
Iteratively Reweighted Least Squares for Basis Pursuit with Global Linear Convergence Rate

Christian Kümmerle^{*§}

Department of Applied Mathematics & Statistics
Johns Hopkins University

Claudio Mayrink Verdun^{†§}

Department of Mathematics and Department of Electrical and Computer Engineering
Technical University of Munich

Dominik Stöger^{‡§}

Department of Mathematics
KU Eichstätt–Ingolstadt

Abstract

The recovery of sparse data is at the core of many applications in machine learning and signal processing. While such problems can be tackled using ℓ_1 -regularization as in the LASSO estimator and in the Basis Pursuit approach, specialized algorithms are typically required to solve the corresponding high-dimensional non-smooth optimization for large instances. Iteratively Reweighted Least Squares (IRLS) is a widely used algorithm for this purpose due to its excellent numerical performance. However, while existing theory is able to guarantee convergence of this algorithm to the minimizer, it does not provide a global convergence rate. In this paper, we prove that a variant of IRLS converges *with a global linear rate* to a sparse solution, i.e., with a linear error decrease occurring immediately from any initialization, if the measurements fulfill the usual null space property assumption. We support our theory by numerical experiments showing that our linear rate captures the correct dimension dependence. We anticipate that our theoretical findings will lead to new insights for many other use cases of the IRLS algorithm, such as in low-rank matrix recovery.

1 Introduction

The field of sparse recovery deals with the problem of recovering an (approximately) sparse vector x from only few linear measurements, presented by an underdetermined system of linear equations of the form $y = Ax$. One approach to solve this problem is to consider the ℓ_0 -minimization under linear constraints, which is NP-hard in general [24, 54]. For computational reasons, instead of ℓ_0 -minimization, it is common practice to consider its convex relaxation

$$\min_{x \in \mathbb{R}^N} \|x\|_1 \quad \text{subject to } Ax = y, \quad (P_1)$$

*kummerle@jhu.edu

†verdun@ma.tum.de

‡dominik.stoeger@ku.de

§The name order of the authors is alphabetical.

where $A \in \mathbb{R}^{m \times N}$, $b \in \mathbb{R}^m$ are given, which is referred to as ℓ_1 (-norm) *minimization* [20, 27, 64] or *basis pursuit* [18, 19] in the literature.

Unlike ℓ_0 -minimization, the optimization program (P_1) is computationally tractable in general, and a close relationship of their minimizers has been recognized and well-studied in the theory of compressive sensing [12, 13, 26, 32]. In statistics and machine learning, an unconstrained variant of (P_1), often called *LASSO*, amounts to the most well-studied tractable estimator for variable selection in high-dimensional inference [37, 48, 65]. ℓ_1 -minimization has many other applications and it was even called the *modern least squares* [14].

The tractability of (P_1) becomes evident from the fact that it can be reformulated as a linear program [66]. However, as many problems of interest in applications are high-dimensional and therefore challenging for standard linear programming methods, many specialized solvers for (P_1) have been proposed, such as the Homotopy Method [28], primal-dual methods [15, 52], Alternating Direction Method of Multipliers [9], Bregman iterative regularization [71] and Semismooth Newton Augmented Lagrangian Methods [46] and Iteratively Reweighted Least Squares (IRLS), the latter of which is in the focus of this paper.

Iteratively Reweighted Least Squares corresponds to a family of algorithms that minimizes non-smooth objective functions by solving a sequence of quadratic problems, with its idea going back to a method proposed by Weiszfeld for the *Fermat-Weber problem* [7, 69]. A variety of different problems such as robust regression in statistics [39, 53], total variation regularization in image processing [2, 33, 55], joint learning of neural networks [72], robust subspace recovery [45] and the recovery of low-rank matrices [31, 42, 43, 51] can be solved efficiently by IRLS in practice, as it relies on simple linear algebra to solve the linear systems arising from the quadratic problems at each iteration, without the need of a careful initialization or intricate parameter tuning. On the other hand, the analysis of IRLS methods is typically challenging: General convergence results are often weak, and stronger convergence results are only available in particular cases; see Section 2.3 for more details.

IRLS for sparse recovery. In the sparse recovery context, the first variants of IRLS were introduced in [34, 59] for the ℓ_p -quasinorm minimization problem (P_p) with $0 < p \leq 1$ that is similar to (P_1), but with $\|x\|_p$ instead of $\|x\|_1$ as an objective. In [16], modifications of the method of [34, 59] using specific smoothing parameter update rules were observed to exhibit excellent numerical performance for solving (P_p), retrieving the underlying sparse vector when most of the methods fail. A useful fact is that IRLS is one of the few methods (ADMM being the other one [9]) that provides a framework to solve both constrained and unconstrained formulations of ℓ_p -minimization problems.

A major step forward in the theoretical understanding of IRLS was achieved in [23], where the authors showed that a variant of IRLS for (P_1) converges globally to an ℓ_1 -minimizer if the measurement operator A fulfills the null space property of sufficient order, which essentially ensures that an ℓ_1 -minimizer is actually sparse. However, since this proof relies on the existence of a convergent subsequence, their proof does not reveal any rate for *global* convergence. The analysis of [23] provides, furthermore, a *locally* linear convergence rate, but this local linear rate has the drawback that it only applies *if the support of the true signal has been discovered*, which is arguably the difficult part of ℓ_0 -minimization—cf. Proposition 3.1 below and Section 4.1.

While several extensions and modifications of the IRLS algorithm in [23] have been proposed (see, e.g., [3, 30]), this following fundamental algorithmic question has remained unanswered:

What is the global convergence rate of the IRLS algorithm for ℓ_1 -minimization?

Our contribution. We resolve this question, formally stated in [62], and present a new IRLS algorithm that *converges linearly* to a sparse ground truth, *starting from any initialization*, as stated in Theorem 3.2. Our algorithm returns a feasible solution with δ -accuracy, i.e., $\|x_* - x^k\|_1 \leq \delta$, where x_* is the underlying s -sparse vector, in $k = O(N\sqrt{(\log N)/m} \log(1/\delta))$ iterations. Analogous to [23], it is assumed that the measurement matrix A satisfies the so-called null space property [21]. We also provide a similar result for approximately sparse vectors. Our proof relies on a novel quantification of the descent of a carefully chosen objective function in the direction of the ground truth. Additionally, we support the theoretical claims by numerical simulations indicating that we capture the correct dimension dependence. We believe that the new analysis techniques in this paper are of independent interest and will pave the way for establishing global convergence rates for other variants of IRLS such as in low-rank matrix recovery [31].

Notation. We denote the cardinality of a set I by $|I|$ and the support of a vector $x \in \mathbb{R}^N$, i.e., the index set of its nonzero entries, by $\text{supp}(x) = \{j \in [N] : x_j \neq 0\}$. We call a vector s -sparse if at most s of its entries are nonzero. We denote by x_I the restriction of x onto the coordinates indexed by I , and use the notation $I^c := [N] \setminus I$ to denote the complement of a set I . Furthermore, $\sigma_s(x)_{\ell_1}$ denotes the ℓ_1 -error of the best s -term approximation of a vector $x \in \mathbb{R}^N$, i.e., $\sigma_s(x)_{\ell_1} = \inf\{\|x - z\|_1 : z \in \mathbb{R}^N \text{ is } s\text{-sparse}\}$.

2 IRLS for sparse recovery

We now present a simple derivation of the Iteratively Reweighted Least Squares (IRLS) algorithm for ℓ_1 -minimization which is studied in this paper. IRLS algorithms can be interpreted as a variant of a Majorize-Minimize (MM) algorithm [63], as we will lay out in the following. It mitigates the non-smoothness of the $\|\cdot\|_1$ -norm by using the smoothed objective function $\mathcal{J}_\varepsilon : \mathbb{R}^N \rightarrow \mathbb{R}$, which is defined, for a given $\varepsilon > 0$, by

$$\mathcal{J}_\varepsilon(x) := \sum_{i=1}^N j_\varepsilon(x_i) \quad \text{with} \quad j_\varepsilon(x) := \begin{cases} |x|, & \text{if } |x| > \varepsilon, \\ \frac{1}{2} \left(\frac{x^2}{\varepsilon} + \varepsilon \right), & \text{if } |x| \leq \varepsilon. \end{cases} \quad (1)$$

The function \mathcal{J}_ε can be considered as a scaled Huber loss function which is widely used in robust regression analysis [40, 50]. Moreover, the function \mathcal{J}_ε is continuously differentiable and fulfills $|x| \leq j_\varepsilon(x) \leq |x| + \varepsilon$ for each $x \in \mathbb{R}$. Instead of minimizing the function \mathcal{J}_ε directly, the idea of IRLS is to minimize instead a suitable chosen quadratic function $Q_\varepsilon(\cdot, x)$, which majorizes \mathcal{J}_ε such that $Q_\varepsilon(z, x) \geq \mathcal{J}_\varepsilon(z)$ for all $z \in \mathbb{R}^N$. This function is furthermore chosen such that $Q_\varepsilon(x, x) = \mathcal{J}_\varepsilon(x)$ holds, which implies that $\min_{z \in \mathbb{R}^n} Q_\varepsilon(z, x) \leq \mathcal{J}_\varepsilon(x)$. The latter inequality implies that by minimizing $Q_\varepsilon(\cdot, x)$, IRLS actually achieves an improvement in the value of \mathcal{J}_ε as well. More specifically, $Q_\varepsilon(\cdot, x)$ is defined by

$$\begin{aligned} Q_\varepsilon(z, x) &:= \mathcal{J}_\varepsilon(x) + \langle \nabla \mathcal{J}_\varepsilon(x), z - x \rangle + \frac{1}{2} \langle (z - x), \text{diag}(w_\varepsilon(x))(z - x) \rangle \\ &= \mathcal{J}_\varepsilon(x) + \frac{1}{2} \langle z, \text{diag}(w_\varepsilon(x))z \rangle - \frac{1}{2} \langle x, \text{diag}(w_\varepsilon(x))x \rangle, \end{aligned} \quad (2)$$

where $\nabla \mathcal{J}_\varepsilon(x) = \left(\left\{ \begin{array}{l} \frac{x_i}{|x_i|}, \\ \frac{x_i}{\varepsilon}, \end{array} \right. \right)_{i=1}^N$ is the gradient of \mathcal{J}_ε at x and the weight vector

$w_\varepsilon(x) \in \mathbb{R}^N$ is a vector of *weights* such that $w_\varepsilon(x)_i := [\max(|x_i|, \varepsilon)]^{-1}$ for $i \in [N]$. The following lemma shows that $Q_\varepsilon(\cdot, \cdot)$ has indeed the above-mentioned properties. We refer to the supplementary material for a proof.

Lemma 2.1. *Let $\varepsilon > 0$, let $\mathcal{J}_\varepsilon : \mathbb{R}^N \rightarrow \mathbb{R}$ be defined as in (1) and $Q_\varepsilon : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$ as defined in (2). Then, for any $z, x \in \mathbb{R}^N$, the following affirmations hold:*

$$i. \text{diag}(w_\varepsilon(x))x = \nabla \mathcal{J}_\varepsilon(x), \quad ii. Q_\varepsilon(x, x) = \mathcal{J}_\varepsilon(x), \quad iii. Q_\varepsilon(z, x) \geq \mathcal{J}_\varepsilon(z).$$

As can be seen from the equality in (2), minimizing $Q_\varepsilon(\cdot, x)$ corresponds to a minimization of a (re-)weighted least squares objective $\langle \cdot, \text{diag}(w_\varepsilon(x))\cdot \rangle$, which lends its name to the method. Note that unlike a classical MM approach, however, IRLS comes with an *update step of the smoothing parameter* ε at each iteration. We provide an outline of the method in Algorithm 1.

The weighted least squares update (3) can be computed such that $x^{k+1} = W_k^{-1} A^* (A W_k^{-1} A^*)^{-1} (y)$ with $W_k = \text{diag}(w_k)$, with the solution of the $(m \times m)$ linear system $(A W_k^{-1} A^*)z = y$ as a main computational step. This linear system is positive definite and suitable for the use of iterative solvers. In [30], an analysis of how accurately the linear system of a similar IRLS method needs to be solved to ensure overall convergence. We note that for small ε_k , the Sherman-Woodbury formula [70] can be used so that the calculation of x^{k+1} boils down to solving a smaller linear system that is well-conditioned, c.f. the supplementary material for details. This numerically advantageous property is not shared by the methods of [3, 23, 30], as our smoothing update (4) is slightly different from the ones proposed in these papers. We refer to Section 2.2 for a discussion.

The update step of the smoothing parameter ε (4) for the IRLS algorithm under consideration requires an a priori estimate of the sparsity of the ground truth of the signal, a piece of information that is

Algorithm 1 Iteratively Reweighted Least Squares for ℓ_1 -minimization

Input: Measurement matrix $A \in \mathbb{R}^{m \times N}$, data vector $y \in \mathbb{R}^m$,
initial weight vector $w_0 \in \mathbb{R}^N$ (default: $w_0 = (1, 1, \dots, 1)$).

Set $\varepsilon_0 = \infty$.

for $k = 0, 1, 2, \dots$ **do**

$$x^{k+1} := \arg \min_{z \in \mathbb{R}^N} \langle z, \text{diag}(w_k) z \rangle \quad \text{subject to} \quad Az = y, \quad (3)$$

$$\varepsilon_{k+1} := \min \left(\varepsilon_k, \frac{\sigma_s(x^{k+1})_{\ell_1}}{N} \right), \quad (4)$$

$$(w_{k+1})_i := \frac{1}{\max(|x_i^{k+1}|, \varepsilon_{k+1})} \quad \text{for each } i \in [N], \quad (5)$$

end for

return Sequence $(x^k)_{k \geq 1}$.

also needed by most of the methods for sparse reconstruction. In practice, an overestimation of s is not a problem for similar numerical results if the overestimation remains within small multiples of the sparsity of the signal. We note, however, that there are also versions of IRLS which do not require a-priori knowledge of s , e.g. [30, 68], as the update rule for the smoothing parameter is chosen differently. An interesting future research direction is to extend the analysis presented here to IRLS with such a smoothing parameter update.

