

Chapter 10

First Order Algorithms in Variational Image Processing

M. Burger, A. Sawatzky, and G. Steidl

Abstract The success of non-smooth variational models in image processing is heavily based on efficient algorithms. Taking into account the specific structure of the models as sum of different convex terms, splitting algorithms are an appropriate choice. Their strength consists in the splitting of the original problem into a sequence of smaller proximal problems which are easy and fast to compute.

Operator splitting methods were first applied to linear, single-valued operators for solving partial differential equations in the 60th of the last century. More than 20 years later these methods were generalized in the convex analysis community to the solution of inclusion problems, where the linear operators have to be replaced by nonlinear, set-valued, monotone operators. Again after more than 20 years splitting methods became popular in image processing. In particular, operator splittings in combination with (augmented) Lagrangian methods and primal-dual methods have been applied very successfully.

In this chapter we give an overview of first order algorithms recently used to solve convex non-smooth variational problems in image processing. We present computational studies providing a comparison of different methods and also illustrating their success in applications.

M. Burger (✉) • A. Sawatzky
Department of Mathematics and Computer Science, University of Münster,
48149 Münster, Germany
e-mail: martin.burger@wwu.de

G. Steidl
Department of Mathematics, University of Kaiserslautern, Paul Ehrlich Str. 31,
67653 Kaiserslautern, Germany

1 Introduction

Variational methods in imaging are nowadays developing towards a quite universal and flexible tool, allowing for highly successful approaches on tasks like denoising, deblurring, inpainting, segmentation, super-resolution, disparity, and optical flow estimation. The overall structure of such approaches is of the form

$$\mathcal{D}(Ku) + \alpha\mathcal{R}(u) \rightarrow \min_u,$$

where the functional \mathcal{D} is a data fidelity term also depending on some input data f and measuring the deviation of Ku from such and \mathcal{R} is a regularization functional. Moreover K is a (often linear) forward operator modeling the dependence of data on an underlying image, and α is a positive regularization parameter. While \mathcal{D} is often smooth and (strictly) convex, the current practice almost exclusively uses nonsmooth regularization functionals. The majority of successful techniques is using nonsmooth and convex functionals like the total variation and generalizations thereof, cf. [28, 31, 41], or ℓ_1 -norms of coefficients arising from scalar products with some frame system, cf. [78] and references therein.

The efficient solution of such variational problems in imaging demands for appropriate algorithms. Taking into account the specific structure as a sum of very different terms to be minimized, splitting algorithms are a quite canonical choice. Consequently this field has revived the interest in techniques like operator splittings or augmented Lagrangians. In this chapter we shall provide an overview of methods currently developed and recent results as well as some computational studies providing a comparison of different methods and also illustrating their success in applications.

We start with a very general viewpoint in the first sections, discussing basic notations, properties of proximal maps, firmly non-expansive and averaging operators, which form the basis of further convergence arguments. Then we proceed to a discussion of several state-of-the-art algorithms and their (theoretical) convergence properties. In this chapter we focus on the so-called first order methods involving only subgradients of the functional, but no higher order derivatives. After a section discussing issues related to the use of analogous iterative schemes for ill-posed problems, we present some practical convergence studies in numerical examples related to PET and spectral CT reconstruction.

2 Notation

In the following we summarize the notations and definitions that will be used throughout the present chapter:

- $x_+ := \max\{x, 0\}$, $x \in \mathbb{R}^d$, whereby the maximum operation has to be interpreted componentwise.

- ι_C is the indicator function of a set $C \subseteq \mathbb{R}^d$ given by

$$\iota_C(x) := \begin{cases} 0 & \text{if } x \in C, \\ +\infty & \text{otherwise.} \end{cases}$$

- $\Gamma_0(\mathbb{R}^d)$ is a set of proper, convex, and lower semi-continuous functions mapping from \mathbb{R}^d into the extended real numbers $\mathbb{R} \cup \{+\infty\}$.
- $\text{dom} f := \{x \in \mathbb{R}^d : f(x) < +\infty\}$ denotes the *effective domain* of f .
- $\partial f(x_0) := \{p \in \mathbb{R}^d : f(x) - f(x_0) \geq \langle p, x - x_0 \rangle \forall x \in \mathbb{R}^d\}$ denotes the *subdifferential* of $f \in \Gamma_0(\mathbb{R}^d)$ at $x_0 \in \text{dom} f$ and is the set consisting of the *subgradients* of f at x_0 . If $f \in \Gamma_0(\mathbb{R}^d)$ is differentiable at x_0 , then $\partial f(x_0) = \{\nabla f(x_0)\}$. Conversely, if $\partial f(x_0)$ contains only one element then f is differentiable at x_0 and this element is just the gradient of f at x_0 . By *Fermat's rule*, \hat{x} is a global minimizer of $f \in \Gamma_0(\mathbb{R}^d)$ if and only if

$$0 \in \partial f(\hat{x}).$$

- $f^*(p) := \sup_{x \in \mathbb{R}^d} \{\langle p, x \rangle - f(x)\}$ is the (Fenchel) *conjugate* of f . For proper f , we have $f^* = f$ if and only if $f(x) = \frac{1}{2} \|x\|_2^2$. If $f \in \Gamma_0(\mathbb{R}^d)$ is *positively homogeneous*, i.e., $f(\alpha x) = \alpha f(x)$ for all $\alpha > 0$, then

$$f^*(x^*) = \iota_{C_f}(x^*), \quad C_f := \{x^* \in \mathbb{R}^d : \langle x^*, x \rangle \leq f(x) \forall x \in \mathbb{R}^d\}.$$

In particular, the conjugate functions of ℓ_p -norms, $p \in [1, +\infty]$, are given by

$$\|\cdot\|_p^*(x^*) = \iota_{B_q(1)}(x^*) \tag{10.1}$$

where $\frac{1}{p} + \frac{1}{q} = 1$ and as usual $p = 1$ corresponds to $q = \infty$ and conversely, and $B_q(\lambda) := \{x \in \mathbb{R}^d : \|x\|_q \leq \lambda\}$ denotes the ball of radius $\lambda > 0$ with respect to the ℓ_q -norm.

3 Proximal Operator

The algorithms proposed in this chapter to solve various problems in digital image analysis and restoration have in common that they basically reduce to the evaluation of a series of proximal problems. Therefore we start with these kind of problems. For a comprehensive overview on proximal algorithms we refer to [139].

3.1 Definition and Basic Properties

For $f \in \Gamma_0(\mathbb{R}^d)$ and $\lambda > 0$, the *proximal operator* $\text{prox}_{\lambda f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ of λf is defined by

$$\text{prox}_{\lambda f}(x) := \underset{y \in \mathbb{R}^d}{\text{argmin}} \left\{ \frac{1}{2\lambda} \|x - y\|_2^2 + f(y) \right\}. \tag{10.2}$$

It compromises between minimizing f and being near to x , where λ is the trade-off parameter between these terms. The *Moreau envelope* or *Moreau-Yoshida regularization* ${}^\lambda f : \mathbb{R}^d \rightarrow \mathbb{R}$ is given by

$${}^\lambda f(x) := \min_{y \in \mathbb{R}^d} \left\{ \frac{1}{2\lambda} \|x - y\|_2^2 + f(y) \right\}.$$

A straightforward calculation shows that ${}^\lambda f = (f^* + \frac{\lambda}{2} \|\cdot\|_2^2)^*$. The following theorem ensures that the minimizer in (10.2) exists, is unique, and can be characterized by a variational inequality. The Moreau envelope can be considered as a smooth approximation of f . For the proof we refer to [8].

