

A superlinearly convergent R -regularized Newton scheme for variational models with concave sparsity-promoting priors

Michael Hintermüller · Tao Wu

Received: 19 December 2012 / Published online: 13 July 2013
© Springer Science+Business Media New York 2013

Abstract A general class of variational models with concave priors is considered for obtaining certain sparse solutions, for which nonsmoothness and non-Lipschitz continuity of the objective functions pose significant challenges from an analytical as well as numerical point of view. For computing a stationary point of the underlying variational problem, a Newton-type scheme with provable convergence properties is proposed. The possible non-positive definiteness of the generalized Hessian is handled by a tailored regularization technique, which is motivated by reweighting as well as the classical trust-region method. Our numerical experiments demonstrate selected applications in image processing, support vector machines, and optimal control of partial differential equations.

Keywords Sparsity · Concave priors · Nonconvex minimization · Semismooth Newton method · Superlinear convergence

1 Introduction

Nonsmooth and nonconvex variational models have recently attracted significant attention in mathematical image (or signal) processing, compressed sensing etc. for computing solutions with certain sparsity properties [3, 5, 11, 12, 21, 22, 24, 26, 29, 31, 32]. Particular concave priors, which have been used for this purpose, include

M. Hintermüller (✉)

Department of Mathematics, Humboldt-University of Berlin, Unter den Linden 6, 10099 Berlin, Germany
e-mail: hint@math.hu-berlin.de

T. Wu

Institute for Mathematics and Scientific Computing, Karl-Franzens-University of Graz, Heinrichstrasse 36, 8010 Graz, Austria
e-mail: tao.wu@edu.uni-graz.at

the bridge priori [26], the fraction prior [21], the logarithmic prior [29], to mention only a few. They reflect a nonconvex compromise between minimizing nonsmooth but convex models and the combinatorial ℓ^0 -problem, which aims at minimizing the support of the property of interest of the solution [4, 7].

The success of these variational models in sparse optimization is also related to the fact that they are typically not locally Lipschitz continuous; see, e.g., [9, 12, 14, 24]. Besides the nonconvexity, this property represents significant challenges in both theory and numerics. On the analytical side, the derivation of stationarity conditions becomes delicate as Euler–Lagrange principles, as known for smooth or nonsmooth but still locally Lipschitz problems, cannot be applied. This difficulty, on the other hand, is also reflected in the development of robust solution algorithms with first-order methods, which are prone to slow convergence, and generalized second-order schemes suffering from the nonconvexity and non-Lipschitz behavior of the objective.

Specific stationarity conditions for certain classes of nonconvex and non-Lipschitz problems can be found for instance in [14]. Concerning solution algorithms, we note that in [31] a graduated nonconvexity approach is employed, where each subproblem is solved by an interior point type algorithm, which unfortunately increases the variable space. An alternative can be found in [32], but in that paper non-Lipschitz continuous models are not admissible. The authors in [35] provide an existence result and a surrogate-functional based algorithm for ℓ^q -norm minimization ($0 < q < 1$) in sequence space, where the sparse regularization term is restricted to be separable. In [11] iteratively reweighted techniques are employed, which converge only linearly. On the other hand, the iteratively reweighted least squares method in [17] converges superlinear for the ℓ^q -norm, but requires certain sparsity assumptions on the underlying solution.

The development in this paper extends earlier work by the authors [22] to a more general problem class (e.g. using a variety of different priors), which allows applications in image processing, sparse support vector machines, sparse optimal control of partial differential equations and possibly more within a unified framework. Algorithmically, the method relies on a local smoothing, which is treated asymptotically for vanishing smoothing parameter by a so-called *smoothing scheme* in the spirit of [12]. For each fixed smoothing parameter, a semismooth Newton solver is developed which relies on a specific *R-regularization* of the generalized Hessian in case it is not positive definite. With this regularization, which vanishes as a solution of the associate Euler–Lagrange equation is approached, it can be shown that the method converges superlinearly without any sparsity requirements and for a rather general class of nonconvex priors. The globalization of the Newton scheme is achieved by a Wolfe-type line search procedure. It is demonstrated in the numerical experiments that this new approach outperforms well-known methods ranging from general-purpose nonlinear optimization algorithms such as the BFGS quasi-Newton method and special-purpose solvers for sparse solutions such as the iteratively reweighted least squares method.

The rest of the paper is organized as follows. Section 2 formulates the variational models under consideration and introduces the corresponding Huber relaxation. A superlinearly convergent *R-regularized* Newton scheme, together with its convergence analysis, is presented in Sect. 3. Finally, Sect. 4 demonstrates the practical behavior of

the proposed algorithm and covers a number of applications of the sparsity-promoting variational models.

2 Variational model with concave priors

We consider the following general variational model:

$$\min_{u \in \mathbb{R}^{|\Omega_u|}} f(u) = \Theta(u) + \alpha\Psi(u), \tag{2.1}$$

where Ω_u denotes the multidimensional index set for u . We assume that the fidelity term $\Theta : \mathbb{R}^{|\Omega_u|} \rightarrow \mathbb{R}$ is a coercive and strictly convex C^2 -function. Thus, the Hessian $\nabla^2\Theta(\cdot)$ exists and is positive definite everywhere in $\mathbb{R}^{|\Omega_u|}$.

The prior term Ψ under consideration is of the form

$$\Psi(u) = \sum_{j \in \Omega_p} \psi(|(Gu)_j|),$$

where $G : \mathbb{R}^{|\Omega_u|} \rightarrow \mathbb{R}^{|\Omega_p|}$ is a bounded linear operator and Ω_p is the multidimensional index set for a transformed vector Gu . The scalar function $\psi : [0, \infty) \rightarrow [0, \infty)$ is supposed to satisfy the following hypotheses:

1. (continuity) ψ is continuous on $[0, \infty)$.
2. (regularity) ψ is C^2 on $(0, \infty)$.
3. (mononicity) ψ is strictly increasing on $[0, \infty)$.
4. (concavity) ψ is concave on $[0, \infty)$.

The motivation for a concave prior ψ is to sparsify the solution u under a certain transform G ; see e.g. [29]. Typical choices for G include the identity [19], the gradient operator [22], or some overcomplete dictionary [27]. In particular, we are interested in those situations where $\psi(|\cdot|)$ is non-smooth or even non-Lipschitz at 0. Particular examples for ψ , which have been considered in either a statistical or variational framework, are specified below.

Example 2.1 (Concave priors)

- Bridge prior [24, 26]: $\psi(s) = s^q/q, 0 < q < 1$.
- Fraction prior [21]: $\psi(s) = qs/(1 + qs), q > 0$.
- Logarithmic prior [29]: $\psi(s) = \log(1 + qs), q > 0$.

The proof of existence of a solution for (2.1) is straightforward due to the fact that the objective f is continuous, coercive, and bounded from below.

Theorem 2.2 (Existence of solution) *There exists a global minimizer for (2.1).*

In order to characterize a stationary point for (2.1), we introduce an auxiliary variable $p \in \mathbb{R}^{|\Omega_p|}$ and derive the Euler–Lagrange equation as follows.

Theorem 2.3 (Necessary optimality condition) *For any global minimizer of (2.1) there exists some $p \in \mathbb{R}^{|\Omega_p|}$ such that*

$$\begin{cases} \nabla\Theta(u) + \alpha G^\top p = 0, \\ \varphi(|(Gu)_j|)p_j = (Gu)_j, \quad \text{for all } j \in \Omega_p \text{ with } (Gu)_j \neq 0, \end{cases} \tag{2.2}$$

where $\varphi(s) := s/\psi'(s)$ for any $s \in (0, \infty)$.

Note that since (2.1) is a nonconvex minimization problem, in general there exist more than one stationary point satisfying the Euler–Lagrange system (2.2).

2.1 Huber relaxation

In order to handle the non-smoothness (or even non-Lipschitz continuity) of $\psi(|\cdot|)$ numerically, we introduce a Huber-type local smoothing [25] by defining

$$\psi_\gamma(s) = \begin{cases} \psi(s) - \left(\psi(\gamma) - \frac{\gamma\psi'(\gamma)}{2}\right), & \text{if } |s| \geq \gamma, \\ \frac{\psi'(\gamma)}{2\gamma}s^2, & \text{if } |s| < \gamma, \end{cases}$$

where $\gamma > 0$ is the associated Huber parameter. Then we replace ψ in (2.1) by the C^1 function ψ_γ and formulate the *Huberized* variational model as:

$$\min_{u \in \mathbb{R}^{|\Omega_u|}} f_\gamma(u) = \Theta(u) + \alpha \sum_{j \in \Omega_p} \psi_\gamma(|(Gu)_j|). \tag{2.3}$$

The corresponding Euler–Lagrange equation for (2.3), which we call the Huberized Euler–Lagrange equation, is given by

$$\nabla f_\gamma(u) = \nabla\Theta(u) + \alpha G^\top \left(\varphi(\max(|Gu|, \gamma))^{-1} Gu\right) = 0, \tag{2.4}$$

or equivalently posed with an auxiliary variable p as follows:

$$\text{res}(u, p; \gamma) := \begin{bmatrix} \nabla\Theta(u) + \alpha G^\top p \\ \varphi(\max(|Gu|, \gamma))p - Gu \end{bmatrix} = 0. \tag{2.5}$$

Here and below we shall adopt the following conventions for componentwise operations on vectors. For a vector $v \in \mathbb{R}^n$, $|v|$ denotes the componentwise absolute value, i.e. $|v| = (|v_1|, \dots, |v_n|)$, vw denotes the componentwise product (v_1w_1, \dots, v_nw_n) with another vector $w \in \mathbb{R}^n$, $v^q = (v_1^q, \dots, v_n^q)$ for an exponent $q \in \mathbb{R}$, $\max(v, \gamma) = (\max(v_1, \gamma), \dots, \max(v_n, \gamma))$ for $\gamma \in \mathbb{R}$, and $\varphi(v) = (\varphi(v_1), \dots, \varphi(v_n))$ is also taken componentwisely. In addition, $D(v)$ denotes a diagonal matrix with diagonal elements given by v , and $\|\cdot\|$ denotes the 2-norm for a vector and the spectral norm for a matrix.