A consequence of Lemma 2.1, step (4), the fact that $\varepsilon \mapsto \mathcal{J}_\varepsilon(z)$ is monotonously non-decreasing, and that $k \mapsto \varepsilon_k$ is non-increasing is that $k \mapsto \mathcal{J}_\varepsilon(z)$ is non-increasing in k . This implies that the iterates x^k, x^{k+1} of Algorithm 1 fulfill

$$\mathcal{J}_{\varepsilon_{k+1}}(x^{k+1}) \leq \mathcal{J}_{\varepsilon_k}(x^{k+1}) \leq Q_{\varepsilon_k}(x^{k+1}, x^k) \leq Q_{\varepsilon_k}(x^k, x^k) = \mathcal{J}_{\varepsilon_k}(x^k). \quad (6)$$

This shows in particular that the sequence $\{\mathcal{J}_{\varepsilon_k}(x^k)\}_{k=0}^\infty$ is non-increasing. For this reason, it can be shown that each accumulation point of the sequence of iterates $(x^k)_{k \geq 0}$ is a (first-order) stationary point of the smoothed ℓ_1 -objective $\mathcal{J}_{\bar{\varepsilon}}(\cdot)$ subject to the measurement constraint imposed by A and y , where $\bar{\varepsilon} = \lim_{k \rightarrow \infty} \varepsilon_k$ (see [23, Theorem 5.3]).

2.1 Null space property

As in [23], the analysis we present is based on the assumption that the measurement matrix A satisfies the so-called null space property [21, 36], which is a key concept in the compressed sensing literature (see, e.g., [32, Chapter 4] for an overview).

Definition 2.2. A matrix $A \in \mathbb{R}^{m \times N}$ is said to satisfy the ℓ_1 -null space property (ℓ_1 -NSP) of order $s \in \mathbb{N}$ with constant $0 < \rho_s < 1$ if for any set $S \subset [N]$ of cardinality $|S| \leq s$, it holds that $\|v_S\|_1 \leq \rho_s \|v_{S^c}\|_1$, for all $v \in \ker(A)$.

In [32, Chapter 4], the property of Definition 2.2 was called *stable* null space property. The importance of the null space property is due to the fact that it gives a necessary and sufficient criterion for the success of basis pursuit for sparse recovery, as the following theorem shows.

Theorem 2.3 ([32, Theorem 4.5]). Given a matrix $A \in \mathbb{R}^{m \times N}$, every vector $x \in \mathbb{R}^N$ such that $\|x\|_0 \leq s$ is the unique solution of (P_1) with $Ax = y$ if and only if A satisfies the null space property of order s for some $0 < \rho_s < 1$.

The ℓ_1 -NSP is implied by the restricted isometry property (see, e.g., [11]), which is fulfilled by a large class of random matrices with high probability. For example, this includes matrices with (sub-)gaussian entries and random partial Fourier matrices [6, 60].

2.2 Existing theory

A predecessor of IRLS for the sparse recovery problem (P_1) , and more generally, for ℓ_p -quasinorm minimization with $0 < p \leq 1$, is the *FOCal Underdetermined System Solver* (FOCUSS) as proposed by Gorodnitsky, Rao and Kreutz-Delgado [34, 59]. Asymptotic convergence of FOCUSS to a stationary point from any initialization was claimed in [59], but the proof was not entirely accurate, as

pointed out by [17]. One limitation of FOCUSS is that, unlike in IRLS as presented in Algorithm 1, no smoothing parameter ε is used, which leads to ill-conditioned linear systems.

To mitigate this, [16] proposed an IRLS method that uses smoothing parameters ε (such as used in Q_ε defined above) that are updated iteratively. It was observed that this leads to a better condition number for the linear systems to be solved in each step of IRLS and, furthermore, that this smoothing strategy has the advantage of finding sparser vectors if the weights of IRLS are chosen to minimize a non-convex ℓ_p -quasinorm for $p < 1$.

Further progress for IRLS designed to minimize an ℓ_1 -norm was achieved in the seminal paper [23]. In [23], it was shown that if the measurement operator fulfills a suitable ℓ_1 -null space property as in Definition 2.2, an IRLS method with iteratively updated smoothing converges to an s -sparse solution, coinciding with the ℓ_1 -minimizer, if there exists one that is compatible with the measurements. This method uses not exactly the update rule of (4), but rather updates the smoothing parameter such that $\varepsilon_{k+1} = \min(\varepsilon_k, R(x^{k+1})_{s+1}/N)$, where $R(x^{k+1})_{s+1}$ is the $(s+1)^{\text{st}}$ -largest element of the set $\{|x_j^{k+1}|, j \in [N]\}$. Furthermore, a *local linear convergence rate* of IRLS was established [23, Theorem 6.1] under same conditions.

However, the analysis of [23] has its limitations: First, there is a gap in the assumption of their convergence results between the sparsity s of a vector to be recovered and the order \widehat{s} of the NSP of the measurement operator. Recently, this gap was circumvented in [3] with an IRLS algorithm that uses a smoothing update rule based on an ℓ_1 -norm, namely, $\varepsilon_{k+1} = \min(\varepsilon_k, \eta(1 - \rho_s)\sigma_s(x^{k+1})_{\ell_1}/N)$, where $\eta \in (0, 1)$, and ρ_s is the NSP constant of the order s of the NSP fulfilled by the measurement matrix A —this rule is quite similar to the rule (4) that we use in Algorithm 1. In particular, [3, Theorem III.6] establishes convergence with local linear rate similar to [23] without the gap mentioned above. The main limitation, however, of the theory of [23] (which is shared by [3]) is that the linear convergence rate only holds *locally*, i.e., in a situation where the support of the sparse vector has already been identified, see also Section 3 and Section 4.1 for a discussion.

We finally mention three relevant papers for the theoretical understanding of IRLS. [5] established the correspondence between the IRLS algorithms and the Expectation-Maximization algorithm for constrained maximum likelihood estimation under a Gaussian scale mixture distribution. By doing so, they established similar results as those from [23], i.e., the global convergence of IRLS with local linear convergence rate (as can be seen from their equation (38), which similar to (7) below) but by using different techniques based on such correspondence. [62] explores the relationship of IRLS for ℓ_1 -minimization and a slime mold dynamics, interpreting both as an instance of the same meta-algorithm. Without requiring any connection between sparse recovery and ℓ_1 -minimization, [29] shows that an IRLS-like algorithm for (P_1) , requires $O(N^{1/3} \log(1/\delta)/\delta^{2/3} + \log(N)/\delta^2)$ iterations to obtain a multiplicative error of $1 + \delta$ on the minimizer $\|x\|_1$. Unlike our result Theorem 3.2, this corresponds not to a linear, but to a sublinear convergence rate.

2.3 Related work

As mentioned in the introduction, IRLS has a long history and has appeared under different names within different communities, e.g., similar algorithms are usually called *half-quadratic algorithms* in image processing [2, 41] and the *Kačanov method* in numerical PDEs [25]. Probably the most common usage of IRLS has been in robust regression [35, 39], c.f. [10] for a survey that also covers applications in approximation theory. For p -norm regression, [1] proposed a version of IRLS for which convergence results for $p \in [2, \infty)$ were established, solving a problem that was open for over thirty years. Also, for robust regression, by using an ℓ_1 -objective on the residual, [53] showed recently global convergence of IRLS with a linear rate, with high probability for sub-Gaussian data. We note that our proof strategy is different from the one of [53] due to a structural difference of (P_1) from robust regression.

In [57], the authors provide a general framework for formulating IRLS algorithms for the optimization of a quite general class of non-convex and non-smooth functions, however, without updated smoothing. They use techniques developed in [4] to show convergence of the sequence of iterates to a critical point under the Kurdyka-Łojasiewicz property [8]. However, no results about convergence rates were presented.

For the sparse recovery problem, the topic discussed in this paper, the references [30, 44, 68] analyzed IRLS for an unconstrained version of (P_1) , which is usually a preferable formulation if

the measurements are corrupted by noise. Additionally, the work [30] addressed the question of how to solve the successive quadratic optimization problems. The authors developed a theory that shows, under the NSP, how accurately the quadratic subproblems need to be solved via the conjugate gradient method in order to preserve the convergence results established in [23].

Finally, for the related problems of low-rank matrix recovery and completion, IRLS strategies have emerged as one of the most successful methods in terms of data-efficiency and scalability [31, 42, 43, 51].

While we were writing this paper, the manuscript [58] appeared providing new insights about IRLS. It describes a surprisingly simple reparametrization of the IRLS formulation for ℓ_p -minimization (with $p \in (2/3, 1)$) that leads to a smooth bilevel optimization problem without any spurious minima, i.e., the stationary points of this new formulation are either global minima or strict saddles. It is an interesting future direction to explore the connection between this new approach and our global convergence theory.

3 IRLS for Basis Pursuit with Global Linear Rate

As discussed in Section 2.2, the main theoretical advancements for IRLS for the sparse recovery problem were achieved in the work [23].

Proposition 3.1. [23, Theorem 6.1] *Assume that $A \in \mathbb{R}^{m \times N}$ satisfies the NSP of order $\hat{s} > s$ with constant $\rho_{\hat{s}}$ such that $0 < \rho_{\hat{s}} < 1 - \frac{2}{\hat{s}+2}$ and $\hat{s} > s + \frac{2\rho_{\hat{s}}}{1-\rho_{\hat{s}}}$ hold. Let $x_* \in \mathbb{R}^N$ be an s -sparse vector and set $y = Ax_*$. Assume that there exists an integer $k_0 \geq 1$ and a positive number $\xi > 0$ such that*

$$\xi := \frac{\|x^{k_0} - x_*\|_1}{\min_{i \in S} |(x_*)_i|} < 1. \quad (7)$$

Then the iterates $\{x^{k_0}, x^{k_0+1}, x^{k_0+2}, \dots\}$ of the IRLS method in [23] converge linearly to x_ , i.e., for all $k \geq k_0$, the k th iteration of IRLS satisfies*

$$\|x^{k+1} - x_*\|_1 \leq \frac{\rho_{\hat{s}}(1 + \rho_{\hat{s}})}{1 - \xi} \left(1 + \frac{1}{\hat{s} - 1 - s}\right) \|x^k - x_*\|_1. \quad (8)$$

The main contribution of this paper is that we overcome a local assumption such as (7) and show that IRLS as defined by Algorithm 1 exhibits a global linear convergence rate, i.e., there is a linear convergence rate starting from any initialization, as early as in the first iteration.

Exactly sparse case. Our first main result, Theorem 3.2, deals with the scenario that the ground truth vector x_* is exactly s -sparse. Our second result, presented in the supplementary material, generalizes the first one to the more realistic situation of approximately sparse vectors.

Theorem 3.2. *Consider the problem of recovering an unknown s -sparse vector $x_* \in \mathbb{R}^N$ from known measurements of the form $y = Ax_*$. Assume that the measurement matrix $A \in \mathbb{R}^{m \times N}$ fulfills the ℓ_1 -NSP of order s with constant $\rho_s < 1/2$. Let the IRLS iterates $\{x^k\}_k$ and $\{\varepsilon_k\}_k$ be defined by the IRLS algorithm (3) and (4) with initialization x^0 . Then, for all $k \in \mathbb{N}$, it holds that*

$$\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1 \leq \left(1 - \frac{c}{\rho_1 N}\right)^k (\mathcal{J}_{\varepsilon_0}(x^0) - \|x_*\|_1) \quad (9)$$

as well as

$$\|x^k - x_*\|_1 \leq 9 \left(1 - \frac{c}{\rho_1 N}\right)^k \|x^0 - x_*\|_1. \quad (10)$$

Here $c = 1/768$ is an absolute constant and $\rho_1 < 1/2$ denotes the ℓ_1 -NSP constant of order 1.

Inequality (9) says that the difference $\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1$ converges linearly with a uniform upper bound of $1 - \frac{c}{\rho_1 N}$ on the linear convergence factor. As our proof, which is detailed in the supplementary material, shows, this implies inequality (10), which implies that also $\|x_* - x^k\|_1$ exhibits linear convergence in the number of iterations k . In particular, this means that for some error tolerance $\delta > 0$, we obtain $\|x_* - x^k\|_1 \leq \delta$ after $O\left(\rho_1 N \log\left(\frac{\|x_* - x^0\|_1}{\delta}\right)\right)$ iterations.

Remark 3.3. Note that it follows directly from Definition 2.2 that the constant ρ_1 of the ℓ_1 -NSP of order 1 satisfies $\rho_1 \leq \rho_s \leq 1$, which implies that δ -accuracy is obtained after $O\left(N \log\left(\frac{\|x_* - x^0\|_1}{\delta}\right)\right)$ iterations. This bound can be improved in many scenarios where one can obtain more explicit bounds on ρ_1 , for example, when A is a Gaussian matrix. Namely, inspecting [32, p. 142 and Thm. 9.2], we observe in this scenario that $\rho_1 \lesssim \sqrt{(\log N)/m}$ with high probability. Hence, in this scenario, at most $O\left(N \sqrt{\frac{\log N}{m}} \log\left(\frac{\|x_* - x^0\|_1}{\delta}\right)\right)$ iterations are needed to achieve δ -accuracy.

The key idea in our proof is to use fact that the quadratic functional $Q_{\varepsilon_k}(\cdot, x^k)$ approximates the ℓ_1 -norm in a neighborhood of the current iterate x^k . For this reason, we also expect that for $t > 0$ sufficiently small, we have that $Q_{\varepsilon_k}(x^k + tv^k, x^k) < Q_{\varepsilon_k}(x^k, x^k)$ if $v^k = x_* - x^k$ is the vector between x^k and the ground truth x_* . Then by choosing t properly, we can guarantee a sufficient decrease of the functional $\mathcal{J}_{\varepsilon_k}(x^k)$ in each iteration.

In Section 4, we conduct experiments that indeed verify the linear convergence in (9) and (10). Moreover, we study numerically whether one can observe a dependence of the convergence rate on the problem parameters N , s and m . We construct a worst-case example which indicates that the convergence rate indeed may depend on the dimension N in a way as described by (9). In a certain sense, this indicates that there are two convergence phases, a global one, where a dimension-dependent constant cannot be avoided and a local convergence phase, where a local convergence result such as described in Proposition 3.1 kicks in.