Theorem 1. *Let $f \in \Gamma_0(\mathbb{R}^d)$. Then,*

- i) *For any $x \in \mathbb{R}^d$, there exists a unique minimizer $\hat{x} = \text{prox}_{\lambda f}(x)$ of (10.2).*
- ii) *The variational inequality*

$$\frac{1}{\lambda} \langle x - \hat{x}, y - \hat{x} \rangle + f(\hat{x}) - f(y) \leq 0 \quad \forall y \in \mathbb{R}^d. \tag{10.3}$$

is necessary and sufficient for \hat{x} to be the minimizer of (10.2).

- iii) *\hat{x} is a minimizer of f if and only if it is a fixed point of $\text{prox}_{\lambda f}$, i.e.,*

$$\hat{x} = \text{prox}_{\lambda f}(\hat{x}).$$

- iv) *The Moreau envelope ${}^\lambda f$ is continuously differentiable with gradient*

$$\nabla({}^\lambda f)(x) = \frac{1}{\lambda} (x - \text{prox}_{\lambda f}(x)). \tag{10.4}$$

- v) *The set of minimizers of f and ${}^\lambda f$ are the same.*

Rewriting iv) as $\text{prox}_{\lambda f}(x) = x - \lambda \nabla({}^\lambda f)(x)$ we can interpret $\text{prox}_{\lambda f}(x)$ as a gradient descent step with step size λ for minimizing ${}^\lambda f$.

Example 1. Consider the univariate function $f(y) := |y|$ and

$$\text{prox}_{\lambda f}(x) = \operatorname{argmin}_{y \in \mathbb{R}} \left\{ \frac{1}{2\lambda} (x - y)^2 + |y| \right\}.$$

Then, a straightforward computation yields that $\text{prox}_{\lambda f}$ is the *soft-shrinkage* function S_λ with threshold λ (see Figure 10.1) defined by

$$S_\lambda(x) := (x - \lambda)_+ - (-x - \lambda)_+ = \begin{cases} x - \lambda & \text{for } x > \lambda, \\ 0 & \text{for } x \in [-\lambda, \lambda], \\ x + \lambda & \text{for } x < -\lambda. \end{cases}$$

Setting $\hat{x} := S_\lambda(x) = \text{prox}_{\lambda f}(x)$, we get

$$\lambda f(x) = |\hat{x}| + \frac{1}{2\lambda}(x - \hat{x})^2 = \begin{cases} x - \frac{\lambda}{2} & \text{for } x > \lambda, \\ \frac{1}{2\lambda}x^2 & \text{for } x \in [-\lambda, \lambda], \\ -x - \frac{\lambda}{2} & \text{for } x < -\lambda. \end{cases}$$

This function λf is known as *Huber function* (see Figure 10.1).




Fig. 10.1: Left: Soft-shrinkage function $\text{prox}_{\lambda f} = S_\lambda$ for $f(y) = |y|$. Right: Moreau envelope λf .

Theorem 2 (Moreau Decomposition). For $f \in \Gamma_0(\mathbb{R}^d)$ the following decomposition holds:

$$\begin{aligned} \text{prox}_f(x) + \text{prox}_{f^*}(x) &= x, \\ {}^1f(x) + {}^1f^*(x) &= \frac{1}{2}\|x\|_2^2. \end{aligned}$$

For a proof we refer to [148, Theorem 31.5].

Remark 1 (Proximal Operator and Resolvent). The subdifferential operator is a set-valued function $\partial f : \mathbb{R}^d \rightarrow 2^{\mathbb{R}^d}$. For $f \in \Gamma_0(\mathbb{R}^d)$, we have by Fermat’s rule and subdifferential calculus that $\hat{x} = \text{prox}_{\lambda f}(x)$ if and only if

$$\begin{aligned} 0 &\in \hat{x} - x + \lambda \partial f(\hat{x}), \\ x &\in (I + \lambda \partial f)(\hat{x}), \end{aligned}$$

which implies by the uniqueness of the proximum that $\hat{x} = (I + \lambda \partial f)^{-1}(x)$. In particular, $J_{\lambda \partial f} := (I + \lambda \partial f)^{-1}$ is a single-valued operator which is called the *resolvent* of the set-valued operator $\lambda \partial f$. In summary, the proximal operator of λf coincides with the resolvent of $\lambda \partial f$, i.e.,

$$\text{prox}_{\lambda f} = J_{\lambda \partial f}.$$

The proximal operator (10.2) and the proximal algorithms described in Section 5 can be generalized by introducing a symmetric, positive definite matrix $Q \in \mathbb{R}^{d,d}$ as follows:

$$\text{prox}_{Q,\lambda f} := \underset{y \in \mathbb{R}^d}{\text{argmin}} \left\{ \frac{1}{2\lambda} \|x - y\|_Q^2 + f(y) \right\}, \quad (10.5)$$

where $\|x\|_Q^2 := x^T Q x$, see, e.g., [52, 57, 190].

3.2 Special Proximal Operators

Algorithms involving the solution of proximal problems are only efficient if the corresponding proximal operators can be evaluated in an efficient way. In the following we collect frequently appearing proximal mappings in image processing. For epigraphical projections see [12, 50, 94].

3.2.1 Orthogonal Projections

The proximal operator generalizes the orthogonal projection operator. The orthogonal projection of $x \in \mathbb{R}^d$ onto a nonempty, closed, convex set C is given by

$$\Pi_C(x) := \underset{y \in C}{\text{argmin}} \|x - y\|_2$$

and can be rewritten for any $\lambda > 0$ as

$$\Pi_C(x) = \underset{y \in \mathbb{R}^d}{\text{argmin}} \left\{ \frac{1}{2\lambda} \|x - y\|_2^2 + \iota_C(y) \right\} = \text{prox}_{\lambda \iota_C}(x).$$

Some special sets C are considered next.

Affine set

$C := \{y \in \mathbb{R}^d : Ay = b\}$ with $A \in \mathbb{R}^{m,d}$, $b \in \mathbb{R}^m$.

In case of $\|x - y\|_2 \rightarrow \min_y$ subject to $Ay = b$ we substitute $z := x - y$ which leads to

$$\|z\|_2 \rightarrow \min_z \quad \text{subject to} \quad Az = r := Ax - b.$$

This can be directly solved, see [20], and leads after back-substitution to

$$\Pi_C(x) = x - A^\dagger(Ax - b),$$

where A^\dagger denotes the Moore-Penrose inverse of A .

Halfspace

$C := \{y \in \mathbb{R}^d : a^\top y \leq b\}$ with $a \in \mathbb{R}^d, b \in \mathbb{R}$.

A straightforward computation gives

$$\Pi_C(x) = x - \frac{(a^\top x - b)_+}{\|a\|_2^2} a.$$

Box and Nonnegative Orthant

$C := \{y \in \mathbb{R}^d : l \leq y \leq u\}$ with $l, u \in \mathbb{R}^d$.

The proximal operator can be applied componentwise and gives

$$(\Pi_C(x))_k = \begin{cases} l_k & \text{if } x_k < l_k, \\ x_k & \text{if } l_k \leq x_k \leq u_k, \\ u_k & \text{if } x_k > u_k. \end{cases}$$

For $l = 0$ and $u = +\infty$ we get the orthogonal projection onto the non-negative orthant

$$\Pi_C(x) = x_+.$$

Probability Simplex

$C := \{y \in \mathbb{R}^d : \mathbf{1}^\top y = \sum_{k=1}^d y_k = 1, y \geq 0\}$.