Since the argument of φ is bounded below by the positive number γ , the quantity $\varphi(\cdot)$ is well defined. Furthermore, φ satisfies the following properties: (1) φ is continuously differentiable on $(0, \infty)$; (2) $\varphi'(s) \geq \psi'(\gamma)^{-1} > 0$ for any $s \in [\gamma, \infty)$.

Consequently, by the composition rule of semismooth functions [20, Theorem 19], the residual function $\text{res}(\cdot, \cdot; \gamma)$ is semismooth at any $(u, p) \in \mathbb{R}^{|\Omega_u|} \times \mathbb{R}^{|\Omega_p|}$. This allows us to apply the generalized Newton method [23, 34] to (2.5) as we shall see in the next section.

3 Superlinearly convergent R -regularized Newton scheme

In this section, we propose a tailored approach for finding a stationary point for (2.3). We start by investigating a structured regularization scheme in the generalized Newton method.

3.1 R -Regularized Newton scheme

Let (u^k, p^k) be the current iterate and the active set characteristic $\chi_{\mathcal{A}^k} \in \mathbb{R}^{|\Omega_p|}$ be defined as

$$(\chi_{\mathcal{A}^k})_j = \begin{cases} 1, & \text{if } |(Gu^k)_j| \geq \gamma, \\ 0, & \text{if } |(Gu^k)_j| < \gamma. \end{cases}$$

The set $\mathcal{A}^k := \{j \in \Omega_p : |(Gu^k)_j| \geq \gamma\}$ is referred to as the active set. In view of the max-function, we shall apply the generalized Newton method [23, 34] to (2.5). This leads us to the following linear system

$$\begin{aligned} & \begin{bmatrix} \nabla^2 \Theta(u^k) & \alpha G^\top \\ -D(1 - \chi_{\mathcal{A}^k} p^k \frac{\varphi'(m^k) Gu^k}{m^k}) G & D(\varphi(m^k)) \end{bmatrix} \begin{bmatrix} \delta u^k \\ \delta p^k \end{bmatrix} \\ &= \begin{bmatrix} -\nabla \Theta(u^k) - \alpha G^\top p^k \\ -\varphi(m^k) p^k + Gu^k \end{bmatrix}, \end{aligned}$$

with

$$m^k := \max(|Gu^k|, \gamma).$$

After eliminating δp^k , we are left with

$$H^k \delta u^k = -g^k,$$

where

$$\begin{aligned} H^k &= H(u^k, \chi_{\mathcal{A}^k} p^k) \\ &= \nabla^2 \Theta(u^k) + \alpha G^\top D\left(\varphi(m^k)^{-1} \left(1 - \chi_{\mathcal{A}^k} p^k \frac{\varphi'(m^k)(Gu^k)}{m^k}\right)\right) G, \end{aligned} \tag{3.1}$$

$$g^k = \nabla f_\gamma(u^k) = \nabla \Theta(u^k) + \alpha G^\top (\varphi(m^k)^{-1} Gu^k).$$

Based on an observation of the structure of the Hessian matrix H^k (see (3.2) below), we are motivated to define the R -regularization of H^k at $(u^k, \chi_{\mathcal{A}^k} p^k)$ as

$$R^k = R(u^k, \chi_{\mathcal{A}^k} p^k) = \alpha G^\top D \left(\chi_{\mathcal{A}^k} p^k \frac{\varphi'(m^k)(Gu^k)}{\varphi(m^k)m^k} \right) G.$$

Then the resulting R -regularized Newton scheme arises as

$$(H^k + \beta R^k) \delta u^k = -g^k.$$

In particular, if we take $\beta = 1$, then the R -regularized Newton scheme becomes

$$(H^k + R^k) \delta u^k = (\nabla^2 \Theta(u^k) + \alpha G^\top D(\varphi(m^k)^{-1})G) \delta u^k = -g^k. \tag{3.2}$$

Note that the fully R -regularized Hessian $H^k + R^k$ is strictly positive definite and thus guarantees a descent direction δu^k for f_γ at u^k .

It is worth mentioning that a Hessian regularization as in (3.2) can be alternatively derived from applying the generalized Newton method to the *reweighted* Euler–Lagrange equation

$$\nabla \Theta(u) + \alpha G^\top (\varphi(\max(|Gu^k|, \gamma))^{-1} u) = 0. \tag{3.3}$$

The reweighting techniques can be traced back to [10, 37], where they are used in a fixed-point iterative scheme for nonlinear optimization. Here we interpret this technique as a regularization of the Hessian in Newton’s method. This motivates why we have coined the term “ R -regularization”.

3.2 Infeasible Newton technique

In order to ensure fast local convergence of the overall Newton scheme presented in Sect. 3.3, we introduce several modifications in the construction of H^k and R^k .

We start by replacing $\chi_{\mathcal{A}^k} p^k$ by \tilde{p}^k in formula (3.1), where

$$\tilde{p}^k := \frac{\chi_{\mathcal{A}^k}(m^k/\varphi(m^k))p^k}{\max(m^k/\varphi(m^k), |p^k|)}.$$

This choice of \tilde{p}^k satisfies the feasibility condition

$$|(\tilde{p}^k)_j| \leq |(Gu^k)_j|/\varphi(|(Gu^k)_j|),$$

on the index subset $\{j \in \Omega_p : |(Gu^k)_j| \geq \gamma\}$. As a consequence, the modified Hessian \tilde{H}^k appears as

$$\tilde{H}^k = H(u^k, \tilde{p}^k) = \nabla^2 \Theta(u^k) + \alpha G^\top D \left(\varphi(m^k)^{-1} \left(1 - \tilde{p}^k \frac{\varphi'(m^k)(Gu^k)}{m^k} \right) \right) G.$$

One of our motivations for such a replacement is that the sequence (\tilde{p}^k) will be uniformly bounded provided that the sequence (u^k) is uniformly bounded, which

will be useful in the derivation of global convergence; see Theorem 3.7 below. In addition, the Hessian modification becomes asymptotically invariant, as shown in the following lemma.

Lemma 3.1 *Assume that $\lim_{k \rightarrow \infty} (u^k, p^k) = (u^*, p^*)$ with the limiting pair (u^*, p^*) satisfying the Euler–Lagrange equation (2.5). Then we have*

$$\lim_{k \rightarrow \infty} \|\tilde{H}^k - H^k\| = 0.$$

Proof Based on the structures of \tilde{H}^k and H^k , it suffices to show $\lim_{k \rightarrow \infty} \|\tilde{p}^k - \chi_{\mathcal{A}^k} p^k\| = 0$. Given the assumption, we have for all $j \in \Omega_p$ that $|p^*| = |Gu^*|/\varphi(\max(|Gu^*|, \gamma))$ and therefore

$$\begin{aligned} |\tilde{p}^k - \chi_{\mathcal{A}^k} p^k| &\leq |p^k| \left| \frac{m^k/\varphi(m^k)}{\max(m^k/\varphi(m^k), |p^k|)} - 1 \right| \\ &\rightarrow |p^*| \left| \frac{\max(|Gu^*|, \gamma)/\varphi(\max(|Gu^*|, \gamma))}{\max(\max(|Gu^*|, \gamma)/\varphi(\max(|Gu^*|, \gamma)), |p^*|)} - 1 \right| = 0 \end{aligned}$$

as $k \rightarrow \infty$. Thus the conclusion follows. □

Besides the Hessian modification, we correspondingly define the modified R -regularization by

$$\tilde{R}^k = R(u^k, \tilde{p}^k) = \alpha G^\top D \left(\tilde{p}^k \frac{\varphi'(m^k)(Gu^k)}{\varphi(m^k)m^k} \right) G + \varepsilon I, \tag{3.4}$$

with an arbitrarily fixed parameter $0 < \varepsilon \ll \alpha$.

Lemma 3.2 *Let the assumptions of Lemma 3.1 hold true. Then we have $\lambda_{\min}(\tilde{R}^k) \geq \varepsilon/2$ for all sufficiently large k .*

Proof As the conclusion is trivial given $G = 0$, we proceed with any given $G \neq 0$. It follows from the assumption that

$$\begin{aligned} &\min \left\{ \left(\tilde{p}^k \frac{\varphi'(m^k)(Gu^k)}{\varphi(m^k)m^k} \right)_j : j \in \Omega_p \right\} \\ &= \min \left\{ \left(\frac{\varphi'(m^k)(Gu^k)\chi_{\mathcal{A}^k}(m^k/\varphi(m^k))p^k}{\varphi(m^k)m^k \max(m^k/\varphi(m^k), |p^k|)} \right)_j : j \in \Omega_p \right\} \\ &\geq \min \left(\left\{ \left(\frac{\varphi'(m^k)(Gu^k)(m^k/\varphi(m^k))p^k}{\varphi(m^k)m^k \max(m^k/\varphi(m^k), |p^k|)} \right)_j : j \in \Omega_p \right\} \cup \{0\} \right) \\ &\xrightarrow{k \rightarrow \infty} \min \left(\left\{ \left(\frac{\varphi'(\max(|Gu^*|, \gamma))(Gu^*)^2}{\max(|Gu^*|, \gamma)(\varphi(\max(|Gu^*|, \gamma)))^2} \right)_j : j \in \Omega_p \right\} \cup \{0\} \right) \geq 0. \end{aligned}$$

Therefore, we have for all sufficiently large k that

$$\min \left\{ \tilde{p}^k \frac{\varphi'(m^k)(Gu^k)}{\varphi(m^k)m^k} : j \in \Omega_p \right\} \geq -\frac{\varepsilon}{2\|G\|^2},$$

and

$$v^\top \tilde{R}^k v \geq -\frac{\varepsilon}{2\|G\|^2} \|Gv\|^2 + \varepsilon \|v\|^2 \geq \frac{\varepsilon}{2} \|v\|^2,$$

for any vector $v \in \mathbb{R}^{|\Omega_p|}$. Thus we conclude that $\lambda_{\min}(\tilde{R}^k) \geq \varepsilon/2$. □

The ε -term in (3.4) is important as indicated by Lemma 3.2 since it guarantees \tilde{R}^k to be strictly positive definite when the iterate is sufficiently close to a solution. However, note that ε can be arbitrarily small and therefore \tilde{R}^k is allowed to have nonpositive eigenvalues during the Newton iterations. In fact, choosing a large ε that dominates the R -regularization term is not desirable in the numerical implementation.