Finally, let us mention that we have undertaken no efforts to optimize the constant $c = 1/768$ in Theorem 3.2. Nevertheless, we note that the constant c can be replaced by the sharper constant c_{ρ_s} as defined in Proposition B.3.

4 Numerical experiments

In this section, we support our theory with numerical experiments. First, we examine whether IRLS indeed exhibits two distinct convergence phases, a ‘‘global’’ one, as described in this paper, and a local one, as described in [3, 23], corresponding to different linear convergence rate factors. Second, we explore to which extent the dimension dependence in the convergence rates (9) and (10) indicated by Theorem 3.2 is necessary, or if we rather can expect a dimension-free linear convergence rate factor. All experiments are conducted on an iMac computer with a 4 GHz Quad-Core Intel Core i7 CPU, using MATLAB R2020b.

4.1 Local and global convergence phase

We first note that the local convergence result of [23, Theorem 6.1] depends on the locality condition $\xi(k) := \frac{\|x^k - x_*\|_1}{\min_{i \in S} |(x_*)_i|} < 1$, cf. (7). Under this condition (and an appropriate null space condition), Proposition 3.1 stated above implies that $\|x^{k+1} - x_*\|_1 \leq \mu \|x^k - x_*\|_1$ with an absolute constant $\mu < 1$ which, in particular, does *not* depend on the dimension N , m , and s . This corresponds to a locally linear rate for IRLS. A very similar condition to (7) is required by the comparable and more recent local convergence statement [3, Theorem III.6, inequality (III.14)] for the IRLS variant considered in [3].

However, a closer look at the locality condition (7) reveal that its *basin of attraction* is very restrictive: This condition means that the *support identification* problem underlying the sparse recovery *has already been solved*, as can be seen from the following proposition, whose proof we provide in the supplementary material.

Proposition 4.1. *Let $x^k, x_* \in \mathbb{R}^N$, let $S \subset [N]$ be the support set of x_* of size $|S| = s$. If (7) holds, i.e., if $\|x^k - x_*\|_1 < \min_{i \in S} |(x_*)_i|$, then the set $S_k \subset [N]$ of the s largest coordinates of x^k coincides with S .*

We now explore the behavior of the IRLS algorithm for ℓ_1 -minimization, Algorithm 1, and the sharpness of Proposition 4.1 in experiments that build on those of [23, Section 8.1]. For this purpose, for $N = 8000$, we sample independently a 200-sparse vector $x_* \in \mathbb{R}^N$ with random support $S \subset [N]$, $s = 200 = |S|$, chosen uniformly at random such that $(x_*)_S$ is chosen according the Haar measure




Figure 1: Instantaneous linear convergence rates of IRLS for ℓ_1 -minimization ($N = 8000$): Linear convergence factors $\mu_{\ell_1}(k) := \|x^k - x_*\|_1 / \|x^{k-1} - x_*\|_1$ (in blue), filled blue circle if $S_k = S$ with S_k of Proposition 4.1 (support identification), and error parameter $\zeta(k) := \|x^k - x_*\|_1 / \min_{i \in S} |(x_*)_i|$ (in red). Horizontal (red) line: Threshold $\zeta = 1$. Vertical (red) line: First iterate k with $\zeta(k) < 1$.

on the sphere of a 200-dimensional unit ℓ_2 -ball, and a measurement matrix $A \in \mathbb{R}^{m \times N}$ with i.i.d. Gaussian entries such that $A_{ij} \sim \mathcal{N}(0, 1/m)$, while setting $m = \lfloor 2s \log(N/s) \rfloor$. Such a matrix is known to fulfill with high probability the ℓ_1 -null space property of order s with constant $\rho_s < 1$ [32, Theorem 9.29].

In Figure 1, we track the decay of the ℓ_1 -error $\|x^k - x_*\|_1$ of the iterates x^k returned by Algorithm 1 via the values of $\zeta(k) := \|x^k - x_*\|_1 / \min_{i \in S} |(x_*)_i|$, depicted in red, and the behavior of the factor $\mu_{\ell_1}(k) := \|x^k - x_*\|_1 / \|x^{k-1} - x_*\|_1$, depicted in blue. We observe that the condition (7) for local convergence with the fast, dimension-less linear rate (8) is satisfied after $k = 33$ iterations, as indicated by the vertical dashed red line.

In the first few iterations, $\zeta(k)$ is larger than 1 by several orders of magnitudes, suggesting that the local convergence rate results of [3, 23] do *not* apply until the later stages of the simulation: In fact, we observe that the support S of x_* is already perfectly identified via the s largest coordinates of x^k as soon as $k \geq 18$. For iterations $18 \leq k \leq 50$, the linear rate $\mu_{\ell_1}(k)$ remains very stably around ≈ 0.7 , after which an accelerated linear rate can be observed.⁵ Before $k = 18$, for this example, the rate $\mu(k)$ hovers around 0.7 with slight variations. For all iterations k , $\mu(k)$ is smaller than 1, in line with the global linear convergence rate implied by Theorem 3.2.

Repeating a similar experiment for a larger ambient space dimension $N = 16000$ and a smaller measurement-to-sparsity ratio such that $m = \lfloor 1.75s \log(N/s) \rfloor$ results in a qualitatively similar situation, as seen in Figure 2(a): In Figure 2(a), we add also a plot of the linear convergence factor $\mu(k) := \frac{\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1}{\mathcal{J}_{\varepsilon_{k-1}}(x^{k-1}) - \|x_*\|_1}$ that tracks the behavior of the linear convergences in the smoothed ℓ_1 -norm objective \mathcal{J} , cf. (17). In addition to what have been observed in Figure 1, we see that $\mu(k)$ and $\mu_{\ell_1}(k)$ exhibit a very similar behavior for this example.

Hence, these experiments indicate that we can distinguish two phases. In the first, global phase linear convergence already sets in, but the instantaneous linear convergence rate has not yet stabilized. In the second one, when the support identification problem has been solved, the instantaneous linear convergence stabilizes.

Remark 4.2. *There are other methods in the literature, such as proximal algorithms, for which convergence results with a two phase behaviour were already established. For example, [47] showed that a forward-backward method applied to the Lasso problem exhibits local linear convergence, and that after a finite number of iterations, the region of fast convergence is reached. In particular, [47, Proposition 3.6(ii)] provides a bound on this number of iterations, which scales proportionally with $\|x_* - x^0\|_2^2$. On the other hand, (3.2) for IRLS provides a bound on the number of iterations until the fast linear convergence rate is reached that scales proportionally with $\log(\|x_* - x^0\|_2)$, but also proportionally with the dimension N . Moreover, most of these results require stronger assumptions than the NSP, such as the restricted isometry property or a restricted strong convexity/smoothness property.*

⁵The latter phenomenon cannot be observed for the IRLS algorithm of [23] as it uses a slightly different objective function than Algorithm 1.

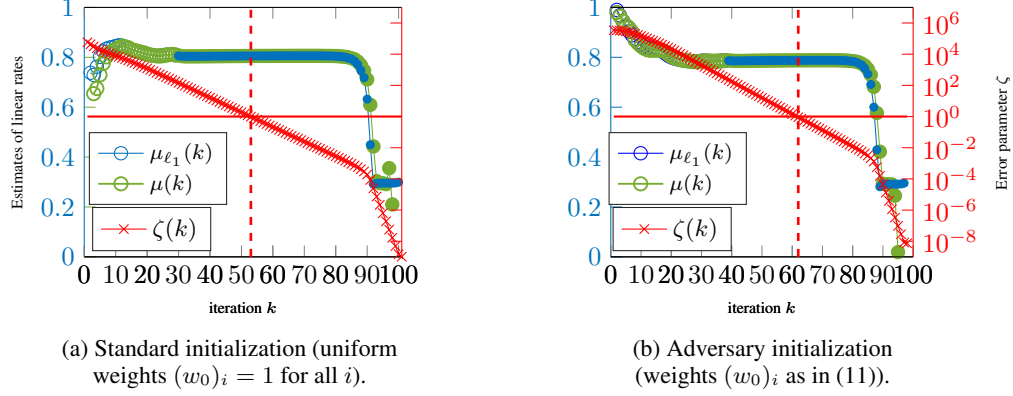


Figure 2: Instantaneous linear convergence rates of IRLS for ℓ_1 -minimization ($N = 16000$): Linear convergence factors $\mu_{\ell_1}(k) := \frac{\|x^k - x_*\|_1}{\|x^{k-1} - x_*\|_1}$ (in blue) and $\mu(k) := \frac{\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1}{\mathcal{J}_{\varepsilon_{k-1}}(x^{k-1}) - \|x_*\|_1}$ (in green), filled circles if $S_k = S$ (perfect support identification), and error parameter $\zeta(k) := \|x^k - x_*\|_1 / \min_{i \in S} |(x_*)_i|$ (in red), horizontal and vertical red lines as in Figure 1.

4.2 Global convergence rate and its dimension dependence

In this section, we explore to which extent the dependence on N in the convergence rates (10) and (14) is necessary or if we can rather expect a dimension-free linear convergence rate factor. To this end, we run a variation of IRLS that initializes the weight vector $w_0 \in \mathbb{R}^N$ not uniformly as in Algorithm 1, but based on an *adversary initialization*, here denoted by z^{adv} . More specifically, we first compute a minimizer

$$z^{\text{adv}} \in \arg \min_{z \in \mathbb{R}^{S^c}: A_{S^c} z = y} \|z\|_1$$

of the ℓ_1 -minimization problem restricted to the off-support coordinates of x_* indexed by S^c and set then $x^0 \in \mathbb{R}^N$ such that $x_{S^c}^0 := z^{\text{adv}}$ and $x_S^0 = 0$. Based on this *initialization* x^0 , we compute $\varepsilon_0 := \frac{\sigma_s(x^0)_{\ell_1}}{N}$ and set the first weight vector such that for all $i \in [N]$,

$$(w_0)_i := \frac{1}{\max(|x_i^0|, \varepsilon_0)}, \quad (11)$$

before proceeding with the IRLS steps (3), (4) and (5) until convergence.

We observe in Figure 2(b) that this initialization, which is *adversary* as it sets very large initial weights on the coordinates of S that correspond to the true support of x_* , eventually results in the same behavior of Algorithm 1 as for the standard initialization by uniform weights, identifying the true support at iteration $k = 39$ compared to $k = 30$. However, in the first few iterations, we see that the instantaneous linear convergence factor $\mu(k)$ is close to 1 with $\mu(1) = 0.980$, decreasing only slowly before stabilizing around 0.79 after around $k = 30$.

While this is just one example, this already indicates that in general, a linear rate such as (8), i.e., without dependence on the dimension N (which has been proven locally in [23, Theorem 6.1] and [3, Theorem III.6]) might not hold in general.

In our next experiment, we further investigate numerically the dimension dependence of the worst-case linear convergence factor $\mu(k) := \frac{\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1}{\mathcal{J}_{\varepsilon_{k-1}}(x^{k-1}) - \|x_*\|_1}$, which is upper bounded by the result of Theorem 3.2. We saw that in the experiment using the adversary initialization mentioned above and depicted in Figure 2(b), the maximal value was attained in the first iteration, i.e., for $\mu(1)$, as the effect of the adversary initialization is most eminent for $k = 1$.

We now run IRLS starting from the adversary initialization for different ambient dimensions $N = 125 \cdot 2^{\ell/2}$ for $\ell = 0, 1, \dots, 14$. For each of the values of N , we sample vectors $x_* \in \mathbb{R}^N$ of sparsity $s = 40$ from the same random model as above, and scale the number of i.i.d. Gaussian measurements with $m = \lfloor 2s \log(N/s) \rfloor$. We average the resulting values for $\mu(1)$ across 500 independent realizations of the experiment.




Figure 3: Comparison of $\frac{N}{100}$ and $\frac{1}{1-\mu(1)}$ (for which Proposition B.3 provides an upper bound of $\frac{\rho_1 N}{c}$) for different dimension parameters N , where $\mu(1) = \frac{\mathcal{J}_{\varepsilon_1}(x^1) - \|x_s\|_1}{\mathcal{J}_{\varepsilon_0}(x^0) - \|x_s\|_1}$ is the linear convergence factor, for IRLS initialized from adversary initialization.

In Figure 3, we see that dependence on N of linear convergence factor $\mu(1)$ that is observed for this experiment is quite well described by the upper bound (9) provided by our main result Theorem 3.2, as $\frac{1}{1-\mu(1)}$ scales almost linearly with N . As a footnote in Section 3 indicates, the constant ρ_1 of the null space property of order 1 scales with $\sqrt{\frac{\log N}{m}}$, and therefore a precise dependence on all the parameters such as m and s might be more complicated than what can be observed in this experiment. Nevertheless, we interpret Figure 3 as strong evidence that the linear convergence rate factor of Proposition B.3 is tight in its dependence on N , and that a dimension-less factor μ cannot be expected in general.

Conclusion

In this paper, we solved an open problem in the algorithmic theory for sparse recovery. In particular, we established a new variant of the IRLS algorithm for Basis Pursuit or ℓ_1 -minimization for which we show a global linear convergence rate under a suitable and sharp assumption, namely, the null-space property. Moreover, we have corroborated our theory with numerical experiments that, first, discussed the difference between the local and global convergence phase and, second, that elucidated the optimality of the dimension dependence of convergence rate given by our main theorem.

We think that the results in this paper give rise to a number of interesting research directions for follow-up work. While the numerical experiments in Section 4 substantiate the hypothesis that the dependence of the convergence rate on N and ρ_1 in our theory is not an artifact of our proof, we also observed in this section that for a *generic initialization* no such dependence can be observed. In view of this, it is interesting to investigate whether a dimension-independent global convergence rate is possible, for example via a *smoothed analysis* [22, 61]. Furthermore, there are currently no convergence rates available for IRLS optimizing a nuclear norm-type objective, which is of great interest for low-rank matrix recovery [31, 43, 51], and we expect that our analysis may be generalizable to this setting as well.