Here we have

$$\Pi_C(x) = (x - \mu \mathbf{1})_+,$$

where $\mu \in \mathbb{R}$ has to be determined such that $h(\mu) := \mathbf{1}^\top (x - \mu \mathbf{1})_+ = 1$. Now μ can be found, e.g., by bisection with starting interval $[\max_k x_k - 1, \max_k x_k]$ or by a method similar to those described in subSection 3.2.2 for projections onto the ℓ_1 -ball. Note that h is a linear spline function with knots x_1, \dots, x_d so that μ is completely determined if we know the neighbor values x_k of μ .

3.2.2 Vector Norms

We consider the proximal operator of $f = \|\cdot\|_p$, $p \in [1, +\infty]$. By the Moreau decomposition in Theorem 2, regarding $(\lambda f)^* = \lambda f^*(\cdot/\lambda)$ and by (10.1) we obtain

$$\begin{aligned} \text{prox}_{\lambda f}(x) &= x - \text{prox}_{\lambda f^*(\cdot/\lambda)}(x) \\ &= x - \Pi_{B_q(\lambda)}(x), \end{aligned}$$

where $\frac{1}{p} + \frac{1}{q} = 1$. Thus the proximal operator can be simply computed by the projections onto the ℓ_q -ball. In particular, it follows for $p = 1, 2, \infty$:

$p = 1, q = \infty$:

For $k = 1, \dots, d$,

$$(\Pi_{B_\infty(\lambda)}(x))_k = \begin{cases} x_k & \text{if } |x_k| \leq \lambda, \\ \lambda \operatorname{sgn}(x_k) & \text{if } |x_k| > \lambda, \end{cases} \quad \text{and} \quad \operatorname{prox}_{\lambda \|\cdot\|_1}(x) = S_\lambda(x),$$

where $S_\lambda(x), x \in \mathbb{R}^d$, denotes the componentwise soft-shrinkage with threshold λ .

$p = q = 2$:

$$\Pi_{B_{2,\lambda}}(x) = \begin{cases} x & \text{if } \|x\|_2 \leq \lambda, \\ \lambda \frac{x}{\|x\|_2} & \text{if } \|x\|_2 > \lambda, \end{cases} \quad \text{and} \quad \operatorname{prox}_{\lambda \|\cdot\|_2}(x) = \begin{cases} 0 & \text{if } \|x\|_2 \leq \lambda, \\ x(1 - \frac{\lambda}{\|x\|_2}) & \text{if } \|x\|_2 > \lambda. \end{cases}$$

$p = \infty, q = 1$:

$$\Pi_{B_{1,\lambda}}(x) = \begin{cases} x & \text{if } \|x\|_1 \leq \lambda, \\ S_\mu(x) & \text{if } \|x\|_1 > \lambda, \end{cases}$$

and

$$\operatorname{prox}_{\lambda \|\cdot\|_\infty}(x) = \begin{cases} 0 & \text{if } \|x\|_1 \leq \lambda, \\ x - S_\mu(x) & \text{if } \|x\|_1 > \lambda, \end{cases}$$

with $\mu := \frac{|x_{\pi(1)}| + \dots + |x_{\pi(m)}| - \lambda}{m}$, where $|x_{\pi(1)}| \geq \dots \geq |x_{\pi(d)}| \geq 0$ are the sorted absolute values of the components of x and $m \leq d$ is the largest index such that $|x_{\pi(m)}|$ is positive and $\frac{|x_{\pi(1)}| + \dots + |x_{\pi(m)}| - \lambda}{m} \leq |x_{\pi(m)}|$, see also [62, 67]. Another method follows similar lines as the projection onto the probability simplex in the previous subsection.

Further, grouped/mixed ℓ_2 - ℓ_p -norms are defined for $x = (x_1, \dots, x_n)^\top \in \mathbb{R}^{dn}$ with $x_j := (x_{jk})_{k=1}^d \in \mathbb{R}^d, j = 1, \dots, n$ by

$$\|x\|_{2,p} := \|(\|x_j\|_2)_{j=1}^n\|_p.$$

For the ℓ_2 - ℓ_1 -norm we see that

$$\operatorname{prox}_{\lambda \|\cdot\|_{2,1}}(x) = \operatorname{argmin}_{y \in \mathbb{R}^{dn}} \left\{ \frac{1}{2\lambda} \|x - y\|_2^2 + \|y\|_{2,1} \right\}$$

can be computed separately for each j which results by the above considerations for the ℓ_2 -norm for each j in

$$\text{prox}_{\lambda \|\cdot\|_2}(x_j) = \begin{cases} 0 & \text{if } \|x_j\|_2 \leq \lambda, \\ x_j(1 - \frac{\lambda}{\|x_j\|_2}) & \text{if } \|x_j\|_2 > \lambda. \end{cases}$$

The procedure for evaluating $\text{prox}_{\lambda \|\cdot\|_{2,1}}$ is sometimes called *coupled or grouped shrinkage*.

Finally, we provide the following rule from [56, Prop. 3.6].

Lemma 1. *Let $f = g + \mu \|\cdot\|$, where $g \in \Gamma_0(\mathbb{R})$ is differentiable at 0 with $g'(0) = 0$. Then $\text{prox}_{\lambda f} = \text{prox}_{\lambda g} \circ S_{\lambda \mu}$.*

Example 2. Consider the *elastic net* regularizer $f(x) := \frac{1}{2}\|x\|_2^2 + \mu\|x\|_1$, see [192]. Setting the gradient in the proximal operator of $g := \frac{1}{2}\|\cdot\|_2^2$ to zero we obtain

$$\text{prox}_{\lambda g}(x) = \frac{1}{1 + \lambda}x.$$

The whole proximal operator of f can be then evaluated componentwise and we see by Lemma 1 that

$$\text{prox}_{\lambda f}(x) = \text{prox}_{\lambda g}(S_{\lambda \mu}(x)) = \frac{1}{1 + \lambda}S_{\lambda \mu}(x).$$

3.2.3 Matrix Norms

Next we deal with proximation problems involving matrix norms. For $X \in \mathbb{R}^{m,n}$, we are looking for

$$\text{prox}_{\lambda \|\cdot\|}(X) = \underset{Y \in \mathbb{R}^{m,n}}{\text{argmin}} \left\{ \frac{1}{2\lambda} \|X - Y\|_{\mathcal{F}}^2 + \|Y\| \right\}, \tag{10.6}$$

where $\|\cdot\|_{\mathcal{F}}$ is the Frobenius norm and $\|\cdot\|$ is any unitarily invariant matrix norm, i.e., $\|X\| = \|UXV^T\|$ for all unitary matrices $U \in \mathbb{R}^{m,m}, V \in \mathbb{R}^{n,n}$. Von Neumann (1937) [176] has characterized the unitarily invariant matrix norms as those matrix norms which can be written in the form

$$\|X\| = g(\sigma(X)),$$

where $\sigma(X)$ is the vector of singular values of X and g is a symmetric *gauge function*, see [182]. Recall that $g : \mathbb{R}^d \rightarrow \mathbb{R}_+$ is a symmetric gauge function if it is a positively homogeneous convex function which vanishes at the origin and fulfills

$$g(x) = g(\varepsilon_1 x_{k_1}, \dots, \varepsilon_k x_{k_d})$$

for all $\varepsilon_k \in \{-1, 1\}$ and all permutations k_1, \dots, k_d of indices. An analogous result was given by Davis [63] for symmetric matrices, where V^T is replaced by U^T and the singular values by the eigenvalues.