Thus far, we arrive at the overall modified R -regularized Newton scheme

$$(\tilde{H}^k + \beta \tilde{R}^k) \delta u^k = -g^k. \tag{3.5}$$

The fully R -regularized scheme, i.e. with $\beta = 1$, generates a descent direction satisfying the estimate in the following theorem.

Theorem 3.3 (Sufficient condition for descent property) *Assume that the sequence (u^k) is uniformly bounded and contained in a compact subset E in $\mathbb{R}^{|\Omega_u|}$. Then the solution δu^k of (3.5) with $\beta = 1$ is a descent direction satisfying*

$$-\frac{(g^k)^\top \delta u^k}{\|g^k\| \|\delta u^k\|} \geq \frac{C_l}{C_u + \alpha \lambda_{\max}(G^\top G) / \varphi(\gamma)} =: \bar{\varepsilon}_d,$$

where $0 < C_l \leq C_u$ are two constants depending on Θ and E .

Proof Let $S = \{v \in \mathbb{R}^{|\Omega_u|} : \|v\| = 1\}$ denote the unit sphere. Due to the compactness of $E \times S$ and the continuity of the functional $(u, v) \mapsto v^\top \nabla^2 \Theta(u) v$, the problem

$$C_l := \inf_{(u,v) \in E \times S} v^\top \nabla^2 \Theta(u) v$$

attains the infimum C_l for some $(\underline{u}, \underline{v}) \in E \times S$. Note that $C_l > 0$, since otherwise our assumption that Θ is a strictly convex C^2 function would be violated.

Analogously, there exists a constant C_u such that

$$C_u := \sup_{(u,v) \in E \times S} v^\top \nabla^2 \Theta(u) v.$$

Obviously, we have $C_u \geq C_l$. Then it follows that

$$\begin{aligned} -\frac{(g^k)^\top \delta u^k}{\|g^k\| \|\delta u^k\|} &\geq \frac{\lambda_{\min}(\tilde{H}^k + \tilde{R}^k)}{\lambda_{\max}(\tilde{H}^k + \tilde{R}^k)} \\ &\geq \frac{\lambda_{\min}(\nabla^2 \Theta(u^k))}{\lambda_{\max}(\nabla^2 \Theta(u^k)) + \lambda_{\max}(\alpha G^\top D(\varphi(m^k)^{-1})G)} \\ &\geq \frac{C_l}{C_u + \alpha \lambda_{\max}(G^\top G) / \varphi(\gamma)}. \end{aligned}$$

For the last inequality, we have used the fact that φ is monotonically increasing on $[\gamma, \infty)$. □

3.3 A superlinearly convergent algorithm

According to Theorem 3.3, the R -regularized Newton scheme (3.5) with $\beta = 1$ provides a descent direction. However, a constant R -regularization (with $\beta = 1$), which is equivalent to a fixed-point approach, is known to be only linearly convergent [8, 30, 37].

Ideally, we would like to utilize a sufficient R -regularization when the objective is nonconvex (or the Hessian possesses negative eigenvalues) at the current iterate. As the iterative scheme proceeds, the iterate may eventually be contained in a neighborhood of some local minimizer satisfying some type of a second-order sufficient optimality condition, such that all (generalized) Hessians of the objective are positive definite within that neighborhood. Under such circumstances, we would rather utilize the true Hessian without any R -regularization in the Newton scheme, as it leads to local superlinear convergence.

In order to achieve these goals, we establish an automated scheme for adapting β , the weight of the R -regularization, motivated by the trust-region method; see e.g. [16, 18, 33] for comprehensive discussions on the trust-region method (but in a completely different context). We start by setting up the local quadratic model of f_γ at the current iterate u^k as follows

$$h^k(d) := f_\gamma(u^k) + (g^k)^\top d + \frac{1}{2} d^\top \tilde{H}^k d.$$

Consider now the minimization of $h^k(\cdot)$ subject to a trust-region constraint, i.e.

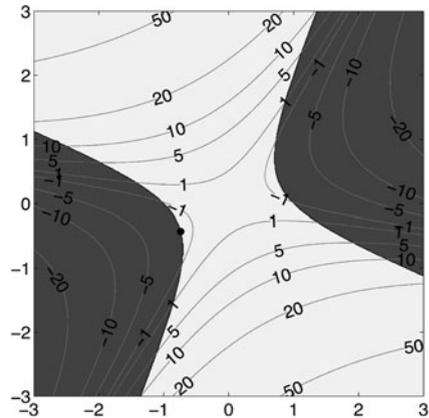
$$\text{minimize } h^k(d), \tag{3.6}$$

$$\text{subject to } d^\top \tilde{R}^k d \leq (\sigma^k)^2, \tag{3.7}$$

where $\sigma > 0$ is the trust-region radius.

Note that the matrix \tilde{H}^k may be indefinite due to the nonconvexity of f_γ . Furthermore, as pointed out after Lemma 3.2, \tilde{R}^k is allowed to have more than one non-positive eigenvalues. Thus, the feasible region induced by (3.7) may be nonconvex and unbounded; see Fig. 1 for an illustration in two dimensions. This is significantly different from the settings in classical trust-region methods [16, 18] where \tilde{R}^k is positive definite and induces a convex, closed and bounded feasible region. Remarkably, \tilde{H}^k and \tilde{R}^k enjoy a special interplay as indicated in the following lemma.

Fig. 1 Illustration of a two-dimensional trust-region subproblem (3.6)–(3.7). The objective function is plotted with contour lines. The feasible region is colored in light gray (contrary to dark gray). The global minimizer $(-0.748, -0.403)$ is marked by the solid dot



Lemma 3.4 *The matrix \tilde{H}^k is positive definite on the subset $\{d \in \mathbb{R}^{|\Omega_u|} : d^\top \tilde{R}^k d \leq 0\}$.*

Proof Let $d \in \mathbb{R}^{|\Omega_u|}$ such that $d \neq 0$ and $d^\top \tilde{R}^k d \leq 0$. Then we have

$$\begin{aligned} d^\top \tilde{H}^k d &= d^\top \nabla^2 \Theta(u^k) d + \alpha (Gd)^\top D(\varphi(m^k)^{-1}) Gd - d^\top \tilde{R}^k d + \varepsilon \|d\|^2 \\ &\geq d^\top \nabla^2 \Theta(u^k) d + \alpha (Gd)^\top D(\varphi(m^k)^{-1}) Gd > 0, \end{aligned}$$

which proves the assertion. □

Such an interplay between \tilde{H}^k and \tilde{R}^k leads us to the existence as well as the characterization of a global minimizer for the trust-region subproblem (3.6)–(3.7), as stated in the following theorem. The proof techniques can be found in Theorems 3.4 and 3.5 in [22].

Theorem 3.5 *There exists a global minimizer d_* for (3.6)–(3.7). Moreover, the necessary and sufficient condition for d_* being optimal is that there exists $\beta_* \geq 0$ such that*

$$(\tilde{H}^k + \beta_* \tilde{R}^k) d_* = -g, \tag{3.8}$$

$$\beta_* - \max(\beta_* + c^{-1} (d_*^\top \tilde{R}^k d_* - (\sigma^k)^2), 0) = 0, \tag{3.9}$$

$$\tilde{H}^k + \beta_* \tilde{R}^k \geq 0, \tag{3.10}$$

for an arbitrarily fixed scalar $c > 0$.

In particular, the complementarity equation (3.9) in Theorem 3.5 provides us a natural fixed-point formula for updating the weight β . Now we are in a position to present our superlinearly convergent R -regularized Newton scheme.

Algorithm 3.6 (Superlinearly convergent R -regularized Newton scheme)

Require: parameters $c > 0$, $0 < \rho_1 \leq \rho_2 < 1$, $0 < \kappa_1 < 1 < \kappa_2$, $0 < \varepsilon \ll \alpha$, $0 < \epsilon_d \leq \bar{\epsilon}_d$, $0 < \tau_1 < 1/2$, $\tau_1 < \tau_2 < 1$.

- 1: Initialize the iterate (u^0, p^0) , the regularization weight $\beta^0 \geq 0$, and the trust-region radius $\sigma^0 > 0$. Set $k := 0$.
- 2: **repeat** {outer loop}
- 3: Generate \tilde{H}^k , \tilde{R}^k , and g^k at the current iterate (u^k, p^k) .
- 4: **repeat** {inner loop}
- 5: Solve the linear equation $(\tilde{H}^k + \beta^k \tilde{R}^k)d^k = -g^k$ for d^k .
- 6: **if** the matrix $\tilde{H}^k + \beta^k \tilde{R}^k$ is singular **or** $-(g^k)^\top d^k / (\|g^k\| \|d^k\|) < \epsilon_d$ **then**
- 7: Set $\beta^k := 1$, and return to Step 5.
- 8: **end if**
- 9: **if** $\beta^k = 1$ **and** $(d^k)^\top \tilde{R}^k d^k > (\sigma^k)^2$ **then**
- 10: Set $\sigma^k := \sqrt{(d^k)^\top \tilde{R}^k d^k}$, and go to Step 15.
- 11: **end if**
- 12: Update $\beta^k := \beta^k + c^{-1}((d^k)^\top \tilde{R}^k d^k - (\sigma^k)^2)$.
- 13: Project β^k onto the interval $[0, 1]$, i.e. set $\beta^k := \max(\min(\beta^k, 1), 0)$.
- 14: **until** the stopping criterion for the inner loop is fulfilled.
- 15: Evaluate $\rho^k := [f_\gamma(u^k) - f_\gamma(u^k + d^k)]/[f_\gamma(u^k) - (f_\gamma(u^k) + (g^k)^\top d^k + (d^k)^\top \tilde{H}^k d^k/2)]$.
- 16: **if** $\rho^k < \rho_1$ **then**
- 17: Set $\sigma^{k+1} := \kappa_1 \sigma^k$.
- 18: **else if** $\rho^k > \rho_2$ **then**
- 19: Set $\sigma^{k+1} := \kappa_2 \sigma^k$.
- 20: **else**
- 21: $\sigma^{k+1} := \sigma^k$.
- 22: **end if**
- 23: Determine the step size a^k along the search direction d^k such that $u^{k+1} = u^k + a^k d^k$ satisfies the following Wolfe–Powell conditions:

$$f_\gamma(u^{k+1}) \leq f_\gamma(u^k) + \tau_1 a^k \nabla f_\gamma(u^k)^\top d^k, \tag{3.11}$$

$$\nabla f_\gamma(u^{k+1})^\top d^k \geq \tau_2 \nabla f_\gamma(u^k)^\top d^k. \tag{3.12}$$

- 24: Generate the next iterate:

$$u^{k+1} := u^k + a^k d^k, \tag{3.13}$$

$$p^{k+1} := \varphi(m^k)^{-1} \left(Gu^k + \left(1 - \tilde{p}^k \frac{\varphi'(m^k)(Gu^k)}{m^k} \right) Gd^k \right).$$

- 25: Initialize the R -regularization weight $\beta^{k+1} := \beta^k$ for the next iteration.
- 26: Set $k := k + 1$.
- 27: **until** the stopping criterion for the outer loop is fulfilled.

