\left(x^{(\ell_k)}\right) \\
&= \limsup_{k \to \infty} \arg\min_{x \in \mathcal{C}} \|x\|_{w^{(\ell_k)}(x^{(\ell_k)})}^2 \quad (72) \\
&= \limsup_{k \to \infty} \arg\min_{x \in \mathcal{C} \cap B} \|x\|_{w^{(\ell_k)}(x^{(\ell_k)})}^2
\end{aligned}
$$

Now, recall that

$$\|x\|_{w^{(\ell_k)}(x^{(\ell_k)})}^2 = \sum_{i=1}^{M} \frac{x_i^2}{\left(\varepsilon^{(\ell_k)^2} + x_i^{(\ell_k)^2}\right)^{1-\nu/2}} \quad (73)$$

It is easy to show that the function $\|x\|_{w^{(\ell_k)}(x^{(\ell_k)})}^2$ is uniformly convergent to

$$\|x\|_{w^{(\ell_k)}(x^{(\ell_k)})}^2 \longrightarrow \sum_{i=1}^{M} \frac{x_i^2}{\left(\bar{\varepsilon}^2 + \overline{x}_i^2\right)^{1-\nu/2}} \quad (74)$$

for all $x \in B$. To see this, note that

$$\left| \frac{x_i^2}{\left(\varepsilon^{(\ell_k)^2} + x_i^{(\ell_k)^2}\right)^{1-\nu/2}} - \frac{x_i^2}{\left(\bar{\varepsilon}^2 + \overline{x}_i^2\right)^{1-\nu/2}} \right| \leq x_i^2 \left(L_{\varepsilon^{(\ell_k)}} |x_i^{(\ell_k)} - \overline{x}_i| + L_{\overline{x}_i} |\varepsilon^{(\ell_k)} - \overline{\varepsilon}| \right) \quad (75)$$

where L_t denotes the Lipschitz constant of the function $(x^2 + t^2)^{\nu/2-1}$. Since $\varepsilon^{(\ell_k)}, \bar{\varepsilon} > 0$, the Lipschitz constants are uniformly bounded. Moreover, since $x \in B$, then each x_i^2 is bounded, hence the uniform convergence of the function $\|x\|_{w^{(\ell_k)}(x^{(\ell_k)})}^2$ is implied. Given the uniform convergence, a result from variational analysis (Theorem 7.33 of [27]) establishes that

$$\limsup_{k\to\infty} \arg\min_{x\in\mathscr{C}\cap B}\|x\|^2_{w^{(\ell_k)}(x^{(\ell_k)})} \subseteq \arg\min_{x\in\mathscr{C}\cap B}\|x\|^2_{w(\overline{x})} \quad (76)$$

where $w_i(\overline{x}):=(\overline{\varepsilon}^2+\overline{x}_i^2)^{\nu/2-1}$. Note that the minimizer of $\|x\|^2_{w(\overline{x})}$ over the convex set $\mathscr{C}\cap B$ is unique. Therefore, the above inclusion is in fact equality. Hence,

$$\overline{x}=\arg\min_{x\in\mathscr{C}\cap B}\|x\|^2_{w(\overline{x})}=\arg\min_{x\in\mathscr{C}}\|x\|^2_{w(\overline{x})}=\mathscr{M}(\overline{x}) \quad (77)$$

by the construction of B. The rest of the proof remains the same by substituting $\overline{\varepsilon}$ for ε.