Finally, it was observed that sparse vectors can be recovered from even fewer measurements via the optimization of a non-convex ℓ_p -quasinorm (with $0 < p < 1$), and that IRLS exhibits excellent performance in this case [16, 23]. While a thorough understanding has remained elusive so far for this non-convex case, we consider our results as a first step towards a global convergence theory for IRLS for the optimization of ℓ_p -quasinorms or similar non-convex surrogate objectives.

Acknowledgments and Disclosure of Funding

The authors want to thank Massimo Fornasier for inspiring discussions and the anonymous reviewers for their detailed comments. C.K. is grateful for support by the NSF under the grant NSF-IIS-1837991. C.M.V. is supported by the German Federal Ministry of Education and Research (BMBF) in the context of the Collaborative Research Project SparseMRI3D+ and by the German Science Foundation (DFG) in the context of the project KR 4512/1-1. D.S. was supported by the Air Force Office of Scientific Research under award #FA9550-18-1-00788.

References

- [1] Deeksha Adil, Richard Peng, and Sushant Sachdeva, *Fast, provably convergent irls algorithm for p-norm linear regression*, Advances in Neural Information Processing Systems (NeurIPS), 2019, pp. 14189–14200.
- [2] Marc Allain, Jerome Idier, and Yves Goussard, *On global and local convergence of half-quadratic algorithms*, IEEE Trans. Image Process. **15** (2006), no. 5, 1130–1142.
- [3] Aleksandr Aravkin, James V Burke, and Daiwei He, *IRLS for Sparse Recovery Revisited: Examples of Failure and a Remedy*, arXiv preprint arXiv:1910.07095 (2019).
- [4] Hedy Attouch, Jerome Bolte, and Benar Fux Svaiter, *Convergence of descent methods for semi-algebraic and tame problems: Proximal algorithms, forward-backward splitting, and regularized gauss-seidel methods*, Math. Programming **137** (2013), 91–129.
- [5] Demba Ba, Behtash Babadi, Patrick L Purdon, and Emery N. Brown, *Convergence and stability of iteratively re-weighted least squares algorithms*, IEEE Transactions on Signal Processing **62** (2013), no. 1, 183–195.
- [6] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin, *A simple proof of the restricted isometry property for random matrices*, Constr. Approx. **28** (2008), 253–263.
- [7] Amir Beck and Shoham Sabach, *Weiszfeld’s method: Old and new results*, J. Optim. Theory Appl. **164** (2015), 1–40.
- [8] Jérôme Bolte, Aris Daniilidis, and Adrian Lewis, *The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems*, SIAM J. Optim. **17** (2007), no. 4, 1205–1223.
- [9] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Foundations and Trends® in Machine Learning **3** (2011), 1–122.
- [10] Charles S. Burrus, *Iterative reweighted least squares*, OpenStax CNX: Available online: <http://cnx.org/contents/92b90377-2b34-49e4-b26f-7fe572db78a1> (2012).
- [11] Jameson Cahill, Xuemei Chen, and Rongrong Wang, *The gap between the null space property and the restricted isometry property*, Linear Algebra Appl. **501** (2016), 363–375.
- [12] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao, *Stable signal recovery from incomplete and inaccurate measurements*, Comm. Pure Appl. Math. **59** (2006), no. 8, 1207–1223.
- [13] Emmanuel J. Candès and Terence Tao, *Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?*, IEEE Trans. Inf. Theory **52** (2006), no. 12, 5406–5425.
- [14] Emmanuel J. Candès, Michael B. Wakin, and Stephen P. Boyd, *Enhancing sparsity by reweighted ℓ_1 minimization*, Journal of Fourier analysis and applications **14** (2008), no. 5-6, 877–905.
- [15] Antonin Chambolle and Thomas Pock, *A first-order primal-dual algorithm for convex problems with applications to imaging*, Journal of mathematical imaging and vision **40** (2011), no. 1, 120–145.
- [16] Rick Chartrand and Wotao Yin, *Iteratively reweighted algorithms for compressive sensing*, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2008, pp. 3869–3872.
- [17] ———, *Nonconvex sparse regularization and splitting algorithms*, Splitting methods in communication, imaging, science, and engineering, 2016, pp. 237–249.
- [18] Scott Shaobing Chen, David L Donoho, and Michael A Saunders, *Atomic decomposition by basis pursuit*, SIAM Rev. **43** (2001), no. 1, 129–159.
- [19] Shaobing Scott Chen and David L. Donoho, *Basis pursuit*, Proc. of 1994 28th Asilomar Conference on Signals, Systems and Computers, 1994, pp. 41–44.
- [20] Jon F. Claerbout and Francis Muir, *Robust modeling with erratic data*, Geophysics **38** (1973), no. 5, 826–844.
- [21] Albert Cohen, Wolfgang Dahmen, and Ronald DeVore, *Compressed sensing and best k-term approximation*, J. Amer. Math. Soc. **22** (2009), 211–231.
- [22] Daniel Dadush and Sophie Huiberts, *A friendly smoothed analysis of the simplex method*, SIAM J. Comput. **49** (2018), no. 5, STOC18–449–STOC18–499.

- [23] Ingrid Daubechies, Ronald DeVore, Massimo Fornasier, and Sinan Güntürk, *Iteratively reweighted least squares minimization for sparse recovery*, Comm. Pure Appl. Math. **63** (2010), 1–38.
- [24] Geoffrey Davis, Stephane Mallat, and Marco Avellaneda, *Adaptive greedy approximations*, Constr. Approx. **13** (1997), no. 1, 57–98.
- [25] Lars Diening, Massimo Fornasier, Tomasi Roland, and Maximilian Wank, *A Relaxed Kačanov iteration for the p -poisson problem*, Numer. Math. **145** (2020), no. 2, 1–34.
- [26] David L. Donoho, *Compressed sensing*, IEEE Trans. Inf. Theor. **52** (2006), no. 4, 1289–1306.
- [27] David L Donoho and Benjamin F Logan, *Signal recovery and the large sieve*, SIAM J. Appl. Math. **52** (1992), no. 2, 577–591.
- [28] David L. Donoho and Yaakov Tsaig, *Fast solution of ℓ_1 -norm minimization problems when the solution may be sparse*, IEEE Trans. Inform. Theory **54** (2008), 4789–4812.
- [29] Alina Ene and Adrian Vladu, *Improved convergence for ℓ_1 and ℓ_∞ regression via iteratively reweighted least squares*, Proceedings of 2019 International Conference on Machine Learning (ICML’19), 2019, pp. 1794–1801.
- [30] Massimo Fornasier, Steffen Peter, Holger Rauhut, and Stephan Worm, *Conjugate gradient acceleration of iteratively re-weighted least squares methods*, Comput. Optim. Appl. **65** (2016), no. 1, 205–259.
- [31] Massimo Fornasier, Holger Rauhut, and Rachel Ward, *Low-rank matrix recovery via iteratively reweighted least squares minimization*, SIAM J. Optim. **21** (2011), 1614–1640.
- [32] Simon Foucart and Holger Rauhut, *A mathematical introduction to compressive sensing*, Applied and Numerical Harmonic Analysis, Springer New York, 2013.
- [33] Donald Geman and George Reynolds, *Constrained Restoration and the Recovery of Discontinuities*, IEEE Trans. Pattern Anal. Mach. Intell. **14** (1992), no. 03, 367–383.
- [34] Irina Gorodnitsky and Bhaskar Rao, *Sparse signal reconstruction from limited data using focuss: A re-weighted minimum norm algorithm*, IEEE Trans. Signal Process. **45** (1997), no. 3, 600–616.
- [35] Peter J. Green, *Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives*, Journal of the Royal Statistical Society: Series B (Methodological) **46** (1984), no. 2, 149–170.
- [36] Remi Gribonval and Morten Nielsen, *Sparse representation in the union of basis*, IEEE Trans. Inform. Theory **49** (2003), 3320–3325.
- [37] Trevor Hastie, Robert Tibshirani, and Martin Wainwright, *Statistical learning with sparsity: The lasso and generalizations*, Chapman and Hall/CRC, 2019.
- [38] Magnus Rudolph Hestenes and Eduard Stiefel, *Methods of conjugate gradients for solving linear systems*, Journal of research of the National Bureau of Standards **49** (1952), no. 1.
- [39] Paul W. Holland and Roy E. Welsch, *Robust regression using iteratively reweighted least-squares*, Communications in Statistics - Theory and Methods **6** (1977), 813–827.
- [40] Peter J. Huber, *Robust estimation of a location parameter*, Ann. Math. Statist. **35** (1964), no. 1, 73–101.
- [41] Jérôme Idier, *Convex half-quadratic criteria and interacting auxiliary variables for image restoration*, IEEE Trans. Image Process. **10** (2001), no. 7, 1001–1009.
- [42] Christian Kümmerle and Claudio Mayrink Verdun, *A scalable second order method for ill-conditioned matrix completion from few samples*, Proceedings of 2021 International Conference on Machine Learning (ICML’21), 2021.
- [43] Christian Kümmerle and Juliane Sigl, *Harmonic mean iteratively reweighted least squares for low-rank matrix recovery*, J. Mach. Learn. Res. **19** (2018), 1815–1863.
- [44] Ming-Jun Lai, Xu Yangyang, and Yin Wotao, *Improved iteratively reweighted least squares for unconstrained smoothed ℓ_q minimization*, SIAM J. Numer. Anal. **2** (2013), no. 51, 927–957.
- [45] Gilad Lerman and Tyler Maunu, *Fast, robust and non-convex subspace recovery*, Information and Inference: A Journal of the IMA **7** (2018), no. 2, 277–336.

- [46] Xudong. Li, Defeng Sun, and Kim-Chuan Toh, *A highly efficient semismooth newton augmented lagrangian method for solving lasso problems*, SIAM J. Optim. **28** (2018), no. 1, 433–458.
- [47] Jingwei Liang, Jalal Fadili, and Gabriel Peyré, *Activity identification and local linear convergence of forward-backward-type methods*, SIAM Journal on Optimization **27** (2017), no. 1, 408–437.
- [48] Nicolai Meinshausen, Peter Bühlmann, et al., *High-dimensional graphs and variable selection with the lasso*, Annals of statistics **34** (2006), no. 3, 1436–1462.
- [49] Gérard Meurant, *The lanczos and conjugate gradient algorithms: From theory to finite precision computations*, Society for Industrial and Applied Mathematics, 2006.
- [50] Gregory P Meyer, *An alternative probabilistic interpretation of the huber loss*, Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2021, pp. 5261–5269.
- [51] Karthik Mohan and Maryam Fazel, *Iterative reweighted algorithms for matrix rank minimization*, J. Mach. Learn. Res. **13** (2012), 3441–3473.
- [52] Cesare Molinari, Mathurin Massias, Lorenzo Rosasco, and Silvia Villa, *Iterative regularization for convex regularizers*, International conference on artificial intelligence and statistics, 2021, pp. 1684–1692.
- [53] Bhaskar Mukhoty, Govind Gopakumar, Prateek Jain, and Purushottam Kar, *Globally-convergent iteratively reweighted least squares for robust regression problems*, Proceedings of Machine Learning Research, 2019, pp. 313–322.
- [54] Balas K. Natarajan, *Sparse approximate solutions to linear systems*, SIAM J. Comput. **24** (1996), 227–234.
- [55] Mila Nikolova and Michael K. Ng, *Analysis of half-quadratic minimization methods for signal and image recovery*, SIAM J. Sci. Comput. **27** (2005), no. 3, 937–966.
- [56] Jorge Nocedal and Stephen Wright, *Numerical optimization*, Springer Science & Business Media, 2006.
- [57] Peter Ochs, Alexey Dosovitskiy, Thomas Brox, and Thomas Pock, *On iteratively reweighted algorithms for nonsmooth nonconvex optimization in computer vision*, SIAM J. Imaging Sci. **8** (2015), no. 1, 331–372.
- [58] Clarice Poon and Gabriel Peyré, *Smooth Bilevel Programming for Sparse Regularization*, Proc. NeurIPS’21, 2021.
- [59] Bhaskar D. Rao and Kenneth Kreutz-Delgado, *An affine scaling methodology for best basis selection*, IEEE Trans. Signal Process. **47** (1999), no. 1, 187–200.
- [60] Mark Rudelson and Roman Vershynin, *On sparse reconstruction from fourier and gaussian measurements*, Comm. Pure Appl. Math. **61** (2008), no. 8, 1025–1045.
- [61] Daniel A. Spielman and Shang-Hua Teng, *Smoothed analysis of termination of linear programming algorithms*, Math. Program., Ser. B **97** (2003), 375–404.
- [62] Damian Straszak and Nisheeth K Vishnoi, *Iteratively reweighted least squares and slime mold dynamics: connection and convergence*, Mathematical Programming (2021), 1–33.
- [63] Ying Sun, Prabhu Babu, and Daniel P. Palomar, *Majorization-minimization algorithms in signal processing, communications, and machine learning*, IEEE Trans. Signal Process. **65** (2017), no. 3, 794–816.
- [64] Howard L. Taylor, Stephen C. Banks, and John F. McCoy, *Deconvolution with the ℓ_1 norm*, Geophysics **44** (1979), no. 1, 39–52.
- [65] Robert Tibshirani, *Regression shrinkage and selection via the lasso*, Journal of the Royal Statistical Society: Series B (Methodological) **58** (1996), no. 1, 267–288.
- [66] Andreas M Tillmann, *Equivalence of linear programming and basis pursuit*, PAMM **15** (2015), no. 1, 735–738.
- [67] Sergey Voronin, *Regularization of linear systems with sparsity constraints with applications to large scale inverse problems*, Ph.D. Thesis, Princeton University, 2012.
- [68] Sergey Voronin and Ingrid Daubechies, *An iteratively reweighted least squares algorithm for sparse regularization*, Functional analysis, harmonic analysis, and image processing: A collection of papers in honor of björn jawerth, 2017, pp. 391–441.
- [69] Endre Weiszfeld, *Sur le point pour lequel la somme des distances de n points donnés est minimum*, Tohoku Mathematical Journal, First Series **43** (1937), 355–386.