Concerning the input parameters involved in the above algorithm, we note that these quantities are presented merely for the generality of the algorithm and do not

require particular tuning for different test runs. Throughout our numerical experiments in Sect. 4, we shall always fix the parameters as follows: $c = 1$, $\rho_1 = 0.25$, $\rho_2 = 0.75$, $\kappa_1 = 0.25$, $\kappa_2 = 2$, $\varepsilon = 10^{-4}\alpha$, $\epsilon_d = 10^{-8}$, $\tau_1 = 0.1$, $\tau_2 = 0.9$.

We remark that Algorithm 3.6 is a hybrid approach combining the trust-region method and the line search technique. The Wolfe–Powell line search, along the search direction d^k obtained from the R -regularized Newton scheme, is responsible for the global convergence of the overall algorithm; see Theorem 3.7 in the following.

Theorem 3.7 (Global convergence) *Let (u^k, p^k) be the sequence generated by Algorithm 3.6. Then we have*

$$\lim_{k \rightarrow +\infty} \|\nabla f_Y(u^k)\| = 0. \tag{3.14}$$

Moreover, if in addition (u^k) is uniformly bounded, then the sequence (u^k, p^k) has an accumulation point (u^*, p^*) satisfying the Euler–Lagrange equation (2.5).

Proof According to Zoutendijk’s theorem for the Wolfe–Powell line search (see, e.g., Theorem 3.2 in [33]), we have $\sum_{k=0}^{\infty} \cos^2 \theta^k \|g^k\|^2 < \infty$, where

$$\cos \theta^k := -\frac{(g^k)^\top d^k}{\|g^k\| \|d^k\|}.$$

Due to the descent property

$$\cos \theta^k \geq \epsilon_d > 0, \tag{3.15}$$

guaranteed by Theorem 3.3 and steps 6–8 in Algorithm 3.6, we have proved (3.14).

Moreover, it follows from the descent property (3.15) that

$$\epsilon_d \|g^k\| \|d^k\| \leq -(g^k)^\top d^k = (d^k)^\top (\tilde{H}^k + \beta^k \tilde{R}^k) d^k \leq \|g^k\| \|d^k\|.$$

Consider $d^k := s^k v^k$ such that $s^k \geq 0$ and $\|v^k\| = 1$ for all k , then we have

$$\epsilon_d \|g^k\| \leq s^k (v^k)^\top (\tilde{H}^k + \beta^k \tilde{R}^k) v^k \leq \|g^k\|.$$

It follows that

$$\lim_{k \rightarrow \infty} s^k (v^k)^\top (\tilde{H}^k + \beta^k \tilde{R}^k) v^k = 0. \tag{3.16}$$

By the uniform boundedness of (u^k) , (v^k) , (\tilde{p}^k) , and (β^k) , there exist $u^*, v^* \in \mathbb{R}^{|\Omega_u|}$, $\tilde{p}^* \in \mathbb{R}^{|\Omega_p|}$, and $\beta^* \in [0, 1]$ such that up to a subsequence $u^k \rightarrow u^*$, $v^k \rightarrow v^*$, $\tilde{p}^k \rightarrow \tilde{p}^*$, and $\beta^k \rightarrow \beta^*$ as $k \rightarrow \infty$. Owing to the continuity of the mappings $(u^k, \tilde{p}^k) \mapsto \tilde{H}^k = H(u^k, \tilde{p}^k)$ and $(u^k, \tilde{p}^k) \mapsto \tilde{R}^k = R(u^k, \tilde{p}^k)$, we also have $\tilde{H}^k \rightarrow \tilde{H}^* := H(u^*, \tilde{p}^*)$ and $\tilde{R}^k \rightarrow \tilde{R}^* := R(u^*, \tilde{p}^*)$ as $k \rightarrow \infty$.

We claim that $\liminf_{k \rightarrow \infty} s^k = 0$. Assume the contrary that (s^k) is uniformly bounded away from 0. Then because of (3.16) we have $(v^*)^\top (\tilde{H}^* + \beta^* \tilde{R}^*) v^* = 0$, or equivalently

$$(\tilde{H}^* + \beta^* \tilde{R}^*) v^* = 0,$$

due to the symmetry of the matrix. This leads to a contradiction as

$$\begin{aligned} \epsilon_d &\leq -\frac{(g^k)^\top d^k}{\|g^k\| \|d^k\|} = \frac{(d^k)^\top (\tilde{H}^k + \beta^k \tilde{R}^k) d^k}{\|(\tilde{H}^k + \beta^k \tilde{R}^k)^{-1} d^k\| \|d^k\|} \\ &\leq \frac{((d^k)^\top (\tilde{H}^k + \beta^k \tilde{R}^k) d^k) \|(\tilde{H}^k + \beta^k \tilde{R}^k) d^k\|}{\|d^k\|^3} \\ &\leq \frac{\|(\tilde{H}^k + \beta^k \tilde{R}^k) d^k\|^2}{\|d^k\|^2} \\ &= \|(\tilde{H}^k + \beta^k \tilde{R}^k) v^k\|^2 \xrightarrow{k \rightarrow \infty} \|(\tilde{H}^* + \beta^* \tilde{R}^*) v^*\|^2 = 0. \end{aligned}$$

We have used the Cauchy–Schwarz inequality in deriving the above inequalities. Thus, we have proved that $\liminf_{k \rightarrow \infty} \|d^k\| = 0$.

Upon extracting another subsequence of (d^k) and using again the same notation for the indices, we have $\lim_{k \rightarrow \infty} d^k = 0$ and then

$$p^* := \lim_{k \rightarrow \infty} p^{k+1} = \varphi(\max(|Gu^*|, \gamma))^{-1} Gu^*,$$

according to the update formula (3.13). Together with the already established fact that

$$0 = \lim_{k \rightarrow \infty} \nabla f_\gamma(u^k) = \nabla f_\gamma(u^*) = \nabla \Theta(u^*) + \alpha G^\top (\varphi(\max(|Gu^*|, \gamma))^{-1} Gu^*),$$

we conclude that (u^*, p^*) satisfies the Euler–Lagrange equation (2.5). □

In addition to the global convergence, the trust-region framework supplies a proper tuning of the R -regularization weight β^k , such that β^k will vanish asymptotically. Thus the algorithm converges locally at a superlinear rate to a local minimizer satisfying the second-order sufficient optimality condition (for semismooth problems); see Theorem 3.8 below. To sketch the proof, note that for sufficiently large k , \tilde{H}^k and \tilde{R}^k both become strictly positive definite. It follows that the alternating iterations on β^k and u^k , i.e. steps 4–14 of Algorithm 3.6, converge and therefore the so-called *Cauchy-point based model reduction criterion* will be satisfied; see [22]. Analogous to the classical trust-region method, the evaluation ratio ρ^k tends to 1 and the trust-region radius σ^k is uniformly bounded away from 0. As a result the weight β^k will vanish in the limit. Finally, the full step size $a^k = 1$ is admissible for all sufficiently k and the step d^k is asymptotically identical to a full semismooth Newton step. We refer to Theorem 3.10 in [22] for a complete proof of the local superlinear convergence.

Theorem 3.8 (Local superlinear convergence) *Let the sequence (u^k, p^k) generated by Algorithm 3.6 converge to some (u^*, p^*) satisfying the Euler–Lagrange equation (2.5). Assume that all generalized Hessians [22] of f_γ at u^* are strictly positive definite. Then we have $\lim_{k \rightarrow \infty} \beta^k = 0$ and the sequence (u^k) converges to u^* superlinearly, i.e.*

$$\|u^{k+1} - u^*\| = o(\|u^k - u^*\|), \quad \text{as } k \rightarrow \infty.$$

4 Numerical examples

In this section, we present a numerical study of Algorithm 3.6. Throughout this section, the linear system in step 5 is handled by the conjugate gradient method with residual tolerance 0.01. Whenever, the matrix $\tilde{H}^k + \beta^k \tilde{R}^k$ is detected to be indefinite or (near-) singular, we immediately utilize step 7 in order to obtain a positive definite linear system. The trust-region subproblem (3.6)–(3.7) is solved only approximately. From our numerical experience, one (inner) iteration on the R -regularization weight β^k seems efficient for the overall algorithm. The regularization parameter α is manually chosen to properly balance the data fidelity and the sparsity-promoting prior. The Huber parameter γ is selected to be sufficiently small, depending on the particular application, in order to obtain a desirable sparse solution. We terminate the overall algorithm once the residual norm $\|\text{res}(u^k, p^k; \gamma)\|$ is reduced by a factor of 10^{-7} relative to its initial value.