Case 2

Suppose that $\overline{\varepsilon}=0$ and that $\mathrm{supp}(x)\overline{}=:T\subseteq\{1, 2, \cdots, M\}$. In this case, if $T \subsetneq \{1, 2, \cdots, M\}$, the limit $\lim_{\ell\to\infty}\mathscr{M}^{(\ell)}(z)$ does not exist. Hence, the proof technique used for Case 1 does no longer hold. However, with a careful examination of the limiting behavior of the mapping $\mathscr{M}^{(\ell)}(z)$, we will show that $x\overline{}$ is a fixed point of the mapping:

$$\mathscr{M}_T(z):= \arg\min_{x\in\mathscr{C}\cap\{x:x_{T^c}=0\}}\sum_{i\in T}\frac{x_i^2}{|z_i|^{2-\nu}} \quad (78)$$

for $|z_i| > 0$ for all $i \in T$. Due to the closedness of $\mathscr{C}$, $x\overline{}\in\mathscr{C}$. So, the set $\mathscr{C}\cap\{x : x_{T^c}=0\}$ is non-empty. If this set is a singleton $\{x\}\overline{}$, then $x\overline{}$ is clearly the fixed point. If not, then there exists $z\in\mathscr{C}\cap\{x : x_{T^c}=0\}$ such that $z \neq x$. Then, the necessary conditions for each minimization at step ℓ_k gives:

$$\sum_{i=1}^M\frac{x_i^{(\ell_k+1)}}{\left(\varepsilon^{(\ell_k)^2}+x_i^{(\ell_k)^2}\right)^{1-\nu/2}}(z_i-x_i^{(\ell_k+1)}) \geq 0. \quad (79)$$

First consider the terms over T. We have:

$$\lim_{k\to\infty}\sum_{i\in T}\frac{x_i^{(\ell_k+1)}}{\left(\varepsilon^{(\ell_k)^2}+x_i^{(\ell_k)^2}\right)^{1-\nu/2}}(z_i-x_i^{(\ell_k+1)})=\sum_{i\in T}\mathrm{sgn}(\overline{x}_i)|\overline{x}_i|^{\nu-1}(z_i-\overline{x}_i) \quad (80)$$

Next, consider the terms over T^c:

$$\lim_{k\to\infty}\sum_{i\in T^c}\frac{-\left(x_i^{(\ell_k+1)}\right)^2}{\left(\varepsilon^{(\ell_k)^2}+x_i^{(\ell_k)^2}\right)^{1-\nu/2}}=\lim_{k\to\infty}\sum_{i\in T^c}-|\overline{x}_i|^{\nu}=0 \quad (81)$$

Hence, we have:

$$\lim_{k\to\infty}\sum_{i=1}^{M}\frac{x_i^{(\ell_k+1)}}{\left(\varepsilon^{(\ell_k)^2}+x_i^{(\ell_k)^2}\right)^{1-\nu/2}}(z_i-x_i^{(\ell_k+1)})=\sum_{i\in T}\mathrm{sgn}(\overline{x}_i)|\overline{x}_i|^{\nu-1}(z_i-\overline{x}_i)\geq 0 \tag{82}$$

which is the necessary and sufficient condition for x being a fixed point of $\mathcal{M}_T(z)$ over all $z\in \mathcal{C}\cap\{x:x_{T^c}=0\}$. Similar to the proof the Theorem 3, it can be shown that x satisfies the necessary optimality conditions for the function:

$$f_\nu^T(z):=\sum_{i\in T}|z_i|^\nu=\|z_T\|_\nu^\nu \tag{83}$$

over the set $\mathcal{C}\cap\{x:x_{T^c}=0\}$. This concludes the proof. Note that the case of $\overline{\varepsilon}=0$ is not favorable in general. Carefully chosen sequences $\{\varepsilon^{(\ell)}\}$ as in [9] together with the assumption that the s-sparse vector $x\in\mathcal{C}$ is unique, can result in convergence of the IRLS to the true s-sparse x. However, for general sequences of $\{\varepsilon^{(\ell)}\}$ with $\lim_{\ell\to\infty}\varepsilon^{(\ell)}=0$, this is not necessarily the case.