We are interested in the *Schatten- p norms* for $p = 1, 2, \infty$ which are defined for $X \in \mathbb{R}^{m,n}$ and $t := \min\{m, n\}$ by

$$\|X\|_* := \sum_{i=1}^t \sigma_i(X) = g_*(\sigma(X)) = \|\sigma(X)\|_1, \quad (\text{Nuclear norm})$$

$$\|X\|_{\mathcal{F}} := \left(\sum_{i=1}^m \sum_{j=1}^n x_{ij}^2\right)^{\frac{1}{2}} = \left(\sum_{i=1}^t \sigma_i(X)^2\right)^{\frac{1}{2}} = g_{\mathcal{F}}(\sigma(X)) = \|\sigma(X)\|_2, \quad (\text{Frobenius norm})$$

$$\|X\|_2 := \max_{i=1, \dots, t} \sigma_i(X) = g_2(\sigma(X)) = \|\sigma(X)\|_{\infty}, \quad (\text{Spectral norm}).$$

The following theorem shows that the solution of (10.6) reduces to a proximal problem for the vector norm of the singular values of X . Another proof for the special case of the nuclear norm can be found in [37].

Theorem 3. *Let $X = U \Sigma_X V^T$ be the singular value decomposition of X and $\|\cdot\|$ a unitarily invariant matrix norm. Then $\text{prox}_{\lambda, \|\cdot\|}(X)$ in (10.6) is given by $\hat{X} = U \Sigma_{\hat{X}} V^T$, where the singular values $\sigma(\hat{X})$ in $\Sigma_{\hat{X}}$ are determined by*

$$\sigma(\hat{X}) := \text{prox}_{\lambda g}(\sigma(X)) = \underset{\sigma \in \mathbb{R}^t}{\text{argmin}} \left\{ \frac{1}{2} \|\sigma(X) - \sigma\|_2^2 + \lambda g(\sigma) \right\} \quad (10.7)$$

with the symmetric gauge function g corresponding to $\|\cdot\|$.

Proof. By Fermat’s rule we know that the solution \hat{X} of (10.6) is determined by

$$0 \in \hat{X} - X + \lambda \partial \|\hat{X}\| \quad (10.8)$$

and from [182] that

$$\partial \|X\| = \text{conv}\{UDV^T : X = U \Sigma_X V^T, D = \text{diag}(d), d \in \partial g(\sigma(X))\}. \quad (10.9)$$

We now construct the unique solution \hat{X} of (10.8). Let $\hat{\sigma}$ be the unique solution of (10.7). By Fermat’s rule $\hat{\sigma}$ satisfies $0 \in \hat{\sigma} - \sigma(X) + \lambda \partial g(\hat{\sigma})$ and consequently there exists $d \in \partial g(\hat{\sigma})$ such that

$$0 = U(\text{diag}(\hat{\sigma}) - \Sigma_X + \lambda \text{diag}(d))V_F^T \Leftrightarrow 0 = U \text{diag}(\hat{\sigma}) V^T - X + \lambda U \text{diag}(d) V^T.$$

By (10.9) we see that $\hat{X} := U \text{diag}(\hat{\sigma}) V^T$ is a solution of (10.8). This completes the proof. \square

For the special matrix norms considered above, we obtain by the previous subsection

$$\begin{aligned}\|\cdot\|_* : \quad \sigma(\hat{X}) &:= \sigma(X) - \Pi_{B_{\infty,\lambda}}(\sigma(X)), \\ \|\cdot\|_{\mathcal{F}} : \quad \sigma(\hat{X}) &:= \sigma(X) - \Pi_{B_{2,\lambda}}(\sigma(X)), \\ \|\cdot\|_2 : \quad \sigma(\hat{X}) &:= \sigma(X) - \Pi_{B_{1,\lambda}}(\sigma(X)).\end{aligned}$$

4 Fixed Point Algorithms and Averaged Operators

An operator $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is *contractive* if it is Lipschitz continuous with Lipschitz constant $L < 1$, i.e., there exists a norm $\|\cdot\|$ on \mathbb{R}^d such that

$$\|Tx - Ty\| \leq L\|x - y\| \quad \forall x, y \in \mathbb{R}^d.$$

In case $L = 1$, the operator is called *nonexpansive*. A function $T : \mathbb{R}^d \supset \Omega \rightarrow \mathbb{R}^d$ is *firmly nonexpansive* if it fulfills for all $x, y \in \mathbb{R}^d$ one of the following equivalent conditions [12]:

$$\begin{aligned}\|Tx - Ty\|_2^2 &\leq \langle x - y, Tx - Ty \rangle, \\ \|Tx - Ty\|_2^2 &\leq \|x - y\|_2^2 - \|(I - T)x - (I - T)y\|_2^2.\end{aligned}\tag{10.10}$$

In particular we see that a firmly nonexpansive function is nonexpansive.

Lemma 2. *For $f \in \Gamma_0(\mathbb{R}^d)$, the proximal operator $\text{prox}_{\lambda f}$ is firmly nonexpansive. In particular the orthogonal projection onto convex sets is firmly nonexpansive.*

Proof. By Theorem 1ii) we have that

$$\frac{1}{\lambda} \langle x - \text{prox}_{\lambda f}(x), z - \text{prox}_{\lambda f}(x) \rangle + f(\text{prox}_{\lambda f}(x)) - f(z) \leq 0 \quad \forall z \in \mathbb{R}^d.$$

With $z := \text{prox}_{\lambda f}(y)$ this gives

$$\langle x - \text{prox}_{\lambda f}(x), \text{prox}_{\lambda f}(y) - \text{prox}_{\lambda f}(x) \rangle + \lambda f(\text{prox}_{\lambda f}(x)) - \lambda f(\text{prox}_{\lambda f}(y)) \leq 0$$

and similarly

$$\langle y - \text{prox}_{\lambda f}(y), \text{prox}_{\lambda f}(x) - \text{prox}_{\lambda f}(y) \rangle + \lambda f(\text{prox}_{\lambda f}(y)) - \lambda f(\text{prox}_{\lambda f}(x)) \leq 0.$$

Adding these inequalities we obtain

$$\begin{aligned}\langle x - \text{prox}_{\lambda f}(x) + \text{prox}_{\lambda f}(y) - y, \text{prox}_{\lambda f}(y) - \text{prox}_{\lambda f}(x) \rangle &\leq 0, \\ \|\text{prox}_{\lambda f}(y) - \text{prox}_{\lambda f}(x)\|_2^2 &\leq \langle y - x, \text{prox}_{\lambda f}(y) - \text{prox}_{\lambda f}(x) \rangle.\end{aligned}$$

□

The Banach fixed point theorem guarantees that a contraction has a unique fixed point and that the *Picard sequence*

$$x^{(r+1)} = Tx^{(r)}\tag{10.11}$$

converges to this fixed point for every initial element $x^{(0)}$. However, in many applications the contraction property is too restrictive in the sense that we often do not have a unique fixed point. Indeed, it is quite natural in many cases that the reached fixed point depends on the starting value $x^{(0)}$. Note that if T is continuous and $(T^r x^{(0)})_{r \in \mathbb{N}}$ is convergent, then it converges to a fixed point of T . In the following, we denote by $\text{Fix}(T)$ the *set of fixed points* of T . Unfortunately, we do not have convergence of $(T^r x^{(0)})_{r \in \mathbb{N}}$ just for nonexpansive operators as the following example shows.