The remainder of the section is organized as follows. In Sect. 4.1, we apply our algorithm to the “classical” problem of a sparse solution to an undetermined linear system. Therein we provide detailed studies and comparisons on various aspects of the algorithm in practice. Sections 4.2–4.4 present selected applications of the general variational framework (2.1) in image processing, feature selection, and optimal control. All experiments in this section are performed under MATLAB R2011b on a 2.66 GHz Intel Core Laptop with 4 GB RAM. The CPU time reported in this section is measured in seconds.

4.1 Sparse solution to underdetermined linear systems

Many problems such as compressed sensing and basis pursuit can be regarded as special cases of finding a sparse solution to an underdetermined linear system; see, e.g., [4, 13]. Here we consider a benchmark example of this type taken from [1]. The true data \hat{u} is a sparse vector of n entries, among which only 5 % are nonzero (either 1 or -1 with equal probability). The measurement z is constructed by multiplying \hat{u} by a randomly generated sensing matrix A of size $(n/4)$ -by- n with orthonormal rows and then adding to the product a white Gaussian noise of standard deviation 0.005. In order to find an approximate sparse solution, we shall apply Algorithm 3.6 to the following minimization problem:

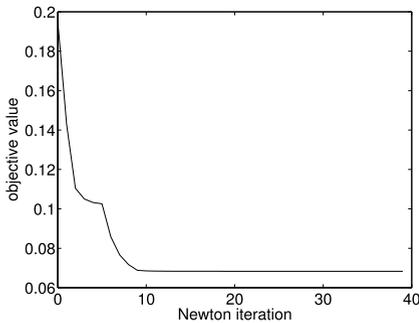
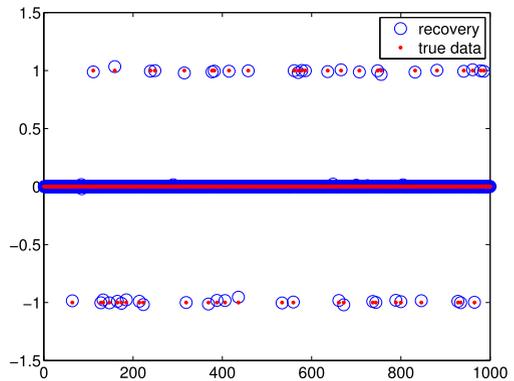
$$\min_{u \in \mathbb{R}^n} f_\gamma(u) = \frac{1}{2} \|Au - z\|^2 + \alpha \sum_{i=1}^n \psi_\gamma(|u_i|),$$

with the chosen parameters $\alpha = 0.001$, $\psi(t) = t^q/q$, $q = 3/4$, $\gamma = 0.001$, and $u^0 = A^\top z$. The restoration result for a trial run with $n = 1000$ is illustrated in Fig. 2, where the red dots denote the true data and a faithful recovery, marked by blue circles, can be observed.

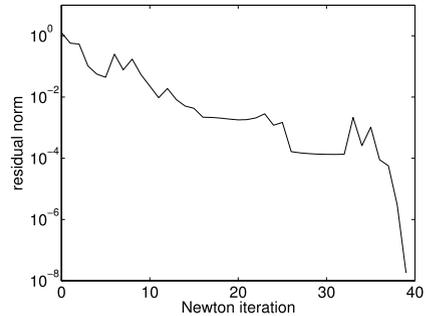
Figure 3 investigates the practical behavior of the algorithm. The history of the objective values over the iterations is plotted in (a). Owing to the descent property of the search direction and the Wolfe–Powell line search along the descent direction, the objective value is monotonically decreasing, which is verified by (a). The residual norm $\|\text{res}(u^k, p^k; \gamma)\|$ displayed in (b) confirms the global convergence claimed

in Theorem 3.7. Moreover, a fast local convergence is observed near the solution. For this fact note that the vertical axis has a logarithmic scale. The R -regularization weight β^k at each (outer) Newton iteration is plotted in (c). Note that the sufficient condition ($\beta^k = 1$) is utilized, e.g., at the 6th Newton iteration in order to obtain a descent direction. On the other hand, β^k tends to zero towards the end of the iteration, permitting fast local convergence. We adopt the implementation in [33] for the Wolfe–Powell line search, for which the history of step sizes is shown in (d). In our

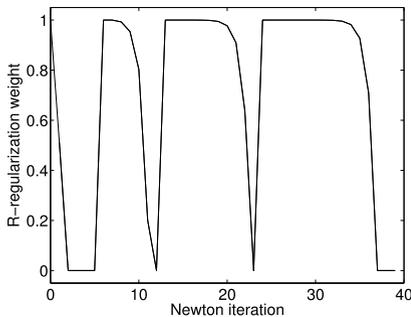
Fig. 2 Sparse solution to underdetermined linear system



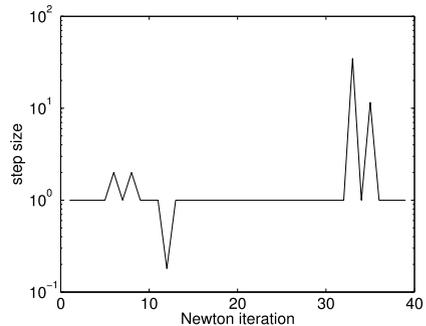
(a) Objective value.



(b) Residual norm.



(c) R -regularization weight.



(d) Step size.

Fig. 3 Behavior of Algorithm 3.6

Table 1 Dependence on initial guess

Initial guess	Objective	Error	CPU
$u^0 = A^\top z$	6.826e-2	1.464e-2	1.99
$u^0 = 0$	6.826e-2	1.467e-2	2.10
Random u^0	6.826e-2	1.502e-2	2.21

Table 2 Comparison with existing algorithms in terms of CPU time

Tolerance	BFGS	IRLS	RRN
1e-4	21.76	2.63	1.73
1e-7	25.62	4.88	2.01

experiments, we find that the Wolfe–Powell rule, which allows a step of size larger than 1, outperforms the backtracking Armijo rule [33] due to the fact that the latter poorly resolves the nonconvex line search problem and therefore causes more Newton iterations.

Since Algorithm 3.6 intends to find a stationary point for the nonconvex minimization problem, it is important to check the quality of such local solutions as well as the dependence of the algorithm on the initial guess. In this experiment, we compare three different choices of initial guesses, namely $u^0 = A^\top z$, $u^0 = 0$, and a randomly chosen u^0 . The corresponding objective value $f_\gamma(u)$, the error measured by $\|u - \hat{u}\|/\|\hat{u}\|$, and the CPU time are reported in Table 1. It is observed that regardless of the initial guess the algorithm always finds a reasonably good numerical solution within nearly the same amount of time.

In Table 2, we compare our R -regularized Newton scheme (RRN) with two other existing algorithms, namely the BFGS quasi-Newton method (see e.g. [33]) and the iteratively reweighted least squares method (see e.g. [11, 17]). For each of the three algorithms under consideration, we test the algorithm with two different tolerance levels on the residual norm. The CPU time is reported in the corresponding entries of Table 2. It is observed that our method always outperform the other two. Notably, the BFGS quasi-Newton method exhibits a fast local convergence in the experiments, but its overall convergence speed suffers from the strong nonlinear nature of the underlying problem. The iteratively reweighted least squares method, a special case of a fixed-point iteration, converges fairly fast at early iterations, but becomes less competitive in case of higher accuracy requirements.

4.2 Image denoising via overcomplete dictionary

In this example, we apply our method to an image denoising problem, where the following $\ell^{1/2}$ -DCT5 model is considered

$$\min_{u \in \mathbb{R}^{|\Omega|}} \frac{1}{2} \|u - z\|^2 + \sum_{l=1}^{24} \sum_{(j_1, j_2) \in \Omega} \alpha_l \psi_\gamma(|(h_l * u)_{j_1, j_2}|). \quad (4.1)$$

Here, z is the observed image (see Fig. 5(b)), which is generated by adding white Gaussian noise of standard deviation 25/255 to the ‘‘Cameraman’’ image (see

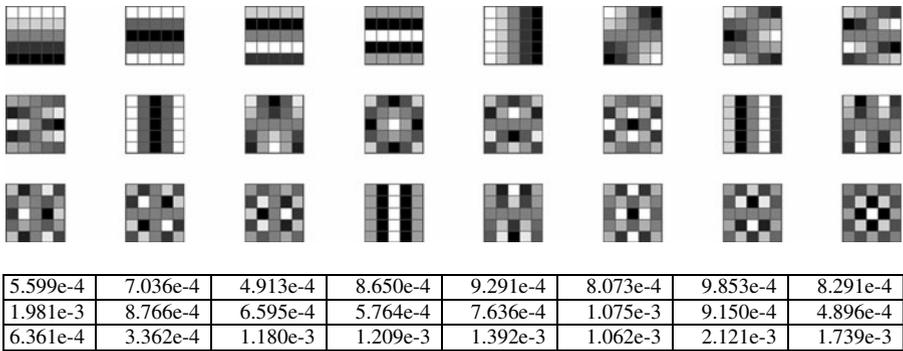


Fig. 4 DCT5 filters and regularization parameters

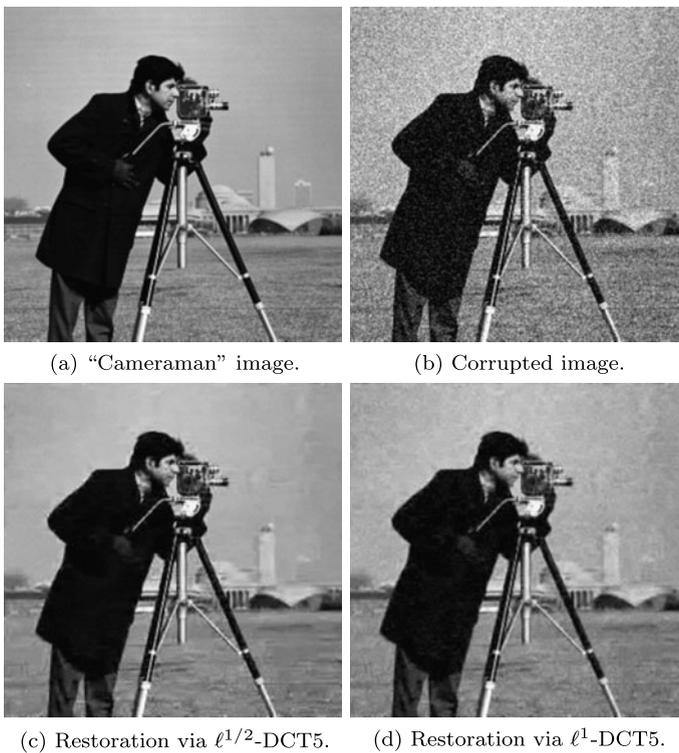


Fig. 5 Image denoising via overcomplete dictionary

Fig. 5(a)). The filters $(h_l)_{l=1}^{24}$ are the two-dimensional 5th-order discrete cosine transform (DCT5) filters, and correspondingly $(\alpha_l)_{l=1}^{24}$ are the regularization parameters trained from a large database of image patches [27]; see Fig. 4 for the illustrations of DCT5 filters and the values of trained regularization parameters. The symbol “*” denotes the conventional two-dimensional convolution. By considering the concave