$\epsilon = 0.1$

$\epsilon = 0$

$\nu = 0.5$

$\nu = 1$
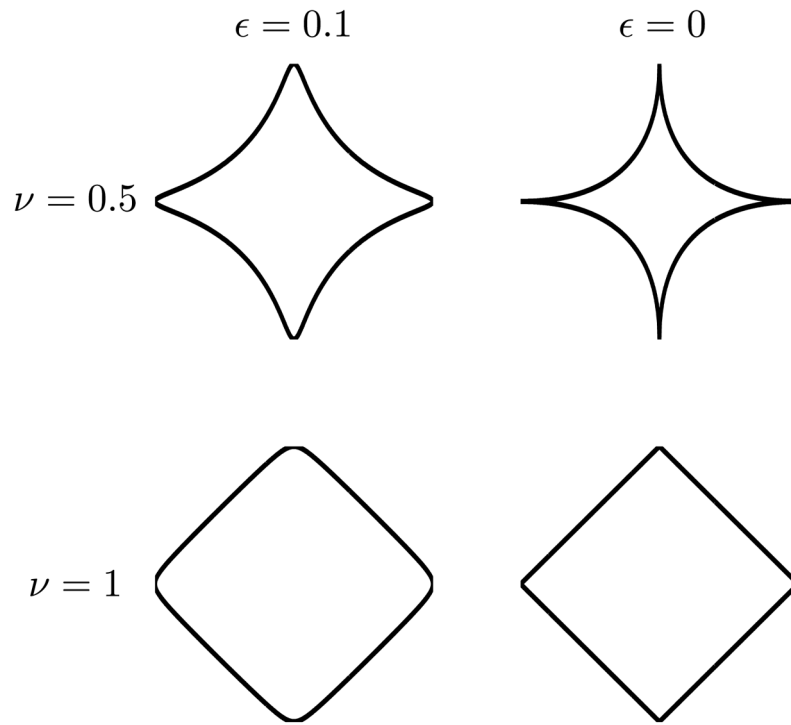
Fig. 1.
Level sets of ℓ_ν balls and their ε-smooth counterparts, which the IRLS(ν, ε) algorithm maximizes subject to fidelity constraints on the signal reconstruction.
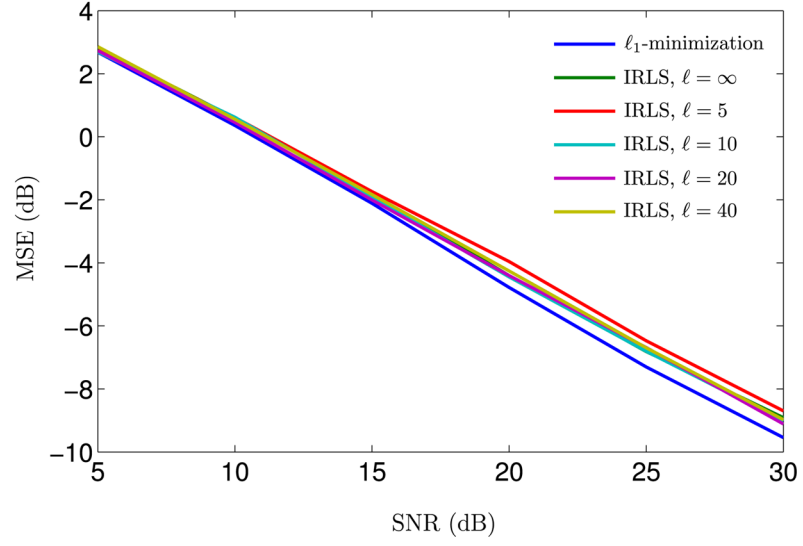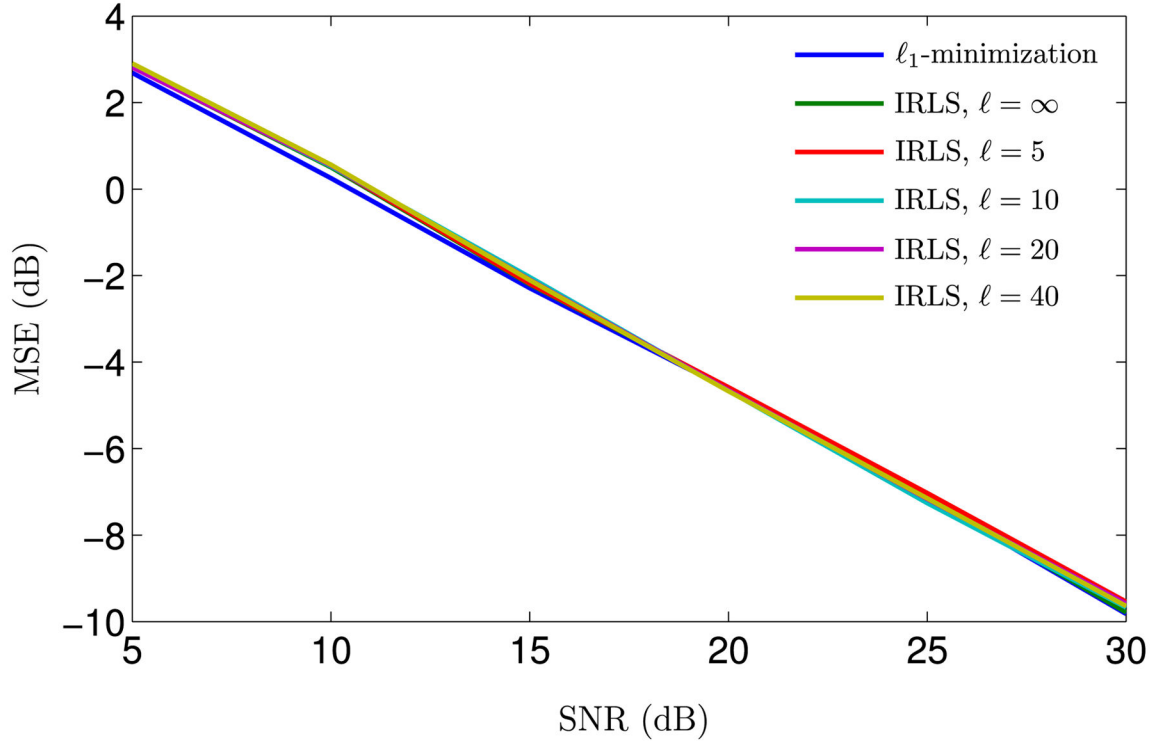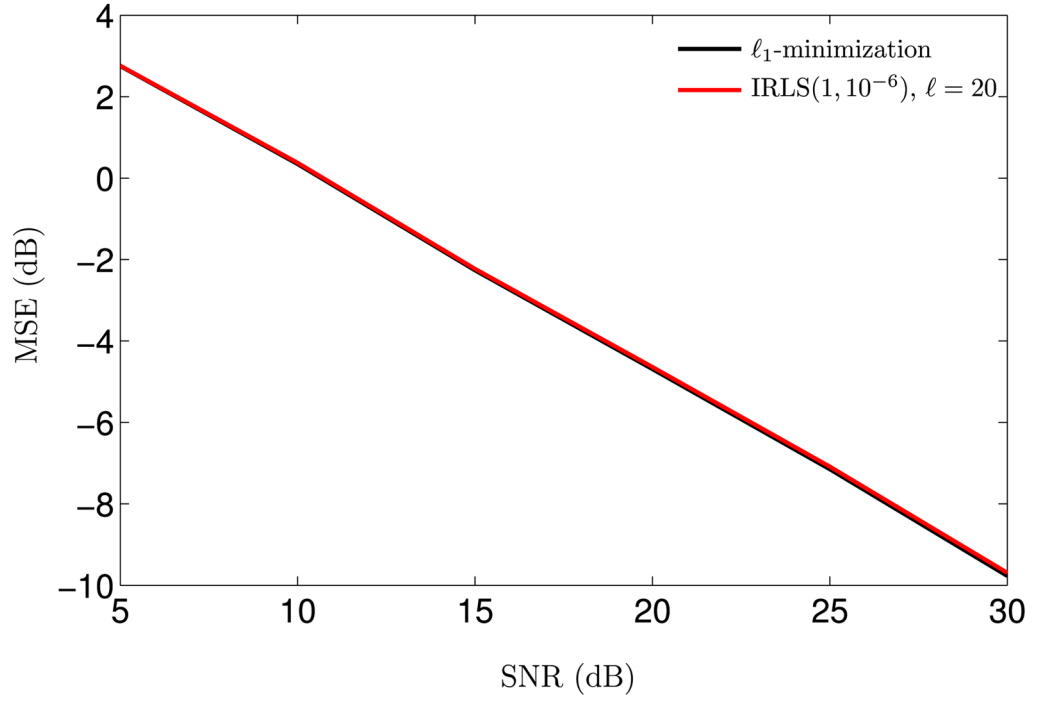
Fig. 2.
Mean square error ($20 \log_{10} \|x - \bar{x}\|_2$) as a function of ℓ and SNR, for $\nu = 1$ and $\varepsilon(5) = 10^{-4}$. The figure shows that the stability of IRLS compares favorably to that of ℓ_1-minimization. More importantly, only a few number of IRLS iterations are required to reach a satisfactory value of the MSE.