Example 3. In \mathbb{R}^2 we consider the reflection operator

$$R := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Obviously, R is nonexpansive and we only have convergence of $(R^r x^{(0)})_{r \in \mathbb{N}}$ if $x^{(0)} \in \text{Fix}(R) = \text{span}\{(1, 0)^\top\}$. A possibility to obtain a ‘better’ operator is to average R , i.e., to build

$$T := \alpha I + (1 - \alpha)R = \begin{pmatrix} 1 & 0 \\ 0 & 2\alpha - 1 \end{pmatrix}, \quad \alpha \in (0, 1).$$

By

$$Tx = x \Leftrightarrow \alpha x + (1 - \alpha)R(x) = x \Leftrightarrow (1 - \alpha)R(x) = (1 - \alpha)x, \quad (10.12)$$

we see that R and T have the same fixed points. Moreover, since $2\alpha - 1 \in (-1, 1)$, the sequence $(T^r x^{(0)})_{r \in \mathbb{N}}$ converges to $(x_1^{(0)}, 0)^\top$ for every $x^{(0)} = (x_1^{(0)}, x_2^{(0)})^\top \in \mathbb{R}^2$.

An operator $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called *averaged* if there exists a nonexpansive mapping R and a constant $\alpha \in (0, 1)$ such that

$$T = \alpha I + (1 - \alpha)R.$$

Following (10.12) we see that

$$\text{Fix}(R) = \text{Fix}(T).$$

Historically, the concept of averaged mappings can be traced back to [112, 120, 156], where the name ‘averaged’ was not used yet. Results on averaged operators can also be found, e.g., in [12, 36, 55].

Lemma 3 (Averaged, (Firmly) Nonexpansive and Contractive Operators).

- i) Every averaged operator is nonexpansive.
- ii) A contractive operator $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ with Lipschitz constant $L < 1$ is averaged with respect to all parameters $\alpha \in (0, (1 - L)/2]$.
- iii) An operator is firmly nonexpansive if and only if it is averaged with $\alpha = \frac{1}{2}$.

Proof. i) Let $T = \alpha I + (1 - \alpha)R$ be averaged. Then the first assertion follows by

$$\|T(x) - T(y)\|_2 \leq \alpha \|x - y\|_2 + (1 - \alpha) \|R(x) - R(y)\|_2 \leq \|x - y\|_2.$$

ii) We define the operator $R := \frac{1}{1-\alpha}(T - \alpha I)$. It holds for all $x, y \in \mathbb{R}^d$ that

$$\begin{aligned} \|Rx - Ry\|_2 &= \frac{1}{1-\alpha} \|(T - \alpha I)x - (T - \alpha I)y\|_2, \\ &\leq \frac{1}{1-\alpha} \|Tx - Ty\|_2 + \frac{\alpha}{1-\alpha} \|x - y\|_2, \\ &\leq \frac{L+\alpha}{1-\alpha} \|x - y\|_2, \end{aligned}$$

so R is nonexpansive if $\alpha \leq (1-L)/2$.

iii) With $R := 2T - I = T - (I - T)$ we obtain the following equalities

$$\begin{aligned} \|Rx - Ry\|_2^2 &= \|Tx - Ty - ((I - T)x - (I - T)y)\|_2^2 \\ &= -\|x - y\|_2^2 + 2\|Tx - Ty\|_2^2 + 2\|(I - T)x - (I - T)y\|_2^2 \end{aligned}$$

and therefore after reordering

$$\begin{aligned} &\|x - y\|_2^2 - \|Tx - Ty\|_2^2 - \|(I - T)x - (I - T)y\|_2^2 \\ &= \|Tx - Ty\|_2^2 + \|(I - T)x - (I - T)y\|_2^2 - \|Rx - Ry\|_2^2 \\ &= \frac{1}{2} (\|x - y\|_2^2 + \|Rx - Ry\|_2^2) - \|Rx - Ry\|_2^2 \\ &= \frac{1}{2} (\|x - y\|_2^2 - \|Rx - Ry\|_2^2). \end{aligned}$$

If R is nonexpansive, then the last expression is ≥ 0 and consequently (10.10) holds true so that T is firmly nonexpansive. Conversely, if T fulfills (10.10), then

$$\frac{1}{2} (\|x - y\|_2^2 - \|Rx - Ry\|_2^2) \geq 0$$

so that R is nonexpansive. This completes the proof. \square

By the following lemma averaged operators are closed under composition.

Lemma 4 (Composition of Averaged Operators).

- i) Suppose that $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is averaged with respect to $\alpha \in (0, 1)$. Then, it is also averaged with respect to any other parameter $\tilde{\alpha} \in (0, \alpha]$.
- ii) Let $T_1, T_2 : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be averaged operators. Then, $T_2 \circ T_1$ is also averaged.

Proof. i) By assumption, $T = \alpha I + (1 - \alpha)R$ with R nonexpansive. We have

$$T = \tilde{\alpha}I + ((\alpha - \tilde{\alpha})I + (1 - \alpha)R) = \tilde{\alpha}I + (1 - \tilde{\alpha}) \underbrace{\left(\frac{\alpha - \tilde{\alpha}}{1 - \tilde{\alpha}}I + \frac{1 - \alpha}{1 - \tilde{\alpha}}R \right)}_{\tilde{R}}$$

and for all $x, y \in \mathbb{R}^d$ it holds that

$$\|\tilde{R}(x) - \tilde{R}(y)\|_2 \leq \frac{\alpha - \tilde{\alpha}}{1 - \tilde{\alpha}} \|x - y\|_2 + \frac{1 - \alpha}{1 - \tilde{\alpha}} \|R(x) - R(y)\|_2 \leq \|x - y\|_2.$$

So, \tilde{R} is nonexpansive.

ii) By assumption there exist nonexpansive operators R_1, R_2 and $\alpha_1, \alpha_2 \in (0, 1)$ such that

$$\begin{aligned} T_2(T_1(x)) &= \alpha_2 T_1(x) + (1 - \alpha_2) R_2(T_1(x)) \\ &= \alpha_2(\alpha_1 x + (1 - \alpha_1) R_1(x)) + (1 - \alpha_2) R_2(T_1(x)) \\ &= \underbrace{\alpha_2 \alpha_1}_{:=\alpha} x + (\alpha_2 - \underbrace{\alpha_2 \alpha_1}_{=\alpha}) R_1(x) + (1 - \alpha_2) R_2(T_1(x)) \\ &= \alpha x + (1 - \alpha) \underbrace{\left(\frac{\alpha_2 - \alpha}{1 - \alpha} R_1(x) + \frac{1 - \alpha_2}{1 - \alpha} R_2(T_1(x)) \right)}_{=:R} \end{aligned}$$

The concatenation of two nonexpansive operators is nonexpansive. Finally, the convex combination of two nonexpansive operators is nonexpansive so that R is indeed nonexpansive. □

An operator $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called *asymptotically regular* if it holds for all $x \in \mathbb{R}^d$ that

$$(T^{r+1}x - T^r x) \rightarrow 0 \quad \text{for } r \rightarrow +\infty.$$

Note that this property does not imply convergence, even boundedness cannot be guaranteed. As an example consider the partial sums of a harmonic sequence.