- [70] Max A. Woodbury, *Inverting modified matrices*, Memorandum report **42** (1950), no. 106, 336.
- [71] Wotao Yin, Stanley Osher, Donald Goldfarb, and Jerome Darbon, *Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing*, SIAM J. Imaging Sci **1** (2008), no. 1, 143–168.
- [72] Zaiwei Zhang, Xiangru Huang, Qixing Huang, Xiao Zhang, and Yuan Li, *Joint learning of neural networks via iterative reweighted least squares.*, Cvpr workshops, 2019, pp. 18–26.

Supplementary material of *Iteratively Reweighted Least Squares for Basis Pursuit with Global Linear Convergence Rate*

In this supplement to the paper, we present in Section A our second main result, Theorem A.1, which generalizes the global linear convergence of Algorithm 1 for sparse vectors, which we presented Theorem 3.2, to approximately sparse vectors. In Section B we provide the proof of our theoretical results. In Section C, we discuss how to solve the (weighted) least squares step in Algorithm 1 in a computationally efficient manner. Finally, Section D contains, for completeness, a proof of two technical results, namely Lemma 2.1 and Proposition 4.1.

Contents

A Linear Convergence of IRLS for approximately sparse vectors	15
B Proofs of the main results	16
C Computationally Efficient Weighted Least Squares Updates	22
D Proofs of technical lemmas	25
D.1 Proof of Lemma 2.1	25
D.2 Proof of Proposition 4.1	26

A Linear Convergence of IRLS for approximately sparse vectors

We now generalize Theorem 3.2 to the scenario where the ground truth x_* is only approximately sparse. By that, we mean that the vector x_* can be well-approximated by an s -sparse vector in the sense that the ℓ_1 -error of the best s -term approximation $\sigma_s(x_*)_{\ell_1} = \inf\{\|x_* - z\|_1 : z \in \mathbb{R}^N \text{ is } s\text{-sparse}\}$ is small, which is a commonly used quantity to measure the model misfit to a sparse vector [32, Section 2.1]. If x_* is approximately sparse in this sense, we can only hope to *approximately* recover x_* by the ℓ_1 -minimization program (P_1) . Indeed, [23, Theorem 5.3(iv)] showed that under a suitable null space property, IRLS for (P_1) finds a vector x , such that $\|x - x_*\|_1$ is at most a constant multiple of the optimal best s -term approximation error $\sigma_s(x_*)_{\ell_1}$. However, as for exactly sparse vectors x_* , only a local, but no global convergence rate was provided in previous literature [23, Theorem 6.4].

The following result shows that in fact, we also can obtain global linear convergence of Algorithm 1 in this case. More precisely, Theorem A.1 implies that $\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1$ decays exponentially fast until a certain accuracy is reached, which is $\sigma_s(x_*)_{\ell_1}$ up to a constant multiple.

Theorem A.1. *Consider the problem of recovering an unknown vector $x_* \in \mathbb{R}^N$ from known measurements of the form $y = Ax_*$. Assume that the measurement matrix $A \in \mathbb{R}^{m \times N}$ fulfills the ℓ_1 -NSP of order s with constant $\rho_s < 1/8$. Let the IRLS iterates $\{x^k\}_k$ and $\{\varepsilon_k\}_k$ be defined by (3) and (4) with initialization x^0 . Then the following three statements hold.*

1. For $k \leq \hat{k} := \min\left\{k \in \mathbb{N} : \sigma_s(x_*)_{\ell_1} > \frac{2}{9} \|(x_* - x^k)_{S^c}\|_1\right\}$ it holds that

$$\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1 \leq \left(1 - \frac{c}{\rho_1 N}\right)^k (\mathcal{J}_{\varepsilon_0}(x^0) - \|x_*\|_1), \quad (12)$$

where S denotes the support of the s largest entries of x_* .

2. For all $1 \leq k \leq \hat{k}$ it holds that

$$\|x^k - x_*\|_1 \leq 6 \left(1 - \frac{c}{\rho_1 N}\right)^k \|x^0 - x_*\|_1 + 10\sigma_s(x_*)_{\ell_1}. \quad (13)$$

3. Moreover, for all integers $k \gtrsim \rho_1 N \log \left(\frac{\|x^0 - x_*\|_1}{\sigma_s(x_*)_{\ell_1}} \right)$ we have that

$$\|x^k - x_*\|_1 \leq 20\sigma_s(x_*)_{\ell_1}. \quad (14)$$

Here $c = 1/3072$ and $\rho_1 < 1/8$ denotes the constant for the ℓ_1 -NSP of order 1.

Remark A.2. Applying Theorem A.1 to the special case $\sigma_s(x_*)_{\ell_1} = 0$, we observe that inequality (10) yields a seemingly sharper result than inequality (13) in Theorem 3.2, which may seem somewhat counterintuitive. However, note that in Theorem 3.2 we require $\rho_s < 1/2$, whereas in Theorem A.1 we have the stronger assumption $\rho_s < 1/8$. Indeed, a closer inspection of the proofs reveals that both the factors 3 and 6 in the inequalities (10) and (13) can be replaced by the factor $\frac{3(1+\rho_s)}{1-\rho_s}$, reconciling those two results.

B Proofs of the main results

In this section, we prove the main results of this paper, Theorem 3.2, and its approximately sparse counterpart, Theorem A.1. To this end, we first state and prove the following technical lemma, which gives an upper and lower bound for $\mathcal{J}_\varepsilon(x) - \|x_*\|_1$, which is the quantity for which we are going to show linear convergence.

Lemma B.1. Let $x_*, x \in \mathbb{R}^N$. Assume that A fulfills the ℓ_1 -NSP of order s with constant $\rho_s < 1$. Furthermore, suppose $Ax_* = Ax$ and that $\varepsilon \leq \frac{1}{N}\sigma_s(x)_{\ell_1}$. Then it holds that

$$\frac{1-\rho_s}{1+\rho_s} \|x - x_*\|_1 - 2\sigma_s(x_*)_{\ell_1} \leq \mathcal{J}_\varepsilon(x) - \|x_*\|_1 \leq 3\sigma_s(x)_{\ell_1}. \quad (15)$$

In order to prove Lemma B.1 we need the following technical lemma.

Lemma B.2. [23, Lemma 4.3] Assume that the matrix $A \in \mathbb{R}^{m \times N}$ fulfills the ℓ_1 -NSP for some s and $\rho_s < 1$. Then for all $x, x_* \in \mathbb{R}^N$ such that $Ax = Ax_*$ it holds that

$$\|x - x_*\|_{\ell_1} \leq \frac{1+\rho_s}{1-\rho_s} (\|x_*\|_1 - \|x\|_1 + 2\sigma_s(x)_{\ell_1}).$$

Proof of Lemma B.1. We observe that $\mathcal{J}_\varepsilon(x) \geq \|x\|_1$ for each $x \in \mathbb{R}^N$, which follows directly from the definition of $\mathcal{J}_\varepsilon(x)$, see Equation (1). Now, fix $x_*, x \in \mathbb{R}^N$, and let S be the set which contains the s largest entries of x in absolute value. Hence, we obtain that

$$\begin{aligned} \mathcal{J}_\varepsilon(x) - \|x_*\|_1 &\geq \|x\|_1 - \|x_*\|_1 \\ &= \|x_{S^c}\|_1 + \|x_S\|_1 - \|x_*\|_1 \\ &\geq \|x_{S^c}\|_1 - \|(x - x_*)_S\|_1 - \|(x_*)_{S^c}\|_1 \\ &\geq \|(x - x_*)_{S^c}\|_1 - \|(x - x_*)_S\|_1 - 2\|(x_*)_{S^c}\|_1, \end{aligned}$$

where in each of the last two inequalities we have applied the reverse triangle inequality. Since $x - x_*$ is contained in the null space of A , it follows from the nullspace property that $\|(x - x_*)_S\|_1 \leq \rho_s \|(x - x_*)_{S^c}\|_1$. Hence, we have shown that

$$\mathcal{J}_\varepsilon(x) - \|x_*\|_1 \geq (1 - \rho_s) \|(x - x_*)_{S^c}\|_1 - 2\|(x_*)_{S^c}\|_1.$$

Since it follows from the null space property that $\|(x - x_*)_{S^c}\|_1 \geq \frac{\|x - x_*\|_1}{1 + \rho_s}$, this shows the first inequality in (15).

Next, we are going to prove the reverse inequality in (15). For that, set $I := \{i \in [N] : |x_i| > \varepsilon\}$. Then we observe that

$$\begin{aligned} \mathcal{J}_\varepsilon(x) - \|x_*\|_1 &= \|x_I\|_1 + \frac{1}{2} \sum_{i \in I^c} \left(\frac{x_i^2}{\varepsilon} + \varepsilon \right) - \|x_*\|_1 \\ &\leq \|x_I\|_1 + |I^c| \varepsilon - \|x_*\|_1 \\ &\leq \|x_I\|_1 + \sigma_s(x)_{\ell_1} - \|x_*\|_1 \\ &\leq \|x\|_1 + \sigma_s(x)_{\ell_1} - \|x_*\|_1. \end{aligned} \quad (16)$$

In the third line we used the assumption $\varepsilon \leq \frac{1}{N}\sigma_s(x)_{\ell_1}$. In order to proceed, we first derive an appropriate upper bound for $\|x\|_1 - \|x_*\|_1$. For that, we note

$$\begin{aligned} \left(\frac{1-\rho_s}{1+\rho_s} + 1\right) (\|x\|_1 - \|x_*\|_1) &\leq \frac{1-\rho_s}{1+\rho_s} \|x - x_*\|_1 - (\|x_*\|_1 - \|x\|_1) \\ &\leq (\|x_*\|_1 - \|x\|_1 + 2\sigma_s(x)_{\ell_1}) - (\|x_*\|_1 - \|x\|_1) \\ &\leq 2\sigma_s(x)_{\ell_1}, \end{aligned}$$

where in the second line we have used Lemma B.2. This shows that $\|x\|_1 - \|x_*\|_1 \leq \frac{2\sigma_s(x)_{\ell_1}}{1+\frac{1-\rho_s}{1+\rho_s}}$.

Combining this with (16), we obtain

$$\mathcal{J}_\varepsilon(x) - \|x_*\|_1 \leq 3\sigma_s(x)_{\ell_1},$$

which finishes the proof of inequality (15). \square

The next key proposition states that the quantity $\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1$ decays linearly under appropriate conditions.

Proposition B.3. *Let $x_* \in \mathbb{R}^N$ be an approximately s -sparse vector with support S . Let $A \in \mathbb{R}^{m \times N}$ and $y = Ax_*$. Assume that A fulfills the ℓ_1 -NSP of order s with constant $\rho_s < 3/4$, if $\sigma_s(x_*)_{\ell_1} = 0$, and $\rho_s < 1/4$ otherwise. Denote by ρ_1 the NSP constant of order 1.*

Let the IRLS iterates $\{x^k\}_k$ and $\{\varepsilon_k\}_k$ be defined by (3) and (4) with initialization x^0 . Then, for all $k \in \mathbb{N}$, such that $\|(x_)_{S^c}\|_1 \leq \frac{2}{9}\|(x_*)_{S^c} - x_{S^c}^\ell\|_1$ for all $\ell < k$, the following holds*

$$\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1 \leq \left(1 - \frac{c_{\rho_s}}{\rho_1 N}\right)^k (\mathcal{J}_{\varepsilon_0}(x^0) - \|x_*\|_1). \quad (17)$$

where the constant c_{ρ_s} is defined by

$$c_{\rho_s} := \begin{cases} \frac{(3/4 - \rho_s)^2}{48} & \text{if } \sigma_s(x_*)_{\ell_1} = 0 \\ \frac{(1/4 - \rho_s)^2}{48} & \text{else} \end{cases}$$

Before proving this statement, let us describe the main ideas of our proof. Recall that x_* has minimal $\|\cdot\|_1$ -norm among all vectors x , which satisfy the constraint $Ax = y$. Hence, setting $v^k = x_* - x^k$ due to convexity of the ℓ_1 -norm we that $\|x^k + tv^k\|_1 < \|x^k\|_1$ for all $0 < t < 1$. Since that the quadratic functional $Q(\cdot, x^k)$ approximates the objective function \mathcal{J}_ε , which is a surrogate for the ℓ_1 -norm, in a neighborhood of the current iterate x^k , we also expect that for $t > 0$ sufficiently small we have that $Q(x^k + tv^k, x^k) < Q(x^k, x^k)$. In order to show that the decrease is sufficiently large, we also need to show that t can be chosen large enough. This will guarantee a sufficient decrease of $\mathcal{J}_{\varepsilon_k}(x^k)$ in each iteration.

Proof of Proposition B.3. In order to show inequality (17) we will prove by induction that for each k , such that $\|(x_*)_{S^c}\|_1 \leq \frac{2}{9}\|(x_*)_{S^c} - x_{S^c}^\ell\|_1$ for all $\ell \leq k$, it holds that

$$\mathcal{J}_{\varepsilon_{k+1}}(x^{k+1}) - \|x_*\|_1 \leq \left(1 - \frac{c_{\rho_s}}{\rho_1 N}\right) (\mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1).$$

Now choose such a $k \geq 1$ and assume that the statement has been shown for all $k' < k$. Set $v^k = x_* - x^k$. For $t \in \mathbb{R}$, we have, by optimality of x^{k+1} in (3), that

$$\mathcal{J}_{\varepsilon_{k+1}}(x^{k+1}) \leq Q_{\varepsilon_k}(x^{k+1}, x^k) \leq Q_{\varepsilon_k}(x^k + tv^k, x^k). \quad (18)$$

Moreover, by the definition of the quadratic objective $Q_{\varepsilon_k}(\cdot, x^k)$ (see (2)), it holds that

$$Q_{\varepsilon_k}(x^k + tv^k, x^k) - \mathcal{J}_{\varepsilon_k}(x^k) = t \langle \nabla \mathcal{J}_{\varepsilon_k}(x^k), v^k \rangle + \frac{t^2}{2} \langle v^k, \text{diag}(w_{\varepsilon_k}(x^k))v^k \rangle. \quad (19)$$

Our goal is to show that by picking t large enough, we can make $Q_{\varepsilon_k}(x^k + tv^k, x^k) - \mathcal{J}_{\varepsilon_k}(x^k) < 0$ sufficiently small. For that, we now control $\langle \nabla \mathcal{J}_{\varepsilon_k}(x^k), v^k \rangle$ and $\langle v^k, \text{diag}(w_{\varepsilon_k}(x^k))v^k \rangle$ separately.