Fig. 3.
Mean square error ($20 \log_{10} \|x \bar{-} x\|_2$) as a function of ℓ and SNR, for $\nu = 1/2$ and $\varepsilon(5) =$ 10^{-4}. This figure highlights the sparsifying properties of $f_\nu(\cdot)$-minimization by IRLS for $\nu <$ 1. As in the previous figure, only a few number of IRLS iterations are required to reach a satisfactory value of the MSE.

Fig. 4.
Mean square error ($20 \log_{10} \|x - x\|_2$) as a function of ℓ and SNR, for $\nu = 1$ and $\varepsilon(5) = 10^{-6}$. The figure shows that, as one decreases ε, the approximation to the ℓ_1-norm improves.
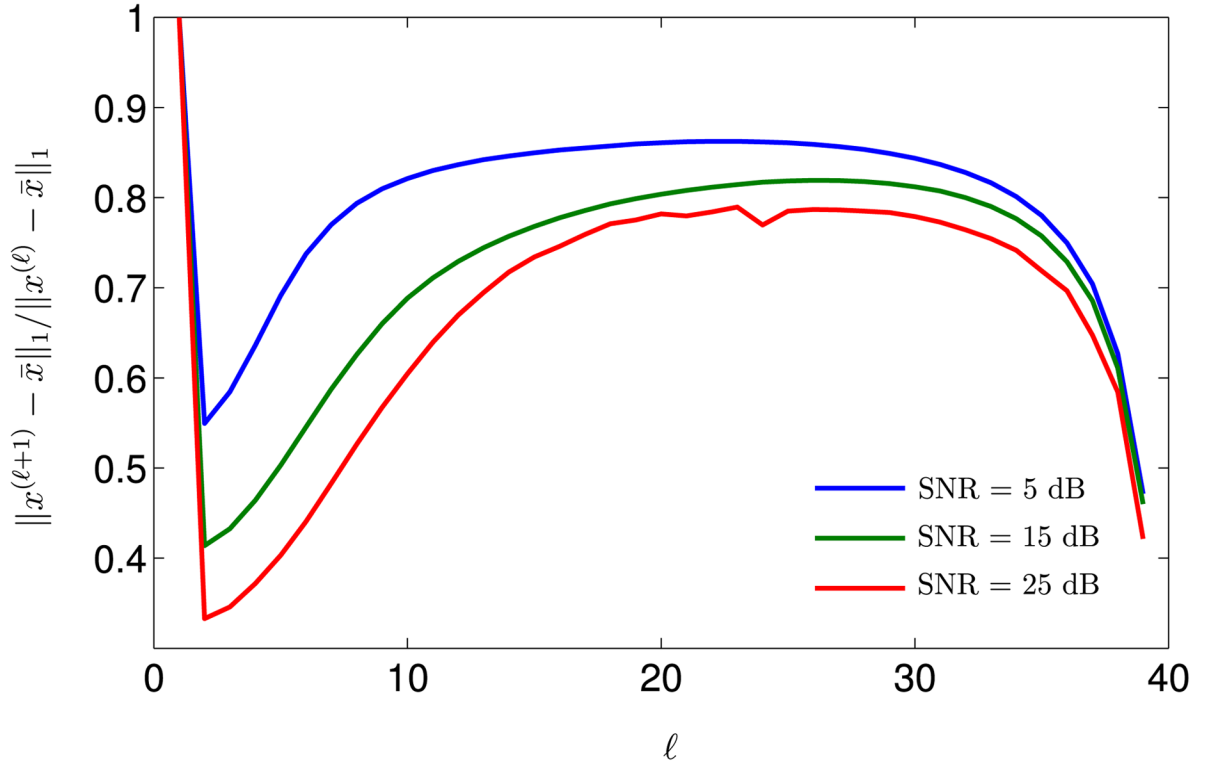
Fig. 5.
Ratio $\|x^{(\ell+1)} - x\|_1^-/\|x^{(\ell)} - x\|_1^-$ as a function of ℓ and SNR, for $\nu = 1$ and $\varepsilon(5) = 10^{-4}$. The figure shows that the IRLS converges with a relatively small number of iterations (ℓ).
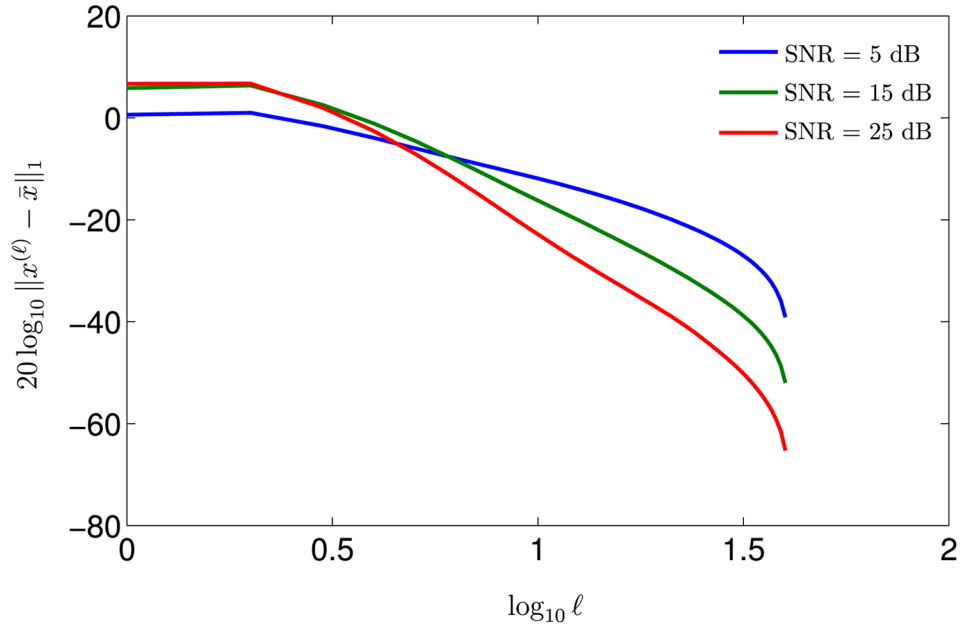
Fig. 6.
$20 \log_{10} \|x^{(\ell)} - \bar{x}\|_1$ as a function of $\log_{10} \ell$ and SNR, for $\nu = 1$ and fixed $\varepsilon(5) = 10^{-4}$. The figure shows that the IRLS converges with a relatively small number of iterations (ℓ).
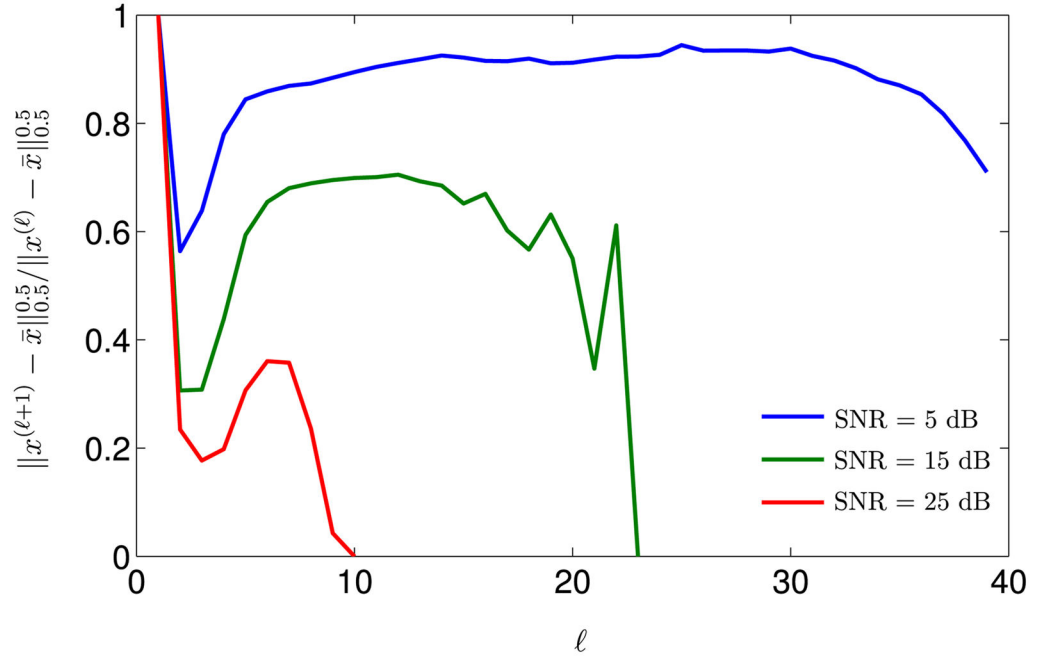
Fig. 7.