Table 3 Dependence on the Huber parameter

γ	0.03	0.02	0.015	0.01
PSNR	26.61	27.47	27.91	28.25
CPU	50.88	61.03	122.7	234.1

Table 4 Comparison of $\ell^{1/2}$ - and ℓ^1 -models

Model	PSNR	#Newton	#CG	CPU
ℓ^1 -DCT5	27.46	11	104	80.37
$\ell^{1/2}$ -DCT5	28.25	24	358	234.1

bridge prior with exponent 1/2 (or $\psi(t) = 2t^{1/2}$), we expect the restored image u to be sparse under the DCT5 transform. In this sense, the variational model (4.1) is an *analysis* approach [6].

We implement Algorithm 3.6 with the initial guess $u^0 = z$ and different choices of the Huber parameter, namely $\gamma = 0.03, 0.02, 0.015$, and 0.01 . The quality of the restored image is measured by the peak signal-to-noise ratio (PSNR). The corresponding PSNR and CPU time with respect to different γ are reported in Table 3. We note the tradeoff in γ -selection that smaller γ typically yields the higher quality on the sparse solution, but costs more CPU time.

We further compare the performance of the $\ell^{1/2}$ -DCT5 model (with $u^0 = z$, $\gamma = 0.01$) and that of ℓ^1 -DCT5 model in [27], for which the corresponding restored images are displayed in (c) and (d) of Fig. 5, respectively. Table 4 reports the quantitative comparison of the two models with respect to the PSNR value, the number of Newton iterations, the total number of conjugate gradient iterations, and the CPU time. It is observed that the ℓ^1 -DCT5 model poorly restores the homogeneous region in order to well preserve the textures in the image. The $\ell^{1/2}$ -DCT5 model is more time-consuming due to solving a nonconvex problem, but yields considerable improvement on the restoration quality.

4.3 Feature selection via sparse support vector machines

We consider an example of feature selection using a support vector machine (SVM) [38], where we aim to identify 10 feature variables out of 200 candidate variables $(x_j)_{j=1}^{200}$. The identification is based on n training samples simulated as follows. For each sample, the outcome $y^i \in \{+1, -1\}$, $i \in \{1, 2, \dots, n\}$, is generated with equal probability. If x_j is a feature variable, then with probability 0.3 the random variable $x_j^i = y^i \mathcal{N}(3, 1)$ is drawn and with probability 0.7 we generate $x_j^i = \mathcal{N}(0, 1)$. If x_j is a noise variable, then $x_j^i = \mathcal{N}(0, 1)$ is independently generated.

The linear SVM uses the classifier $y = \text{sgn}(b + \sum_{j=1}^{200} w_j x_j)$ to predict the outcome for a fresh input x . The unknowns $b \in \mathbb{R}$ and $w \in \mathbb{R}^{200}$ are determined by solving the following minimization problem

$$\min_{b \in \mathbb{R}, w \in \mathbb{R}^{200}} \alpha \sum_{j=1}^{200} \psi_\gamma(|w_j|) + \frac{1}{n} \sum_{i=1}^n L_{\epsilon_{hl}} \left(y^i \left(b + \sum_{j=1}^{200} w_j x_j^i \right) \right), \quad (4.2)$$

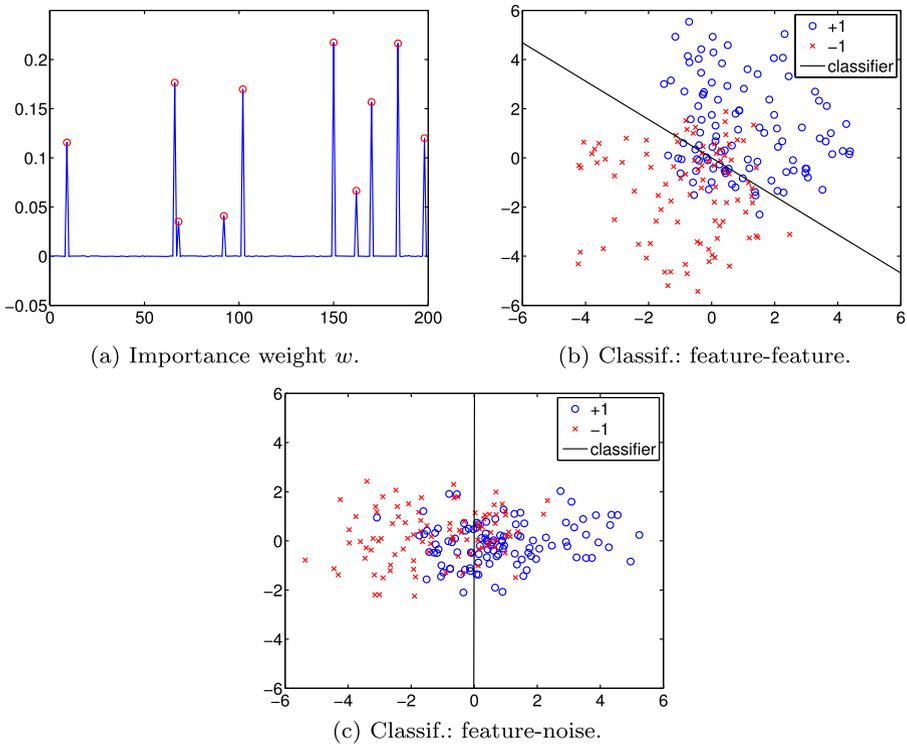


Fig. 6 Feature selection via sparse support vector machine

where $L_{\epsilon_{hl}}(\cdot)$ is a smoothed hinge loss [9] defined by

$$L_{\epsilon_{hl}}(s) = \begin{cases} \max(1 - s, 0), & \text{if } |s - 1| \geq \epsilon_{hl}, \\ (1 + \epsilon_{hl} - s)^2 / (4\epsilon_{hl}), & \text{if } |s - 1| < \epsilon_{hl}, \end{cases}$$

with the smoothing parameter $\epsilon_{hl} = 0.01$. In this experiment, we choose $\alpha = 0.1$, $\psi(t) = \log(1 + 2t)$, and $\gamma = 0.001$.

The computational results for a trial run with $n = 200$ training samples are displayed in Fig. 6. We plot the importance weight w in (a), where in particular the weights for the 10 presumed feature variables are marked by red circles. From this figure, it is observed that the variational model (4.2) has correctly identified the feature variables among all candidate variables. In (b) and (c), we illustrate the computed classifier with respect to the training data projected onto two particular candidate data variables, i.e. $y = \text{sgn}(b + w_{j_1}x_{j_1} + w_{j_2}x_{j_2})$. More specifically, in (b) x_{j_1} and x_{j_2} are two distinct feature variables, and in (c) one is a feature variable and the other is a noise variable. In both figures, the coordinates of the circles indicate the random-variable values of those simulated samples with outcomes $+1$, and the coordinates of the crosses indicate the random-variable values of those simulated samples with outcomes -1 .

4.4 Sparse optimal control

Finally, we demonstrate an application in sparse optimal control, which shows considerable promises in actuator placement problems; see, e.g., [15, 36]. Consider the following stationary control problem:

$$\text{minimize } J(y, u) = \frac{1}{2} \int_{\Omega} |y - z|^2 dx + \frac{\mu}{2} \int_{\Omega} |\nabla u|^2 dx + \alpha \int_{\Omega} \psi(|u|) dx \quad (4.3)$$

$$\text{over } (y, u) \in H_0^1(\Omega) \times U_{ad} \quad (4.4)$$

$$\text{subject to } \int_{\Omega} \nabla y \cdot \nabla v dx = \int_{\Omega} u v dx \text{ for any } v \in H_0^1(\Omega). \quad (4.5)$$

Here Ω is a bounded Lipschitz domain, $\alpha > 0$, $0 < \mu \ll \alpha$ are some given parameters, a desired state is given by $z \in H_0^1(\Omega)$, and U_{ad} is some weakly closed subset in $H_0^1(\Omega)$. A (continuous) concave prior $\psi(\cdot)$ is applied in order to promote the sparsity of the optimal control in the spatial domain.

In general, it is a difficult task to establish the existence of solutions for a nonconvex minimization problem in function space due to the lack of weak (or weak*) lower semicontinuity for the objective; see, e.g., [3, 22]. However, in this special case with the H^1 -regularization (the μ -term), we are able to show the existence of solution as in the following theorem.

Theorem 4.1 *The stationary control problem (4.3)–(4.5) admits a solution.*

Proof By the Lax–Milgram Lemma, the solution mapping $u \mapsto y = (-\Delta)^{-1}u$ for (4.5) is linear and continuous. Thus we only need to consider the reduced problem:

$$\min_{u \in U_{ad}} \hat{J}(u) = \frac{1}{2} \int_{\Omega} |(-\Delta)^{-1}u - z|^2 dx + \frac{\mu}{2} \int_{\Omega} |\nabla u|^2 dx + \alpha \int_{\Omega} \psi(|u|) dx.$$

Since $\hat{J}(\cdot)$ is bounded from below and coercive in $H_0^1(\Omega)$, any infimizing sequence (u^k) is uniformly bounded in $H_0^1(\Omega)$. By the reflexivity of the space $H^1(\Omega)$ and the weak closedness of the admissible set U_{ad} , there exists a subsequence of (u^k) , also denoted by (u^k) , such that $u^k \rightharpoonup u^*$ in $H^1(\Omega)$ as $k \rightarrow \infty$ for some $u^* \in U_{ad}$.