Theorem 4 (Asymptotic Regularity of Averaged Operators).

Let $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be an averaged operator with respect to the nonexpansive mapping R and the parameter $\alpha \in (0, 1)$. Assume that $\text{Fix}(T) \neq \emptyset$. Then, T is asymptotically regular.

Proof. Let $\hat{x} \in \text{Fix}(T)$ and $x^{(r)} = T^r x^{(0)}$ for some starting element $x^{(0)}$. Since T is nonexpansive, i.e., $\|x^{(r+1)} - \hat{x}\|_2 \leq \|x^{(r)} - \hat{x}\|_2$ we obtain

$$\lim_{r \rightarrow \infty} \|x^{(r)} - \hat{x}\|_2 = d \geq 0. \tag{10.13}$$

Using $\text{Fix}(T) = \text{Fix}(R)$ it follows

$$\limsup_{r \rightarrow \infty} \|R(x^{(r)}) - \hat{x}\|_2 = \limsup_{r \rightarrow \infty} \|R(x^{(r)}) - R(\hat{x})\|_2 \leq \lim_{r \rightarrow \infty} \|x^{(r)} - \hat{x}\|_2 = d. \tag{10.14}$$

Assume that $\|x^{(r+1)} - x^{(r)}\|_2 \not\rightarrow 0$ for $r \rightarrow \infty$. Then, there exists a subsequence $(x^{(r_l)})_{l \in \mathbb{N}}$ such that

$$\|x^{(r_l+1)} - x^{(r_l)}\|_2 \geq \varepsilon$$

for some $\varepsilon > 0$. By (10.13) the sequence $(x^{(r_l)})_{l \in \mathbb{N}}$ is bounded. Hence there exists a convergent subsequence $(x^{(r_{l_j})})$ such that

$$\lim_{j \rightarrow \infty} x^{(r_{l_j})} = a,$$

where $a \in S(\hat{x}, d) := \{x \in \mathbb{R}^d : \|x - \hat{x}\|_2 = d\}$ by (10.13). On the other hand, we have by the continuity of R and (10.14) that

$$\lim_{j \rightarrow \infty} R(x^{(r_{l_j})}) = b, \quad b \in B(\hat{x}, d).$$

Since $\varepsilon \leq \|x^{(r_{l_j+1})} - x^{(r_{l_j})}\|_2 = \|(\alpha - 1)x^{(r_{l_j})} + (1 - \alpha)R(x^{(r_{l_j})})\|_2$ we conclude by taking the limit $j \rightarrow \infty$ that $a \neq b$. By the continuity of T and (10.13) we obtain

$$\lim_{j \rightarrow \infty} T(x^{(r_{l_j})}) = c, \quad c \in S(\hat{x}, d).$$

However, by the strict convexity of $\|\cdot\|_2^2$ this yields the contradiction

$$\begin{aligned} \|c - \hat{x}\|_2^2 &= \lim_{j \rightarrow \infty} \|T(x^{(r_{l_j})}) - \hat{x}\|_2^2 = \lim_{j \rightarrow \infty} \|\alpha(x^{(r_{l_j})} - \hat{x}) + (1 - \alpha)(R(x^{(r_{l_j})}) - \hat{x})\|_2^2 \\ &= \|\alpha(a - \hat{x}) + (1 - \alpha)(b - \hat{x})\|_2^2 < \alpha\|a - \hat{x}\|_2^2 + (1 - \alpha)\|b - \hat{x}\|_2^2 \\ &\leq d^2. \end{aligned}$$

□

The following theorem was first proved for operators on Hilbert spaces by Opial [133, Theorem 1] based on the results in [29], where convergence must be replaced by weak convergence in general Hilbert spaces. A shorter proof can be found in the appendix of [61]. For finite dimensional spaces the proof simplifies as follows.

Theorem 5 (Opial’s Convergence Theorem).

Let $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ fulfill the following conditions: $\text{Fix}(T) \neq \emptyset$, T is nonexpansive and asymptotically regular. Then, for every $x^{(0)} \in \mathbb{R}^d$, the sequence of Picard iterates $(x^{(r)})_{r \in \mathbb{N}}$ generated by $x^{(r+1)} = Tx^{(r)}$ converges to an element of $\text{Fix}(T)$.

Proof. Since T is nonexpansive, we have for any $\hat{x} \in \text{Fix}(T)$ and any $x^{(0)} \in \mathbb{R}^d$ that

$$\|T^{r+1}x^{(0)} - \hat{x}\|_2 \leq \|T^r x^{(0)} - \hat{x}\|_2.$$

Hence $(T^r x^{(0)})_{r \in \mathbb{N}}$ is bounded and there exists a subsequence $(T^{r_l} x^{(0)})_{l \in \mathbb{N}} = (x^{(r_l)})_{l \in \mathbb{N}}$ which converges to some \tilde{x} . If we can show that $\tilde{x} \in \text{Fix}(T)$ we are done because in this case

$$\|T^r x^{(0)} - \tilde{x}\|_2 \leq \|T^{r_l} x^{(0)} - \tilde{x}\|_2, \quad r \geq r_l$$

and thus the whole sequence converges to \tilde{x} .

Since T is asymptotically regular it follows that

$$(T - I)(T^{r_l}x^{(0)}) = T^{r_l+1}x^{(0)} - T^{r_l}x^{(0)} \rightarrow 0$$

and since $(T^{r_l}x^{(0)})_{l \in \mathbb{N}}$ converges to \tilde{x} and T is continuous we get that $(T - I)(\tilde{x}) = 0$, i.e., $\tilde{x} \in \text{Fix}(T)$. □

Combining the above Theorems 4 and 5 we obtain the following main result.

Theorem 6 (Convergence of Averaged Operator Iterations). *Let $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be an averaged operator such that $\text{Fix}(T) \neq \emptyset$. Then, for every $x^{(0)} \in \mathbb{R}^d$, the sequence $(T^r x^{(0)})_{r \in \mathbb{N}}$ converges to a fixed point of T .*

5 Proximal Algorithms

5.1 Proximal Point Algorithm

By Theorem 1 iii) the minimizer of a function $f \in \Gamma_0(\mathbb{R}^d)$, which we suppose to exist, is characterized by the fixed point equation

$$\hat{x} = \text{prox}_{\lambda f}(\hat{x}).$$

The corresponding Picard iteration gives rise to the following *proximal point algorithm* which dates back to [121, 147]. Since $\text{prox}_{\lambda f}$ is firmly nonexpansive by Lemma 2 and thus averaged, the algorithm converges by Theorem 6 for any initial value $x^{(0)} \in \mathbb{R}^d$ to a minimizer of f if there exists one.

Algorithm 1 Proximal Point Algorithm (PPA)

Initialization: $x^{(0)} \in \mathbb{R}^d, \lambda > 0$

Iterations: For $r = 0, 1, \dots$

$$x^{(r+1)} = \text{prox}_{\lambda f}(x^{(r)}) = \underset{x \in \mathbb{R}^d}{\text{argmin}} \left\{ \frac{1}{2\lambda} \|x^{(r)} - x\|_2^2 + f(x) \right\}$$

The PPA can be generalized for the sum $\sum_{i=1}^n f_i$ of functions $f_i \in \Gamma_0(\mathbb{R}^d), i = 1, \dots, n$. Popular generalizations are the so-called cyclic PPA [18] and the parallel PPA [53].