Part I: Bounding the linear term $\langle \nabla \mathcal{J}_{\varepsilon_k}(x^k), v^k \rangle$:

Let $I := \{i \in [N] : |x_i^k| > \varepsilon_k\}$ and denote by S the set which contains the s largest entries of x_* in absolute value. In the case that x_* is sparse, S is given by the support of x_* , i.e. $S = \text{supp}(x_*)$. Consider

$$\langle \nabla \mathcal{J}_{\varepsilon_k}(x^k), v^k \rangle = \sum_{i=1}^N \frac{x_i^k}{\max(|x_i^k|, \varepsilon_k)} v_i^k = \sum_{i \in S} \frac{x_i^k}{\max(|x_i^k|, \varepsilon_k)} v_i^k + \sum_{i \in S^c} \frac{x_i^k}{\max(|x_i^k|, \varepsilon_k)} v_i^k.$$

The first summand can be bounded by

$$\begin{aligned} \sum_{i \in S} \frac{x_i^k}{\max(|x_i^k|, \varepsilon_k)} v_i^k &= \sum_{i \in S \cap I} \text{sgn}(x_i^k) v_i^k + \sum_{i \in S \cap I^c} \frac{x_i^k}{\varepsilon_k} v_i^k \\ &\leq \|v_{S \cap I}^k\|_1 + \|v_{S \cap I^c}^k\|_1 \\ &= \|v_S^k\|_1 \\ &\leq \rho_s \|v_{S^c}^k\|_1. \end{aligned}$$

Where the last inequality comes from the NSP for v^k which, in turns, comes from the fact that $v^k = x_* - x^k \in \ker(A)$, since $Ax_* = Ax^k$. For the second summand we have that

$$\begin{aligned} &\sum_{i \in S^c} \frac{x_i^k}{\max(|x_i^k|, \varepsilon_k)} v_i^k \\ &= \sum_{i \in S^c \cap I} \text{sgn}(x_i^k) v_i^k + \sum_{i \in S^c \cap I^c} \frac{x_i^k v_i^k}{\varepsilon_k} \\ &= \sum_{i \in S^c \cap I} \text{sgn}(x_i^k) (x_*)_i - \sum_{i \in S^c \cap I} \text{sgn}(x_i^k) x_i^k + \sum_{i \in S^c \cap I^c} \frac{x_i^k (x_*)_i}{\varepsilon_k} - \sum_{i \in S^c \cap I^c} \frac{(x_i^k)^2}{\varepsilon_k} \\ &\leq \|(x_*)_{S^c \cap I}\|_1 - \|x_{S^c \cap I}^k\|_1 + \|(x_*)_{S^c \cap I^c}\|_1 - \frac{\|x_{S^c \cap I^c}^k\|_2^2}{\varepsilon_k} \\ &= -\|x_{S^c \cap I}^k\|_1 + \|(x_*)_{S^c}\|_1 - \frac{\|x_{S^c \cap I^c}^k\|_2^2}{\varepsilon_k} \\ &= \|(x_*)_{S^c}\|_1 - \|x_{S^c}^k\|_1 + \|x_{S^c \cap I^c}^k\|_1 - \frac{\|x_{S^c \cap I^c}^k\|_2^2}{\varepsilon_k} \\ &\leq 2\|(x_*)_{S^c}\|_1 - \|v_{S^c}^k\|_1 + \|x_{S^c \cap I^c}^k\|_1 - \frac{\|x_{S^c \cap I^c}^k\|_2^2}{\varepsilon_k}. \end{aligned}$$

In order to proceed, we note that from the elementary inequality $ab \leq \frac{1}{2}(a^2 + b^2)$ and from $\|x_{S^c \cap I^c}^k\|_1 \leq \sqrt{N} \|x_{S^c \cap I^c}^k\|_2$, it follows that

$$\|x_{S^c \cap I^c}^k\|_1 \leq \frac{1}{2} \left(\frac{\varepsilon_k \|x_{S^c \cap I^c}^k\|_1^2}{2 \|x_{S^c \cap I^c}^k\|_2^2} + 2 \frac{\|x_{S^c \cap I^c}^k\|_2^2}{\varepsilon_k} \right) \leq \frac{\varepsilon_k N}{4} + \frac{\|x_{S^c \cap I^c}^k\|_2^2}{\varepsilon_k}.$$

Hence, using that $\varepsilon_k \leq \sigma_s(x^k)_{\ell_1}/N$, we have shown that

$$\begin{aligned} \sum_{i \in S^c} \frac{x_i^k}{\max(|x_i^k|, \varepsilon_k)} v_i^k &\leq 2\|(x_*)_{S^c}\|_1 - \|v_{S^c}^k\|_1 + \frac{\varepsilon_k N}{4} \\ &\leq 2\|(x_*)_{S^c}\|_1 - \|v_{S^c}^k\|_1 + \frac{\sigma_s(x^k)_{\ell_1}}{4} \\ &\leq 2\|(x_*)_{S^c}\|_1 - \|v_{S^c}^k\|_1 + \frac{\|x_{S^c}^k\|_1}{4} \\ &\leq 2\|(x_*)_{S^c}\|_1 - \|v_{S^c}^k\|_1 + \frac{\|v_{S^c}^k\|_1}{4} + \frac{\|(x_*)_{S^c}\|_1}{4} \\ &= \frac{9}{4} \|(x_*)_{S^c}\|_1 - \frac{3}{4} \|v_{S^c}^k\|_1, \end{aligned}$$

where we used the triangular inequality for the vector $v^k = x^k - x_*$ on the set S^c and the fact that $\sigma_s(x^k)_{\ell_1} \leq \|x_{S^c}^k\|_1$. Hence, by adding up terms we obtain that

$$\langle \nabla \mathcal{J}_{\varepsilon_k}(x^k), v^k \rangle \leq \frac{9}{4} \| (x_*)_{S^c} \|_1 - \left(\frac{3}{4} - \rho_s \right) \|v_{S^c}^k\|_1 \leq -(\beta - \rho_s) \|v_{S^c}^k\|_1.$$

Here, we have set $\beta = 3/4$ in the case that $\sigma_s(x_*)_{\ell_1} = 0$ and $\beta = 1/4$ else. Moreover, we used the assumption $\| (x_*)_{S^c} \|_1 \leq \frac{2}{9} \|v_{S^c}^k\|_1$.

Part II: Bounding the quadratic term $\langle v^k, \text{diag}(w_{\varepsilon_k}(x^k))v^k \rangle$

In order bound the quadratic term in (19) we first decompose it into two parts

$$\langle v^k, \text{diag}(w_{\varepsilon_k}(x^k))v^k \rangle = \sum_{i=1}^N \frac{(v_i^k)^2}{\max(|x_i^k|, \varepsilon_k)} = \sum_{i \in S} \frac{(v_i^k)^2}{\max(|x_i^k|, \varepsilon_k)} + \sum_{i \in S^c} \frac{(v_i^k)^2}{\max(|x_i^k|, \varepsilon_k)}. \quad (20)$$

For the first summand, we note that

$$\sum_{i \in S} \frac{(v_i^k)^2}{\max(|x_i^k|, \varepsilon_k)} \leq \frac{\|v_{S^c}^k\|_1 \|v_S^k\|_\infty}{\varepsilon_k} \leq \rho_s \frac{\|v_{S^c}^k\|_1 \|v^k\|_\infty}{\varepsilon_k} \leq \frac{\|v_{S^c}^k\|_1 \|v^k\|_\infty}{\varepsilon_k}. \quad (21)$$

For the second summand, it holds that

$$\sum_{i \in S^c} \frac{(v_i^k)^2}{\max(|x_i^k|, \varepsilon_k)} \leq \frac{\|v_{S^c}^k\|_\infty \|v_{S^c}^k\|_1}{\varepsilon_k} \leq \frac{\|v_{S^c}^k\|_1 \|v^k\|_\infty}{\varepsilon_k}, \quad (22)$$

Hence, by adding (21) and (22) up, it follows that

$$\langle v^k, \text{diag}(w_{\varepsilon_k}(x^k))v^k \rangle \leq 2 \frac{\|v_{S^c}^k\|_1 \|v^k\|_\infty}{\varepsilon_k}.$$

Next, we note that

$$\|v^k\|_\infty \leq \rho_1 \|v^k\|_1 \leq \rho_1 (1 + \rho_s) \|v_{S^c}^k\|_1 \leq 2\rho_1 \|v_{S^c}^k\|_1.$$

Hence, we have shown that

$$\langle v^k, \text{diag}(w_{\varepsilon_k}(x^k))v^k \rangle \leq 4\rho_1 \frac{\|v_{S^c}^k\|_1^2}{\varepsilon_k}.$$

Part III: Combining the bounds to obtain decrease in k -th step:

Inserting the bounds of Part I and Part II into (19) we obtain

$$Q_{\varepsilon_k}(x^k + tv^k, x^k) - \mathcal{J}_{\varepsilon_k}(x^k) \leq -tb + t^2a =: h(t), \quad (23)$$

where the function $h : \mathbb{R} \rightarrow \mathbb{R}$ is a quadratic polynomial with coefficients $b = (\beta - \rho_s) \|v_{S^c}^k\|_1$ and $a = 4\rho_1 \frac{\|v_{S^c}^k\|_1^2}{\varepsilon_k}$. We observe that the minimizer of h is given by $t = \frac{b}{2a}$. Inserting this into h , we obtain that

$$h\left(\frac{b}{2a}\right) = -\frac{b^2}{4a} = -\frac{(\beta - \rho_s)^2 \|v_{S^c}^k\|_1^2 \varepsilon_k}{16\rho_1 \|v_{S^c}^k\|_1^2} = -\frac{(\beta - \rho_s)^2}{16\rho_1} \varepsilon_k. \quad (24)$$

Combining this with (6), we obtain, for $t = \frac{b}{2a}$,

$$\begin{aligned} \mathcal{J}_{\varepsilon_{k+1}}(x^{k+1}) - \mathcal{J}_{\varepsilon_k}(x^k) &\leq Q_{\varepsilon_k}(x^{k+1}, x^k) - \mathcal{J}_{\varepsilon_k}(x^k) \\ &\leq Q_{\varepsilon_k}(x^k + tv^k, x^k) - \mathcal{J}_{\varepsilon_k}(x^k) \leq -\frac{(\beta - \rho_s)^2}{16\rho_1} \varepsilon_k. \end{aligned}$$

Hence, by rearranging terms it follows that

$$\mathcal{J}_{\varepsilon_{k+1}}(x^{k+1}) - \|x_*\|_1 \leq \mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1 - \frac{(\beta - \rho_s)^2}{16\rho_1} \varepsilon_k. \quad (25)$$

In order to proceed, we need to bound ε_k from below. For that, we note that

$$\varepsilon_k = \min \left(\varepsilon_{k-1}, \frac{\sigma_s(x^k)_{\ell_1}}{N} \right) = \frac{\sigma_s(x^\ell)_{\ell_1}}{N}$$

for some $\ell \leq k$. By Lemma B.1, we have the following inequality chain

$$N\varepsilon_k = \sigma_s(x^\ell)_{\ell_1} \geq \frac{1}{3} (\mathcal{J}_{\varepsilon^\ell}(x^\ell) - \|x_*\|_1) \geq \frac{1}{3} (\mathcal{J}_{\varepsilon^k}(x^k) - \|x_*\|_1),$$

where in the second inequality we have used that, by induction, $\mathcal{J}_{\varepsilon^k}(x^k) \leq \mathcal{J}_{\varepsilon^\ell}(x^\ell)$. Plugging this into (25) leads to

$$\mathcal{J}_{\varepsilon_{k+1}}(x^{k+1}) - \|x_*\|_1 \leq \left(1 - \frac{(\beta - \rho_s)^2}{48\rho_1 N} \right) (\mathcal{J}_{\varepsilon^k}(x^k) - \|x_*\|_1).$$

This finishes the induction step and concludes the proof of Proposition B.3. \square

From Proposition B.3 we can deduce Theorem 3.2, the first main result of this manuscript.

Proof of Theorem 3.2. Recall that by Proposition B.3 we have for all $k \in \mathbb{N}$ that

$$\mathcal{J}_{\varepsilon^k}(x^k) - \|x_*\|_1 \leq \left(1 - \frac{c_{\rho_s}}{\rho_1 N} \right)^k (\mathcal{J}_{\varepsilon_0}(x^0) - \|x_*\|_1)$$

with a constant $c_{\rho_s} = \frac{(3/4 - \rho_s)^2}{48}$ and where S denotes the set, which contains the s largest entries of x_* in absolute value. By our assumption $\rho_s < 1/2$ it follows that $c_{\rho_s} \geq 1/768$, which implies that inequality (9) holds.

By Lemma B.1 we have that

$$\mathcal{J}_{\varepsilon_0}(x^0) - \|x_*\|_1 \leq 3\sigma_s(x^0)_{\ell_1} \leq 3\|x^0 - x_*\|_1.$$

Next, we note that, again by Lemma B.1, it holds that

$$\frac{1 - \rho_s}{1 + \rho_s} \|x - x_*\|_1 \leq \mathcal{J}_{\varepsilon^k}(x^k) - \|x_*\|_1.$$

Combining the three inequalities in this proof together with the assumption $\rho_s < 1/2$ yields inequality (10), which finishes the proof. \square

Next, we are going to prove the second main result in this manuscript, Theorem A.1, which deals with the approximately sparse case.