As the functional $u \in H_0^1(\Omega) \mapsto \frac{1}{2} \int_{\Omega} |(-\Delta)^{-1}u - z|^2 dx + \frac{\mu}{2} \int_{\Omega} |\nabla u|^2 dx$ is convex and strongly continuous, it is weakly lower semicontinuous, and thus we have

$$\begin{aligned} & \frac{1}{2} \int_{\Omega} |(-\Delta)^{-1}u^* - z|^2 dx + \frac{\mu}{2} \int_{\Omega} |\nabla u^*|^2 dx \\ & \leq \liminf_{k \rightarrow \infty} \frac{1}{2} \int_{\Omega} |(-\Delta)^{-1}u^k - z|^2 dx + \frac{\mu}{2} \int_{\Omega} |\nabla u^k|^2 dx. \end{aligned}$$

On the other hand, the compact embedding of $H_0^1(\Omega)$ into $L^2(\Omega)$ (see, e.g., Theorem 5.3.3 in [2]) implies the strong convergence of (u^k) to u^* in $L^2(\Omega)$, and thus we

have, up to another subsequence, $u^k(x) \rightarrow u^*(x)$ almost everywhere in Ω as $k \rightarrow \infty$. By Fatou’s lemma and the continuity of the scalar function $\psi(|\cdot|)$, we have

$$\int_{\Omega} \psi(|u^*|)dx \leq \liminf_{k \rightarrow \infty} \int_{\Omega} \psi(|u^k|)dx.$$

Altogether, we have $\hat{J}(u^*) \leq \liminf_{k \rightarrow \infty} \hat{J}(u^k)$, indicating that u^* is an optimal solution to the underlying problem. \square

We remark that without the H^1 -regularization term the above proof would no longer be valid due to the lack of coercivity of the reduced objective $\hat{J}(\cdot)$. In addition, H^1 -regularization enforces sufficient regularity on a weakly convergent (sub)sequence that finally yields the almost everywhere pointwise convergence of the infimizing (sub)sequence.

Now we turn our attention to the numerical solution for the following discretized control problem in reduced form:

$$\min_{u \in \mathbb{R}^{|\Omega|}} \sum_{(j_1, j_2) \in \Omega} \frac{1}{2} |(-\Delta)^{-1}u - z|^2 + \alpha \psi_{\gamma}(|u|) + \frac{\mu}{2} |\nabla u|^2.$$

Here $\Omega = \{0, 1, 2, \dots, 2^N\}^2$, $N \in \mathbb{N}$, denotes the 2D index set for the discretized square domain $(0, 1)^2$ with a uniform mesh size $h = 2^{-N}$. The Laplacian Δ with homogenous Dirichlet boundary conditions is discretized by the standard 5-point stencil. The desired state $z \in \mathbb{R}^{|\Omega|}$ is defined by

$$z_{j_1, j_2} = \sin(2\pi h j_1) \sin(2\pi h j_2) \exp(2h j_1)/6,$$

for all $(j_1, j_2) \in \Omega$; see Fig. 7(a). Note that we have taken the admissible set to be universal, i.e. $U_{ad} = \mathbb{R}^{|\Omega|}$. In the following experiments, we shall fix $N = 7$, $\alpha = 10^{-4}$, $\gamma = 0.1$, and $u^0 = -\Delta z$. The associated numerical results are displayed in Fig. 7.

As shown in (b), we compute the optimal control with the prior $\psi(t) = \frac{4}{3}t^{3/4}$ and $\mu = 10^{-12}\alpha$. In fact, in the discrete setting with fixed mesh size, this result is almost identical to the optimal control with $\mu = 0$ displayed in (d). Note that the optimal control is highly sparse in the spatial domain with sparsity rate $|\{(j_1, j_2) \in \Omega : |u_{j_1, j_2}| \geq \gamma\}|/|\Omega|$ equal to 0.47 %. The corresponding realized state $(-\Delta)^{-1}u$ is given in (c), and the mean tracking error $\|(-\Delta)^{-1}u - z\|/|\Omega|$ is equal to 9.5322e-05. For comparison, we also compute the optimal control obtained from the (convex) prior $\psi(t) = t$ (together with $\mu = 0$), for which the realized state and the control are shown in (e) and (f), respectively. The corresponding sparsity rate of the control in (f) is 31.48 % and the mean tracking error is 1.0041e-04. The comparison tells that the optimal control via the concave prior can produce a better realization of the desired state even with much higher spatial sparsity. Nevertheless, we remark that the magnitudes of the nontrivial entries (whose magnitudes are larger than γ) of the control in (d) are typically much larger than those in (f), which indicates that a higher physical capability of the control devices is typically required in order to compensate a reduction on the number of the control devices.

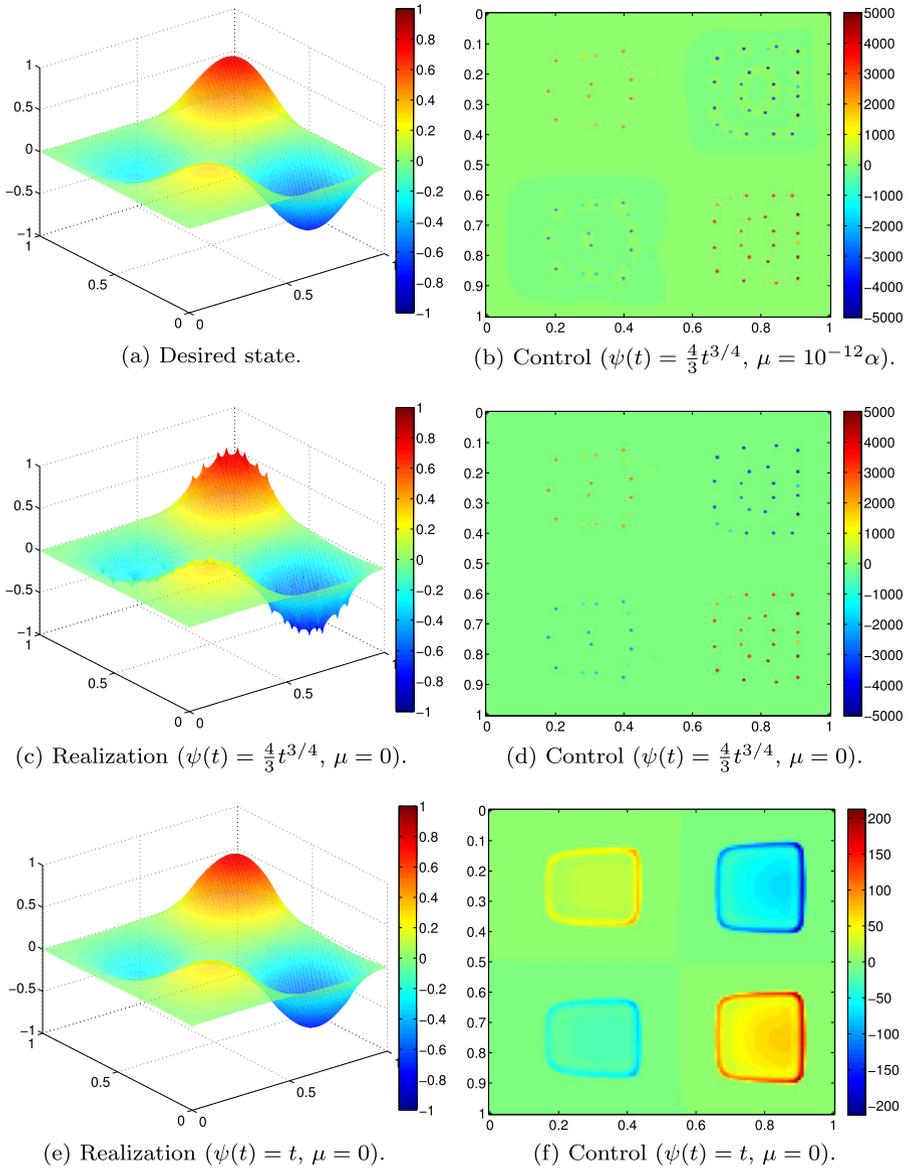


Fig. 7 Sparse optimal control

5 A smoothing scheme and the consistency result

In Sect. 3, we have proposed the R -regularized Newton algorithm for solving (2.3) which is a Huberized version of the original problem (2.1). We further witness from the numerical examples in Sect. 4 that, up to a reasonable choice of the Huber param-

eter γ , Algorithm 3.6 efficiently computes a numerical solution that is often satisfactory for practical concerns.

Nevertheless, we are intrigued by the question how to use the R -regularized Newton algorithm to track the solution to the original non-smooth (or even non-Lipschitz) problem. Motivated by the recent findings in [12], here we provide a smoothing scheme with convergence analysis to accomplish this goal. It is substantiated by the convergence of the smoothing scheme below that the Huberization strategy provides a consistent approximation of the seemingly intractable non-smooth (or even non-Lipschitz) problem.

Algorithm 5.1 (Smoothing scheme)

Require: parameters in Algorithm 3.6 and in addition $0 < \nu < 1, \eta > 0$.

- 1: Initialize the iterate (u^0, p^0) , the regularization scalar $\beta^0 \geq 0$, the trust-region radius $\sigma^0 > 0$, and the Huber parameter $\gamma^0 > 0$. Set $k := 0$.
- 2: **repeat**
- 3: Implement Steps 3–25 in Algorithm 3.6 for the γ^k -relaxed problem.
- 4: **if** $\|\text{res}(u^k, p^k; \gamma^k)\| \geq \eta\gamma^k$ **then**
- 5: Set $\gamma^{k+1} := \gamma^k$.
- 6: **else**
- 7: Set $\gamma^{k+1} := \nu\gamma^k$.
- 8: **end if**
- 9: Set $k := k + 1$.
- 10: **until** the stopping criterion for the smoothing scheme is fulfilled.