5.2 Proximal Gradient Algorithm

We are interested in minimizing functions of the form $f = g + h$, where $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex, differentiable with Lipschitz continuous gradient and Lipschitz constant L , i.e.,

$$\|\nabla g(x) - \nabla g(y)\|_2 \leq L\|x - y\|_2 \quad \forall x, y \in \mathbb{R}^d, \tag{10.15}$$

and $h \in \Gamma_0(\mathbb{R}^d)$. Note that the Lipschitz condition on ∇g implies

$$g(x) \leq g(y) + \langle \nabla g(y), x - y \rangle + \frac{L}{2} \|x - y\|_2^2 \quad \forall x, y \in \mathbb{R}^d, \quad (10.16)$$

see, e.g., [134]. We want to solve

$$\operatorname{argmin}_{x \in \mathbb{R}^d} \{g(x) + h(x)\}. \quad (10.17)$$

By Fermat's rule and subdifferential calculus we know that \hat{x} is a minimizer of (10.17) if and only if

$$\begin{aligned} 0 &\in \nabla g(\hat{x}) + \partial h(\hat{x}), \\ \hat{x} - \eta \nabla g(\hat{x}) &\in \hat{x} + \eta \partial h(\hat{x}), \\ \hat{x} &= (I + \eta \partial h)^{-1} (\hat{x} - \eta \nabla g(\hat{x})) = \operatorname{prox}_{\eta h} (\hat{x} - \eta \nabla g(\hat{x})). \end{aligned} \quad (10.18)$$

This is a fixed point equation for the minimizer \hat{x} of f . The corresponding Picard iteration is known as *proximal gradient algorithm* or as *proximal forward-backward splitting*.

Algorithm 2 Proximal Gradient Algorithm (FBS)

Initialization: $x^{(0)} \in \mathbb{R}^d$, $\eta \in (0, 2/L)$

Iterations: For $r = 0, 1, \dots$

$$x^{(r+1)} = \operatorname{prox}_{\eta h} (x^{(r)} - \eta \nabla g(x^{(r)}))$$

In the special case when $h := \iota_C$ is the indicator function of a nonempty, closed, convex set $C \subset \mathbb{R}^d$, the above algorithm for finding

$$\operatorname{argmin}_{x \in C} g(x)$$

becomes the gradient descent re-projection algorithm, also known as “gradient projection algorithm”.

Algorithm 3 Gradient Descent Re-Projection Algorithm

Initialization: $x^{(0)} \in \mathbb{R}^d$, $\eta \in (0, 2/L)$

Iterations: For $r = 0, 1, \dots$

$$x^{(r+1)} = \Pi_C (x^{(r)} - \eta \nabla g(x^{(r)}))$$

It is also possible to use flexible variables $\eta_r \in (0, \frac{2}{L})$ in the proximal gradient algorithm. For further details, modifications, and extensions see also [72, Chapter 12]. The convergence of the algorithm follows by the next theorem.

Theorem 7 (Convergence of Proximal Gradient Algorithm). *Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a convex, differentiable function on \mathbb{R}^d with Lipschitz continuous gradient and Lipschitz constant L and $h \in \Gamma_0(\mathbb{R}^d)$. Suppose that a solution of (10.17) exists. Then, for every initial point $x^{(0)}$ and $\eta \in (0, \frac{2}{L})$, the sequence $\{x^{(r)}\}_r$ generated by the proximal gradient algorithm converges to a solution of (10.17).*

Proof. We show that $\text{prox}_{\eta h}(I - \eta \nabla g)$ is averaged. Then we are done by Theorem 6. By Lemma 2 we know that $\text{prox}_{\eta h}$ is firmly nonexpansive. By the Baillon-Haddad Theorem [12, Corollary 16.1] the function $\frac{1}{L}\nabla g$ is also firmly nonexpansive, i.e., it is averaged with parameter $\frac{1}{2}$. This means that there exists a nonexpansive mapping R such that $\frac{1}{L}\nabla g = \frac{1}{2}(I + R)$ which implies

$$I - \eta \nabla g = I - \frac{\eta L}{2}(I + R) = (1 - \frac{\eta L}{2})I + \frac{\eta L}{2}(-R).$$

Thus, for $\eta \in (0, \frac{2}{L})$, the operator $I - \eta \nabla g$ is averaged. Since the concatenation of two averaged operators is averaged again we obtain the assertion. \square

Under the above conditions a linear convergence rate can be achieved in the sense that

$$f(x^{(r)}) - f(\hat{x}) = \mathcal{O}(1/r),$$

see, e.g., [49, 125]¹.

Example 4. For solving

$$\underset{x \in \mathbb{R}^d}{\text{argmin}} \left\{ \underbrace{\frac{1}{2} \|Kx - b\|_2^2}_g + \underbrace{\lambda \|x\|_1}_h \right\}$$

we compute $\nabla g(x) = K^T(Kx - b)$ and use that the proximal operator of the ℓ_1 -norm is just the componentwise soft-shrinkage. Then the proximal gradient algorithm becomes

$$x^{(r+1)} = \text{prox}_{\lambda \eta \|\cdot\|_1} \left(x^{(r)} - \eta K^T(Kx^{(r)} - b) \right) = S_{\eta \lambda} \left(x^{(r)} - \eta K^T(Kx^{(r)} - b) \right).$$

This algorithm is known as *iterative soft-thresholding algorithm* (ISTA) and was developed and analyzed through various techniques by many researchers. For a general Hilbert space approach of ISTA, see, e.g., [61].

The FBS algorithm has been recently extended to the case of non-convex functions in [6, 7, 22, 52, 132]. The convergence analysis mainly rely on the assumption that the objective function $f = g + h$ satisfies the Kurdyka-Lojasiewicz inequality which is indeed fulfilled for a wide class of functions as log – exp, semi-algebraic, and subanalytic functions which are of interest in image processing.

¹ There exist different notations for the convergence rate of algorithms in the literature. Sometimes the notation of this chapter is also called “superlinear convergence” while $\|\hat{x} - x^{(r)}\| \leq C\delta^r$, $\delta < 1$ is used for the linear convergence. But if $C = C(r) \rightarrow 0$ as $r \rightarrow +\infty$ in the last formula, this could be also meant by “superlinear convergence”.

5.3 Accelerated Algorithms

For large scale problems as those arising in image processing a major concern is to find efficient algorithms solving the problem in a reasonable time. While each FBS step has low computational complexity, it may suffer from slow linear convergence [49]. Using a simple extrapolation idea with appropriate parameters τ_r , the convergence can often be accelerated:

$$\begin{aligned} y^{(r)} &= x^{(r)} + \tau_r (x^{(r)} - x^{(r-1)}), \\ x^{(r+1)} &= \text{prox}_{\eta h} (y^{(r)} - \eta \nabla g(y^{(r)})). \end{aligned} \quad (10.19)$$

By the next Theorem 8 we will see that $\tau_r = \frac{r-1}{r+2}$ appears to be a good choice. Clearly, we can vary η in each step again. Choosing θ_r such that $\tau_r = \frac{\theta_r(1-\theta_{r-1})}{\theta_{r-1}}$, e.g., $\theta_r = \frac{2}{r+2}$ for the above choice of τ_r , the algorithm can be rewritten as follows:

Algorithm 4 Fast Proximal Gradient Algorithm

Initialization: $x^{(0)} = z^{(0)} \in \mathbb{R}^d$, $\eta \in (0, 1/L)$, $\theta_r = \frac{2}{r+2}$

Iterations: For $r = 0, 1, \dots$

$$y^{(r)} = (1 - \theta_r)x^{(r)} + \theta_r z^{(r)}$$

$$x^{(r+1)} = \text{prox}_{\eta h} (y^{(r)} - \eta \nabla g(y^{(r)}))$$

$$z^{(r+1)} = x^{(r)} + \frac{1}{\theta_r} (x^{(r+1)} - x^{(r)})$$

By the following standard theorem the extrapolation modification of the FBS algorithm ensures a quadratic convergence rate see also [125, 171].