Proof of Theorem A.1. Recall that

$$\hat{k} := \min \left\{ k \in \mathbb{N} : \sigma_s(x_*)_{\ell_1} > \frac{2}{9} \|(x_*)_{S^c} - x_{S^c}^k\|_1 \right\}.$$

Moreover, we note that by Proposition B.3 we have for $k \leq \hat{k}$

$$\mathcal{J}_{\varepsilon^k}(x^k) - \|x_*\|_1 \leq \left(1 - \frac{c_{\rho_s}}{\rho_1 N} \right)^k (\mathcal{J}_{\varepsilon_0}(x^0) - \|x_*\|_1) \quad (26)$$

with a constant $c_{\rho_s} = \frac{(1/4 - \rho_s)^2}{48}$. Hence, by our assumption $\rho_s < 1/8$ we obtain $c_{\rho_s} \geq 1/3072$ and inequality (12) follows, which proves the first statement. In order to prove the second statement, let \tilde{k}

and k be natural numbers, such that $\tilde{k} \leq \hat{k}$ and $k \geq \tilde{k}$ holds. Then we obtain that

$$\begin{aligned}
\frac{1 - \rho_s}{1 + \rho_s} \|x^k - x_*\|_1 - 2\sigma_s(x_*)_{\ell_1} &\leq \mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1 \\
&\leq \mathcal{J}_{\varepsilon_{\tilde{k}}}(x^{\tilde{k}}) - \|x_*\|_1 \\
&\leq \left(1 - \frac{c_{\rho_s}}{\rho_1 N}\right)^{\tilde{k}} (\mathcal{J}_{\varepsilon_0}(x^0) - \|x_*\|_1) \\
&\leq 3 \left(1 - \frac{c_{\rho_s}}{\rho_1 N}\right)^{\tilde{k}} \sigma_s(x^0)_{\ell_1},
\end{aligned}$$

where in the first inequality we applied Lemma B.1. In the second inequality we used that the sequence $\{\mathcal{J}_{\varepsilon_\ell}(x^\ell)\}_\ell$ is monotonically decreasing and in the third inequality we used inequality (26). In the fourth inequality we again used Lemma B.1. By rearranging terms and using the assumption $\rho_s < 1/8$ it follows for all integers \tilde{k} and k such that $\tilde{k} \leq \hat{k}$ and $k \geq \tilde{k}$

$$\|x^k - x_*\|_1 \leq 6 \left(1 - \frac{c_{\rho_s}}{\rho_1 N}\right)^{\tilde{k}} \sigma_s(x^0)_{\ell_1} + 4\sigma_s(x_*)_{\ell_1}. \quad (27)$$

In order to proceed, recall that S denotes the support of the s largest entries of x_* . Then we note that

$$\sigma_s(x^0)_{\ell_1} \leq \|x_{S^c}^0\|_1 \leq \|(x^0 - x_*)_{S^c}\|_1 + \|(x_*)_{S^c}\|_1 \leq \|x^0 - x_*\|_1 + \sigma_s(x_*)_{\ell_1}. \quad (28)$$

Hence, we have shown that for all integers \tilde{k} and k such that $\tilde{k} \leq \hat{k}$ and $k \geq \tilde{k}$ it holds that

$$\|x^k - x_*\|_1 \leq 6 \left(1 - \frac{c_{\rho_s}}{\rho_1 N}\right)^{\tilde{k}} \|x^0 - x_*\|_1 + 10\sigma_s(x_*)_{\ell_1}. \quad (29)$$

By setting $k = \tilde{k}$, we observe that this implies inequality (13), which shows the second statement. In order to prove the third statement, we will distinguish two cases. For the first case, assume that $\hat{k} \geq \left\lceil \frac{\rho_1 N}{c_{\rho_s}} \log \left(\frac{\|x^0 - x_*\|_1}{\sigma_s(x_*)_{\ell_1}} \right) \right\rceil$. Then for $k \geq \tilde{k} := \left\lceil \frac{\rho_1 N}{c_{\rho_s}} \log \left(\frac{\|x^0 - x_*\|_1}{\sigma_s(x_*)_{\ell_1}} \right) \right\rceil$ it follows from inequality (29) that

$$\|x^k - x_*\|_1 \leq 6 \left(1 - \frac{c_{\rho_s}}{\rho_1 N}\right)^{\frac{\rho_1 N}{c_{\rho_s}} \log \left(\frac{\|x^0 - x_*\|_1}{\sigma_s(x_*)_{\ell_1}} \right)} \|x^0 - x_*\|_1 + 10\sigma_s(x_*)_{\ell_1} \leq 20\sigma_s(x_*)_{\ell_1}.$$

where in the second inequality we have used the elementary inequality $\log(1+t) \leq t$ for $t > -1$. This shows the third statement in the first case. To prove the second case, assume that $\hat{k} < \left\lceil \frac{\rho_1 N}{c_{\rho_s}} \log \left(\frac{\|x^0 - x_*\|_1}{\sigma_s(x_*)_{\ell_1}} \right) \right\rceil$. Then we can compute that

$$\begin{aligned}
\frac{1 - \rho_s}{1 + \rho_s} \|x^k - x_*\|_1 - 2\sigma_s(x_*)_{\ell_1} &\leq \mathcal{J}_{\varepsilon_k}(x^k) - \|x_*\|_1 \\
&\leq \mathcal{J}_{\varepsilon_{\hat{k}}}(x^{\hat{k}}) - \|x_*\|_1 \\
&\leq 3\sigma_s(x^{\hat{k}})_{\ell_1} \\
&\leq 3\|x^{\hat{k}} - x_*\|_1 + 3\sigma_s(x_*)_{\ell_1} \\
&\leq 3(1 + \rho_s) \|(x^{\hat{k}} - x_*)_{S^c}\|_1 + 3\sigma_s(x_*)_{\ell_1} \\
&\leq 20\sigma_s(x_*)_{\ell_1}.
\end{aligned}$$

In the first and third inequality we have used Lemma B.1. In the second inequality we have used the monotonicity of the sequence $\{\mathcal{J}_{\varepsilon_k}(x^k)\}_k$. In the fourth inequality we have argued as in inequality (28) and in the fifth inequality we have used the null space property. In the last inequality we have used that by definition of \hat{k} it holds that $\sigma_s(x_*)_{\ell_1} > \frac{2}{9} \|(x_*)_{S^c} - x_{S^c}^{\hat{k}}\|_1$. This shows that also in the second case the third statement holds, which finishes the proof. \square

C Computationally Efficient Weighted Least Squares Updates

As pointed out in Section 2, the weighted least squares step (3) of Algorithm 1 can be computed using the explicit formula provided by the weighted Moore-Penrose inverse such that $x^{k+1} = W_k^{-1} A^* (A W_k^{-1} A^*)^{-1} (y)$ where $W_k = \text{diag}(w_k)$, see e.g. [32, Chapter 15.3]. This leads to the main computational step of solving the positive linear system

$$(A W_k^{-1} A^*) z = y \quad (30)$$

of size $(m \times m)$, which is solved in practice often by an iterative solver such as the conjugate gradient method [38, 49], especially if the measurement matrix A allows for fast matrix-vector multiplications [30, 67]. The authors of [30] provide conditions on the accuracy of conjugate gradient (CG) solvers of the successive linear systems for an IRLS algorithm for basis pursuit using weights that are given by

$$(w_k)_i = 1 / \sqrt{|x_i^k|^2 + \varepsilon_k^2} \quad (31)$$

for each $i \in [n]$ (coinciding with the choice of the weights of [3, 23]).

It was observed that combining IRLS with CG provides a competitive algorithm for sparse recovery if the required accuracy is chosen appropriately [30], however, with the problem that as the IRLS method approaches a sparse solution, and therefore, the smoothing parameter ε_k becomes small, the linear system matrix $A W_k^{-1} A^*$ suffers more and more from ill-conditioning [30, Section 5.2]. To the best of our knowledge, this numerical issue has not been resolved for any IRLS method for sparse recovery in the literature so far.

In this section, we provide an implementation of the weighted least squares step (3) that avoids the ill-conditioning issue of previous IRLS methods and we sketch why the condition number of the resulting linear system can indeed be bounded under standard assumptions.

The starting point of this implementation is the observation that for weight choice (5) used in this paper, i.e., for

$$(w_k)_i := \frac{1}{\max(|x_i^k|, \varepsilon_k)}$$

for each $i \in [N]$, unlike for (31), it is possible to write the inverse weight matrix $W_k^{-1} \in \mathbb{R}^{N \times N}$ such that

$$W_k^{-1} = W_{I_k}^{-1} + \varepsilon_k (\text{Id}_N - P_{I_k}) = (W_{I_k}^{-1} - \varepsilon_k P_{I_k}) + \varepsilon_k \text{Id}, \quad (32)$$

if $I_k := \{i \in [N] : |x_i^k| > \varepsilon_k\}$ denotes the set I of the proof of Proposition B.3, $W_{I_k}^{-1} \in \mathbb{R}^{N \times N}$ is the diagonal matrix with entries $|x_i^k|$ if $i \in I_k$ and 0 otherwise, and where P_{I_k} denotes the projection matrix such that $P_{I_k} x = x_{I_k}$. In particular, we observe that $(W_k)^{-1}$ is the sum of a scaled identity matrix Id and a (diagonal) matrix with only $s_k := |I_k|$ non-zero entries. Furthermore, due to the update rule (4) of the smoothing parameter ε_k , it can be seen that s_k is small and of the order of s if ε_k approaches 0, i.e., if the k -th iterate x^k of Algorithm 1 has only small coordinates outside a set of s large coordinates.

By the Sherman-Morrison-Woodbury matrix inversion formula [70], cf. Lemma C.1 below, it is therefore possible to reformulate the weighted least squares problem such that x^{k+1} can be computed such that solving a positive definite linear system of size $(s_k \times s_k)$, which furthermore is well-conditioned, becomes the main computational step, avoiding solving the ill-conditioned system (30) in this way. We summarize this implementation in Algorithm 2 below.

The implementation uses the matrix $V \in \mathbb{R}^{m \times N}$ with orthonormal columns which denotes a the projection onto the range space of the measurement matrix $A \in \mathbb{R}^{m \times N}$, as well as the left singular vector matrix $U \in \mathbb{R}^{m \times m}$ of A and the diagonal matrix $\Sigma_A \in \mathbb{R}^{m \times m}$ containing the singular values of A . These can be pre-computed before using IRLS, for example via a singular value decomposition of A .⁶ Likewise, the vector

$$\tilde{y} = V \Sigma_A^{-1} (U^* y) \quad (33)$$

⁶In a large scale setting where A is only accessed through matrix-vector multiplications, providing this information is not necessary as respective information can be replaced by using matrix-vector multiplications with A , $(A A^*)^{-1}$ and A^* .

can be pre-computed and can be re-used at each outer iteration of IRLS.

In the following, if $I \subset [N]$, we denote by $M_I \in \mathbb{R}^{m \times |I|}$ the restriction of a matrix $M \in \mathbb{R}^{m \times N}$ to the columns indexed by I , and by $Q_{I_k} \in \mathbb{R}^{N \times I_k}$ the projector matrix such that $P_{I_k} = Q_{I_k} Q_{I_k}^*$. Furthermore, let $D_{I_k}^{-1} \in \mathbb{R}^{I_k \times I_k}$ be a diagonal matrix such that $(D_{I_k}^{-1})_{ii} = |x_i^k|$ for each $i \in I_k$.

Algorithm 2 Practical implementation of weighted least squares step of IRLS for small ε_k

Input: Matrix $V \in \mathbb{R}^{m \times N}$ projecting onto range space of measurement matrix A , $\tilde{y} \in \mathbb{R}^N$ from (33), smoothing parameter ε_k , projection $\gamma_k^{(0)} = Q_{I_k}^* Q_{I_{k-1}}(\gamma_{k-1}) \in \mathbb{R}^{I_k}$ of solution $\gamma_{k-1} \in \mathbb{R}^{I_{k-1}}$ of linear system (34) for previous iteration $k-1$.

1: Compute $h_k^0 = Q_{I_k}^* \tilde{y} - (\varepsilon_k (D_{I_k}^{-1} - \varepsilon_k \text{Id})^{-1} + (V^*)_{I_k}^* (V^*)_{I_k}) \gamma_k^{(0)} \in \mathbb{R}^{I_k}$.

2: Solve

$$(\varepsilon_k (D_{I_k}^{-1} - \varepsilon_k \text{Id})^{-1} + (V^*)_{I_k}^* (V^*)_{I_k}) \Delta \gamma_k = h_k^0 \quad (34)$$

for $\gamma_k \in \mathbb{R}^{I_k}$ by the *conjugate gradient* method [38, 49].

3: Compute $\gamma_k = \gamma_k^{(0)} + \Delta \gamma_k \in \mathbb{R}^{I_k}$.

4: Compute residual $r_{k+1} := \tilde{y} - V(V^*)_{I_k}(\gamma_k) \in \mathbb{R}^N$.

5: Set $x^{k+1} = r_{k+1}$.

6: Set $(x^{k+1})_{I_k} = (x^{k+1})_{I_k} + \gamma_k$.

7: **Output:** $x^{k+1} \in \mathbb{R}^N$ and $\gamma_k \in I_k$.

Next, we verify that x^{k+1} as computed by Algorithm 2 indeed is a solution of the weighted least squares problem (3).

As a preparation, we state the Sherman-Morrison-Woodbury formula for matrix inversion.

Lemma C.1 ([70]). *Let $B \in \mathbb{R}^{n \times n}$, $C \in \mathbb{R}^{k \times k}$, $E \in \mathbb{R}^{n \times k}$ and $F \in \mathbb{R}^{k \times n}$. Then, it holds that*

$$(ECF^* + B)^{-1} = B^{-1} - B^{-1}E(C^{-1} + F^*B^{-1}E)^{-1}F^*B^{-1}$$

With this, we proceed to the proof of the following statement.

Lemma C.2. *If $x^{k+1} \in \mathbb{R}^N$ is the output of Algorithm 2, then x^{k+1} coincides with the solution of the weighted least squares problem (3).*

Proof of Lemma C.2. We recall from above that if $x_*^{k+1} \in \mathbb{R}^N$ is the solution of (3), it holds that $x_*^{k+1} = W_k^{-1}A^*(z)$ where z is as in (30). Using (32), we see that

$$AW_k^{-1}A^* = A((W_{I_k}^{-1} - \varepsilon_k P_{I_k}) + \varepsilon_k \text{Id})A^* = A_{I_k}(D_{I_k}^{-1} - \varepsilon_k \text{Id})A_{I_k}^* + \varepsilon_k AA^*.$$

By identifying $B := \varepsilon_k AA^*$, $C := (D_{I_k}^{-1} - \varepsilon_k \text{Id})$ and $E = F := A_{I_k}$, and by noting that the matrix C is invertible, since, on the set I_k , we have $|x_i^{k+1}| > \varepsilon_k$, we obtain by using Lemma C.1 that

$$(A(W_k)^{-1}A^*)^{-1} = \varepsilon_k^{-1}Z - \varepsilon_k^{-1}ZA_{I_k}G^{-1}A_{I_k}^*Z,$$

using the notation $Z := (AA^*)^{-1}$ and $G := \varepsilon_k C^{-1} + A_{I_k}^*ZA_{I_k}$.