Ratio $\|x^{(\ell+1)} - \overline{x}\|_{0.5}^{0.5} / \|x^{(\ell)} - \overline{x}\|_{0.5}^{0.5}$ as a function of ℓ and SNR, for $\nu = 0.5$ and fixed $\varepsilon(5) = 10^{-4}$. The figure shows that the IRLS converges with a relatively small number of iterations (ℓ).
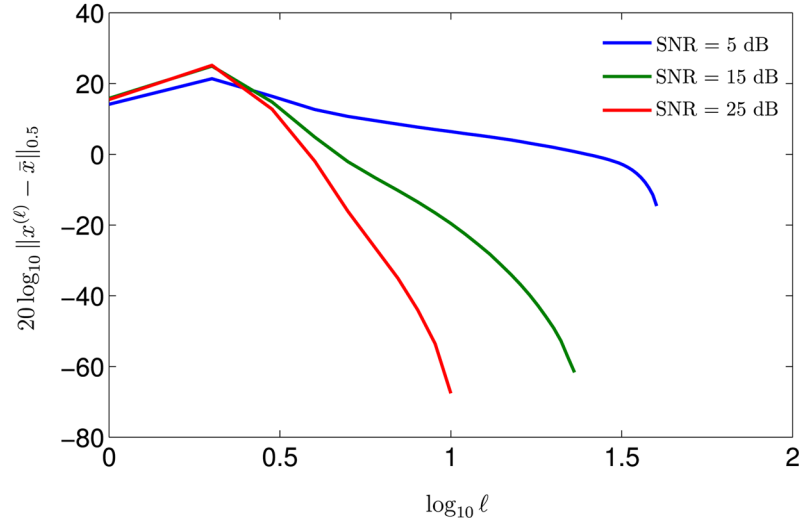
Fig. 8.
$20 \log_{10} \|x^{(\ell)} - \bar{x}\|_{0.5}$ as a function of $\log_{10} \ell$ and SNR, for $\nu = 0.5$ and fixed $\varepsilon = 10^{-4}$. The figure shows that the IRLS converges with a relatively small number of iterations (ℓ).