Lemma 5.2 *Let the sequence (u^k, p^k) be generated by Algorithm 5.1. Assume that (u^k) is uniformly bounded. Then we have*

$$\lim_{k \rightarrow \infty} \gamma^k = 0 \quad \text{and} \quad \liminf_{k \rightarrow \infty} \|\text{res}(u^k, p^k; \gamma^k)\| = 0.$$

Proof Define the index set

$$\mathcal{K} := \{k : \gamma^{k+1} = \nu\gamma^k\}.$$

If \mathcal{K} is finite, then there exists some \bar{k} such that for all $k > \bar{k}$ we have $\gamma^k = \gamma^{\bar{k}}$ and $\|\text{res}(u^k, p^k; \gamma^k)\| \geq \eta\gamma^{\bar{k}}$. This contradicts the global convergence in Theorem 3.7 guaranteeing that $\liminf_{k \rightarrow \infty} \|\text{res}(u^k, p^k; \gamma^k)\| = 0$. Thus, \mathcal{K} is infinite and $\lim_{k \rightarrow \infty} \gamma^k = 0$. Moreover, let $\mathcal{K} = (k^l)_{l=1}^\infty$ with $k^1 < k^2 < \dots$, then we have $\|\text{res}(u^{k^l}, p^{k^l}; \gamma^{k^l})\| \leq \rho\gamma^{k^l} \rightarrow 0$ as $l \rightarrow \infty$. Hence, $\liminf_{k \rightarrow \infty} \|\text{res}(u^k, p^k; \gamma^k)\| = 0$. □

Theorem 5.3 (Consistency) *Assume that the sequence (u^k, p^k) generated by Algorithm 5.1 be uniformly bounded. Then the sequence has an accumulation point (u^*, p^*) that satisfies the Euler–Lagrange equation (2.2).*

Proof Let the sequence (u^k, p^k) be generated by Algorithm 5.1. By Lemma 5.2, there exists a subsequence, also denoted by (u^k, p^k) , such that $\lim_{k \rightarrow \infty} \gamma^k = 0$ and

$\lim_{k \rightarrow \infty} \text{res}(u^k, p^k; \gamma^k) = 0$. By compactness, there exists yet another subsequence, again with the same notation, that converges to some $(u^*, p^*) \in \mathbb{R}^{|\Omega_u|} \times \mathbb{R}^{|\Omega_p|}$. In particular, we have

$$\nabla \Theta(u^*) + \alpha G^\top p^* = \lim_{k \rightarrow \infty} [\nabla \Theta(u^k) + \alpha G^\top p^k] = 0.$$

Moreover, since $\lim_{k \rightarrow \infty} \gamma^k = 0$ according to Lemma 5.2, on the index subset $\{j \in \Omega_p : (Gu^*)_j \neq 0\}$ we have that $\max(|Gu^k|, \gamma^k) = |Gu^k|$ for all sufficiently large k and therefore that

$$\varphi(|Gu^*|)p^* - Gu^* = \lim_{k \rightarrow \infty} [\varphi(\max(|Gu^k|, \gamma^k))p^k - Gu^k] = 0.$$

Thus we conclude that (u^*, p^*) satisfies the Euler–Lagrange equation (2.2). \square

Acknowledgements This research was supported by the Austrian Science Fund (FWF) through START project Y305 “Interfaces and Free Boundaries” and through SFB project F3204 “Mathematical Optimization and Applications in Biomedical Sciences”. The authors would like to thank T. Pock (TU Graz) for communication on the problem concerning the overcomplete dictionary. We also thank the referee for the pointer to a recent thesis work [28], which contains an extensive algorithmic study of nonconvex-regularization based sparse optimization.

References

1. ℓ_1 -magic. <http://users.ece.gatech.edu/~justin/l1magic/>
2. Attouch, H., Buttazzo, G., Michaille, G.: Variational Analysis in Sobolev and BV Spaces: Applications to PDEs and Optimization. SIAM, Philadelphia (2006)
3. Aubert, G., Kornprobst, P.: Mathematical Problems in Image Processing. Springer, New York (2002)
4. Bruckstein, A.M., Donoho, D.L., Elad, M.: From sparse solutions of systems of equations to sparse modeling of signals and images. SIAM Rev. **51**, 34–81 (2009)
5. Burke, J.V., Lewis, A.S., Overton, M.L.: A robust gradient sampling algorithm for nonsmooth, nonconvex optimization. SIAM J. Optim. **15**, 751–779 (2005)
6. Cai, J.F., Osher, S., Shen, Z.: Split Bregman methods and frame based image restoration. Multiscale Model. Simul. **8**, 337–369 (2009)
7. Candès, E.J., Tao, T.: Near optimal signal recovery from random projections: universal encoding strategies? IEEE Trans. Inf. Theory **52**, 5406–5425 (2006)
8. Chan, T.F., Mulet, P.: On the convergence of the lagged diffusivity fixed point method in total variation image restoration. SIAM J. Numer. Anal. **36**, 354–367 (1999)
9. Chapelle, O.: Training a support vector machine in the primal. Neural Comput. **19**, 1155–1178 (2007)
10. Charbonnier, P., Blanc-Féraud, L., Aubert, G., Barlaud, M.: Deterministic edge-preserving regularization in computed imaging. IEEE Trans. Image Process. **6**, 298–311 (1997)
11. Chartrand, R., Yin, W.: Iteratively reweighted algorithms for compressive sensing. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 3869–3872 (2008)
12. Chen, X.: Smoothing methods for nonsmooth, nonconvex minimization. Math. Program., Ser. B **134**, 71–99 (2012)
13. Chen, S.S., Donoho, D.L., Saunders, M.A.: Atomic decomposition by basis pursuit. SIAM Rev. **43**, 129–159 (2001)
14. Chen, X., Xu, F., Ye, Y.: Lower bound theory of nonzero entries in solutions of ℓ_2 - ℓ_p minimization. SIAM J. Sci. Comput. **32**, 2832–2852 (2010)
15. Clason, C., Kunisch, K.: A measure space approach to optimal source placement. Comput. Optim. Appl. **53**, 155–171 (2012)
16. Conn, A.R., Gould, N.I.M., Toint, P.L.: Trust-Region Methods. SIAM, Philadelphia (2000)

17. Daubechies, I., DeVore, R., Fornasier, M., Güntürk, C.: Iteratively reweighted least squares minimization for sparse recovery. *Commun. Pure Appl. Math.* **63**, 1–38 (2010)
18. Dennis, J.E. Jr., Schnabel, R.B.: *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM, Philadelphia (1996)
19. Fan, J., Li, R.: Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Am. Stat. Assoc.* **96**, 1348–1360 (2001)
20. Fischer, A.: Solution of monotone complementarity problems with locally Lipschitzian functions. *Math. Program.* **76**, 513–532 (1997)
21. Geman, D., Reynolds, G.: Constrained restoration and the recovery of discontinuities. *IEEE Trans. Pattern Anal. Mach. Intell.* **14**, 367–383 (1992)
22. Hintermüller, M., Wu, T.: Nonconvex TV^q -models in image restoration: analysis and a trust-region regularization based superlinearly convergent solver. *SIAM J. Imaging Sci.* (to appear)
23. Hintermüller, M., Ito, K., Kunisch, K.: The primal-dual active set strategy as a semismooth Newton method. *SIAM J. Optim.* **13**, 865–888 (2003)
24. Huang, J., Horowitz, J.L., Ma, S.: Asymptotic properties of bridge estimators in sparse high-dimensional regression models. *Ann. Stat.* **36**, 587–613 (2008)
25. Huber, P.J.: Robust estimation of a location parameter. *Ann. Math. Stat.* **53**, 73–101 (1964)
26. Knight, K., Fu, W.: Asymptotics for lasso-type estimators. *Ann. Stat.* **28**, 1356–1378 (2000)
27. Kunisch, K., Pock, T.: A bilevel optimization approach for parameter learning in variational models. *SIAM J. Imaging Sci.* **6**, 938–983 (2013)
28. Lin, Q.: Sparsity and nonconvex nonsmooth optimization. Ph.D. thesis, University of Washington (2009)
29. Nikolova, M.: Analysis of the recovery of edges in images and signals by minimizing nonconvex regularized least-squares. *Multiscale Model. Simul.* **4**, 960–991 (2005)
30. Nikolova, M., Chan, R.H.: The equivalence of half-quadratic minimization and the gradient linearization iteration. *IEEE Trans. Image Process.* **16**, 1623–1627 (2007)
31. Nikolova, M., Ng, M.K., Zhang, S., Ching, W.K.: Efficient reconstruction of piecewise constant images using nonsmooth nonconvex minimization. *SIAM J. Imaging Sci.* **1**, 2–25 (2008)
32. Nikolova, M., Ng, M.K., Tam, C.P.: Fast nonconvex nonsmooth minimization methods for image restoration and reconstruction. *IEEE Trans. Image Process.* **19**, 3073–3088 (2010)
33. Nocedal, J., Wright, S.: *Numerical Optimization*, 2nd edn. Springer, New York (2006)
34. Qi, L., Sun, J.: A nonsmooth version of Newton’s method. *Math. Program.* **58**, 353–367 (1993)
35. Ramlau, R., Zarzer, C.A.: On the minimization of a Tikhonov functional with a non-convex sparsity constraint. *Electron. Trans. Numer. Anal.* **39**, 476–507 (2012)
36. Stadler, G.: Elliptic optimal control problems with L^1 -control cost and applications for the placement of control devices. *Comput. Optim. Appl.* **44**, 159–181 (2009)
37. Vogel, C.R., Oman, M.E.: Iterative methods for total variation denoising. *SIAM J. Sci. Comput.* **17**, 227–238 (1996)
38. Weston, J., Mukherjee, S., Chapelle, O., Pontil, M., Poggio, T., Vapnik, V.: Feature selection for SVMs. *Adv. Neural Inf. Process. Syst.* **13**, 668–674 (2000)