Therefore, denoting the matrix $\varepsilon_k C^{-1} + A_{I_k}^*ZA_{I_k}$ by G , we have

$$z = (A(W_k)^{-1}A^*)^{-1}y = \varepsilon_k^{-1}Z(y) - \varepsilon_k^{-1}ZA_{I_k}G^{-1}A_{I_k}^*Z(y) = \varepsilon_k^{-1}Z(y - A_{I_k}G^{-1}A_{I_k}^*Z(y)) \quad (35)$$

and therefore

$$\begin{aligned} A_{I_k}^*(z) &= \varepsilon_k^{-1}(A_{I_k}^*Z(y) - A_{I_k}^*ZA_{I_k}G^{-1}A_{I_k}^*Z(y)) \\ &= \varepsilon_k^{-1}(A_{I_k}^*Z(y) - A_{I_k}^*ZA_{I_k}(\varepsilon_k C^{-1} + A_{I_k}^*ZA_{I_k})^{-1}A_{I_k}^*Z(y)) \\ &= \varepsilon_k^{-1}\left(A_{I_k}^*Z(y) - (A_{I_k}^*ZA_{I_k} \pm \varepsilon_k C^{-1})(\varepsilon_k C^{-1} + A_{I_k}^*ZA_{I_k})^{-1}A_{I_k}^*Z(y)\right) \\ &= C^{-1}(\varepsilon_k C^{-1} + A_{I_k}^*ZA_{I_k})^{-1}A_{I_k}^*Z(y). \end{aligned} \quad (36)$$

Thus, for the $(k + 1)$ -th iteration of IRLS, we obtain for the solution x_*^{k+1} of (3) the representation

$$\begin{aligned}
x_*^{k+1} &= W_k^{-1} A^*(z) = (\varepsilon_k \text{Id} + Q_{I_k} (D_{I_k}^{-1} - \varepsilon_k \text{Id}) Q_{I_k}^*) A^*(z) \\
&= [\varepsilon_k \text{Id} + Q_{I_k} C Q_{I_k}^*] A^* z = \varepsilon_k A^* z + Q_{I_k} C A_{I_k}^*(z) \\
&= \varepsilon_k A^* z + Q_{I_k} C C^{-1} (\varepsilon_k C^{-1} + A_{I_k}^* Z A_{I_k})^{-1} A_{I_k}^* Z(y) \\
&= \varepsilon_k A^* z + Q_{I_k} G^{-1} A_{I_k}^* Z(y)
\end{aligned} \tag{37}$$

using (36) in the third equality.

Next, if V is as in the input of Algorithm 2, since

$$A_{I_k}^* Z A_{I_k} = Q_{I_k}^* A^* (A A^*)^{-1} A Q_{I_k} = Q_{I_k}^* V V^* Q_{I_k} = (V^*)_{I_k}^* (V^*)_{I_k},$$

we observe that the matrix G from above actually coincides with the linear system matrix of (34).

Furthermore, if γ_k is as in step 3 of Algorithm 2, it satisfies

$$\gamma_k = \gamma_k^{(0)} + \Delta \gamma_k = \gamma_k^{(0)} + G^{-1} \mathbf{h}_k^0 = \gamma_k^{(0)} + G^{-1} (Q_{I_k}^* \tilde{y} - G \gamma_k^{(0)}) = G^{-1} Q_{I_k}^* \tilde{y},$$

where we also observe that

$$Q_{I_k}^* \tilde{y} = Q_{I_k}^* V \Sigma_A^{-1} U^*(y) = Q_{I_k}^* A^* (A A^*)^{-1} (y) = A_{I_k}^* Z(y).$$

Using the latter equation, we can identify z of (35) such that

$$z = \varepsilon_k^{-1} Z(y - A_{I_k} G^{-1} A_{I_k}^* Z(y)) = \varepsilon_k^{-1} Z(y - A_{I_k} G^{-1} Q_{I_k}^* \tilde{y}) = \varepsilon_k^{-1} Z(y - A_{I_k} \gamma_k).$$

Inserting this into (37), we obtain

$$\begin{aligned}
x_*^{k+1} &= \varepsilon_k A^* z + Q_{I_k} G^{-1} A_{I_k}^* Z(y) = \varepsilon_k A^* z + Q_{I_k} G^{-1} Q_{I_k}^* \tilde{y} = \varepsilon_k A^* z + Q_{I_k} \gamma_k \\
&= A^* Z(y - A_{I_k} \gamma_k) + Q_{I_k} \gamma_k \\
&= V \Sigma_A^{-1} U^*(y - U \Sigma_A V^* Q_{I_k} \gamma_k) + Q_{I_k} \gamma_k \\
&= \tilde{y} - V (V^*)_{I_k} \gamma_k + Q_{I_k} \gamma_k \\
&= r_{k+1} + Q_{I_k} \gamma_k,
\end{aligned}$$

where the residual r_{k+1} is as in step 4 of Algorithm 2.

Comparing the equation $x_*^{k+1} = r_{k+1} + Q_{I_k} \gamma_k$ with the steps 5 and 6 of Algorithm 2, we observe that the output x^{k+1} of Algorithm 2 coincides with x_*^{k+1} , which finishes the proof. \square

It is well-known that if the conjugates gradient method is used as an inexact solver using a limited number of iterations only, a sufficient condition for quantifying its accuracy can be achieved by bounding the condition number of the system matrix (see, e.g., [56, Section 5.1]).

In fact, using few iterations of a CG method to solve the system (34) leads typically to quite accurate solutions, in particular if ε_k is small. This can be seen by analyzing the condition number of the matrix G

$$G = \varepsilon_k (D_{I_k}^{-1} - \varepsilon_k \text{Id}_{I_k})^{-1} + A_{I_k}^* Z A_{I_k} =: M_{1,k} + M_{2,k} \tag{38}$$

of eq. (34) in Algorithm 2.

We start by bounding the condition number of $M_{2,k}$. For that, let x such that $\|x\|_2 = 1$. We calculate that

$$\|x^T A_{I_k}^* (A A^*)^{-1} A_{I_k} x\| \leq \|A_{I_k} x\|^2 \|(A A^*)^{-1}\| \leq (1 + \delta)^2 \|(A A^*)^{-1}\|,$$

where in the last inequality we have assumed that A_{I_k} satisfies the restricted isometry property. This implies that

$$\|M_{1,k}\| = \|A_{I_k}^* (A A^*)^{-1} A_{I_k}\| \leq (1 + \delta)^2 \sigma_{\min}(A)^{-2}. \tag{39}$$

In a similar way we can derive that

$$\sigma_{\min}(A_{I_k}^* (A A^*)^{-1} A_{I_k}) \geq (1 - \delta)^2 \|A\|^{-2}. \tag{40}$$

We note that $M_{1,k}$ is a diagonal matrix with entries given by $\varepsilon_k/(x_i^k - \varepsilon_k)$. Now, by additionally assuming that the iteration x_k is already close to the ground truth and, without loss of generality, by assuming that k is such that $\|x^k - x_*\|_\infty \leq 1/4\|x_*\|_\infty$, we have

$$\|M_{1,k}\|_2 = \varepsilon_k/(x_i^k - \varepsilon_k) \leq \varepsilon_k/(\frac{3}{4} \min_i |(x_*)_i| - \varepsilon_k). \quad (41)$$

Note that this term becomes arbitrarily small, when ε_k becomes arbitrarily small.

Hence, for ε_k small enough it follows that

$$\kappa(G) = \frac{\|G\|}{\sigma_{\min}(G)} \stackrel{(a)}{\leq} \frac{\|M_{1,k}\| + \|M_{2,k}\|}{\sigma_{\min}(M_{2,k}) - \|M_{1,k}\|} \stackrel{(b)}{\leq} \frac{2(1+\delta)^2\|A\|^2}{(1-\delta)^2\sigma_{\min}(A)^2} = \frac{2(1+\delta)^2}{(1-\delta)^2} \kappa(A)^2.$$

Here, in (a) we used Weyl's inequality and (b) holds as soon as ε_k is small enough due to inequalities (39), (40), and (41). Hence, we have shown that if A is well-conditioned, then also the matrix G will be well-conditioned and the CG-method will yield very accurate solutions.

D Proofs of technical lemmas

D.1 Proof of Lemma 2.1

In this section, we prove that the chosen quadratic upper bound satisfies the majorization-minimization property.

Proof. We prove each of the three statements separately.

1. Let $x \in \mathbb{R}^N$. Then the i -th coordinate of $\text{diag}(w_\varepsilon(x))x$ is given by

$$(\text{diag}(w_\varepsilon(x))x)_i = \begin{cases} \frac{x_i}{|x_i|} = \text{sgn}(x_i), & \text{if } |x_i| > \varepsilon, \\ \frac{x_i}{\varepsilon}, & \text{if } |x_i| \leq \varepsilon. \end{cases} = j'_\varepsilon(x_i) = (\nabla \mathcal{J}_\varepsilon(x))_i,$$

where $\mathcal{J}_\varepsilon(x)$ is the gradient of \mathcal{J}_ε at x .

2. This follows directly from the definition of $Q_\varepsilon(x, z)$ and by setting $x = z$.
3. We define $I := \{i \in [N] : |x_i| > \varepsilon\}$ and write the difference $Q_\varepsilon(z, x) - \mathcal{J}_\varepsilon(z)$ as

$$\begin{aligned} Q_\varepsilon(z, x) - \mathcal{J}_\varepsilon(z) &= \frac{1}{2} (\langle z, \text{diag}(w_\varepsilon(x))z \rangle - \langle x, \text{diag}(w_\varepsilon(x))x \rangle) \\ &= \sum_{i \in I} \left(\frac{1}{2}|x_i| + \frac{1}{2} \frac{z_i^2}{|x_i|} - j_\varepsilon(z_i) \right) + \sum_{i \in I^c} \left(\frac{1}{2}\varepsilon + \frac{1}{2} \frac{z_i^2}{\varepsilon} - j_\varepsilon(z_i) \right) \end{aligned}$$

and show that each summand of the two sums is non-negative. In particular, if $i \in I$, then assume first that $|z_i| > \varepsilon$. Then

$$\frac{1}{2}|x_i| + \frac{1}{2} \frac{z_i^2}{|x_i|} - j_\varepsilon(z_i) = \frac{1}{2} \left(|x_i| + \frac{z_i^2}{|x_i|} \right) - |z_i| \geq |z_i| - |z_i| = 0$$

due to inequality $a \leq \frac{1}{2}(a^2/b + b)$, which holds for any $b > 0$.

On the other hand, if $|z_i| \leq \varepsilon$, then

$$\begin{aligned} \frac{1}{2}|x_i| + \frac{1}{2} \frac{z_i^2}{|x_i|} - j_\varepsilon(z_i) &= \frac{1}{2}|x_i| + \frac{1}{2} \frac{z_i^2}{|x_i|} - \frac{1}{2} \left(\frac{z_i^2}{\varepsilon} + \varepsilon \right) \\ &= \frac{1}{2} (|x_i| - \varepsilon) + \frac{1}{2} z_i^2 \left(\frac{1}{|x_i|} - \frac{1}{\varepsilon} \right) \\ &\geq \frac{1}{2} (|x_i| - \varepsilon) + \frac{1}{2} \varepsilon^2 \left(\frac{1}{|x_i|} - \frac{1}{\varepsilon} \right) \\ &= \frac{1}{2} \left(|x_i| + \frac{\varepsilon^2}{|x_i|} \right) - \varepsilon \geq \varepsilon - \varepsilon = 0, \end{aligned} \quad (42)$$

where we used that $\frac{1}{|x_i|} - \frac{1}{\varepsilon} < 0$ in the first inequality. In the second inequality, we again used $a \leq \frac{1}{2}(a^2/b + b)$ for any $b > 0$. Now let $i \in I^c$. We again consider the two cases, $|z_i| \leq \varepsilon$ and $|z_i| > \varepsilon$. In the first case we have that $\frac{1}{2}\varepsilon + \frac{1}{2}\frac{z_i^2}{\varepsilon} - j_\varepsilon(z_i) = 0$, and in the second case we have that

$$\frac{1}{2}\varepsilon + \frac{1}{2}\frac{z_i^2}{\varepsilon} - j_\varepsilon(z_i) = \frac{1}{2}\varepsilon + \frac{1}{2}\frac{z_i^2}{\varepsilon} - |z_i| \geq |z_i| - |z_i| = 0,$$

which concludes the proof. □

D.2 Proof of Proposition 4.1

In this section, we prove Proposition 4.1, namely, if x^k , k -th IRLS iteration, is already close enough from the ground truth, then we would observe that the support would have been already identified and, consequently, the hardest part of the sparse recovery problem would have been solved.

Proof. Let $j \in S^c$, where S is the support set of x_* . Then

$$|x_j^k| \leq \sum_{i \in S^c} |x_i^k| < \min_{i \in S} |(x_*)_i| - \sum_{i \in S} |x_i^k - (x_*)_i|,$$

using the assumption $\sum_{i \in S^c} |(x^k)_i| + \sum_{i \in S} |x_i^k - (x_*)_i| = \|x^k - x_*\|_1 < \min_{i \in S} |(x_*)_i|$.

On the other hand, for $j \in S$, we can estimate that

$$|x_j^k| = |x_j^k - (x_*)_j + (x_*)_j| \geq |(x_*)_j| - |x_j^k - (x_*)_j| \geq \min_{i \in S} |(x_*)_i| - \sum_{i \in S} |x_i^k - (x_*)_i|.$$

Taking the previous two inequalities together, we conclude that

$$\max_{j \in S^c} |x_j^k| < \min_{j \in S} |x_j^k|,$$

which finishes the proof. □