Theorem 8. *Let $f = g + h$, where $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex, Lipschitz differentiable function with Lipschitz constant L and $h \in \Gamma_0(\mathbb{R}^d)$. Assume that f has a minimizer \hat{x} . Then the fast proximal gradient algorithm fulfills*

$$f(x^{(r)}) - f(\hat{x}) = \mathcal{O}(1/r^2).$$

Proof. First we consider the progress in one step of the algorithm. By the Lipschitz differentiability of g in (10.16) and since $\eta < \frac{1}{L}$ we know that

$$g(x^{(r+1)}) \leq g(y^{(r)}) + \langle \nabla g(y^{(r)}), x^{(r+1)} - y^{(r)} \rangle + \frac{1}{2\eta} \|x^{(r+1)} - y^{(r)}\|_2^2 \quad (10.20)$$

and by the variational characterization of the proximal operator in Theorem 1ii) for all $u \in \mathbb{R}^d$ that

$$\begin{aligned}
 h(x^{(r+1)}) &\leq h(u) + \frac{1}{\eta} \langle y^{(r)} - \eta \nabla g(y^{(r)}) - x^{(r+1)}, x^{(r+1)} - u \rangle \\
 &\leq h(u) - \langle \nabla g(y^{(r)}), x^{(r+1)} - u \rangle + \frac{1}{\eta} \langle y^{(r)} - x^{(r+1)}, x^{(r+1)} - u \rangle. \quad (10.21)
 \end{aligned}$$

Adding the main inequalities (10.20) and (10.21) and using the convexity of g yields

$$\begin{aligned}
 f(x^{(r+1)}) &\leq f(u) + \underbrace{-g(u) + g(y^{(r)}) + \langle \nabla g(y^{(r)}), u - y^{(r)} \rangle}_{\leq 0} \\
 &\quad + \frac{1}{2\eta} \|x^{(r+1)} - y^{(r)}\|_2^2 + \frac{1}{\eta} \langle y^{(r)} - x^{(r+1)}, x^{(r+1)} - u \rangle \\
 &\leq f(u) + \frac{1}{2\eta} \|x^{(r+1)} - y^{(r)}\|_2^2 + \frac{1}{\eta} \langle y^{(r)} - x^{(r+1)}, x^{(r+1)} - u \rangle.
 \end{aligned}$$

Combining these inequalities for $u := \hat{x}$ and $u := x^{(r)}$ with $\theta_r \in [0, 1]$ gives

$$\begin{aligned}
 &\theta_r \left(f(x^{(r+1)}) - f(\hat{x}) \right) + (1 - \theta_r) \left(f(x^{(r+1)}) - f(x^{(r)}) \right) \\
 &= f(x^{(r+1)}) - f(\hat{x}) + (1 - \theta_r) \left(f(\hat{x}) - f(x^{(r)}) \right) \\
 &\leq \frac{1}{2\eta} \|x^{(r+1)} - y^{(r)}\|_2^2 + \frac{1}{\eta} \langle y^{(r)} - x^{(r+1)}, x^{(r+1)} - \theta_r \hat{x} - (1 - \theta_r) x^{(r)} \rangle \\
 &= \frac{1}{2\eta} \left(\|y^{(r)} - \theta_r \hat{x} - (1 - \theta_r) x^{(r)}\|_2^2 - \|x^{(r+1)} - \theta_r \hat{x} - (1 - \theta_r) x^{(r)}\|_2^2 \right) \\
 &= \frac{\theta_r^2}{2\eta} \left(\|z^{(r)} - \hat{x}\|_2^2 - \|z^{(r+1)} - \hat{x}\|_2^2 \right).
 \end{aligned}$$

Thus, we obtain for a single step

$$\frac{\eta}{\theta_r^2} \left(f(x^{(r+1)}) - f(\hat{x}) \right) + \frac{1}{2} \|z^{(r+1)} - \hat{x}\|_2^2 \leq \frac{\eta(1 - \theta_r)}{\theta_r^2} \left(f(x^{(r)}) - f(\hat{x}) \right) + \frac{1}{2} \|z^{(r)} - \hat{x}\|_2^2.$$

Using the relation recursively on the right-hand side and regarding that $\frac{(1 - \theta_r)}{\theta_r^2} \leq \frac{1}{\theta_{r-1}^2}$ we obtain

$$\frac{\eta}{\theta_r^2} \left(f(x^{(r+1)}) - f(\hat{x}) \right) \leq \frac{\eta(1 - \theta_0)}{\theta_0^2} \left(f(x^{(0)}) - f(\hat{x}) \right) + \frac{1}{2} \|z^{(0)} - \hat{x}\|_2^2 = \frac{1}{2} \|x^{(0)} - \hat{x}\|_2^2.$$

This yields the assertion

$$f(x^{(r+1)}) - f(\hat{x}) \leq \frac{2}{\eta(r+2)^2} \|x^{(0)} - \hat{x}\|_2^2.$$

□

There exist many variants or generalizations of the above algorithm (with certain convergence rates under special assumptions):

- *Nesterov's algorithms* [126, 128], see also [60, 171]; this includes approximation algorithms for nonsmooth g [14, 129] as NESTA,
- *fast iterative shrinkage algorithms* (FISTA) by Beck and Teboulle [13],
- *variable metric strategies* [24, 33, 57, 138], where based on (10.5) step (10.19) is replaced by

$$x^{(r+1)} = \text{prox}_{Q_r, \eta_r h} \left(y^{(r)} - \eta_r Q_r^{-1} \nabla g(y^{(r)}) \right) \quad (10.22)$$

with symmetric, positive definite matrices Q_r .

Line search strategies can be incorporated [89, 93, 127]. Finally we mention Barzilei-Borwein step size rules [11] based on a Quasi-Newton approach and relatives, see [79] for an overview and the cyclic proximal gradient algorithm related to the cyclic Richardson algorithm [165].

6 Primal-Dual Methods

6.1 Basic Relations

The following minimization algorithms closely rely on the primal-dual formulation of problems. We consider functions $f = g + h(A \cdot)$, where $g \in \Gamma_0(\mathbb{R}^d)$, $h \in \Gamma_0(\mathbb{R}^m)$, and $A \in \mathbb{R}^{m,d}$, and ask for the solution of the primal problem

$$(P) \quad \underset{x \in \mathbb{R}^d}{\text{argmin}} \{g(x) + h(Ax)\}, \quad (10.23)$$

that can be rewritten as

$$(P) \quad \underset{x \in \mathbb{R}^d, y \in \mathbb{R}^m}{\text{argmin}} \{g(x) + h(y) \quad \text{s.t.} \quad Ax = y\}. \quad (10.24)$$

The *Lagrangian* of (10.24) is given by

$$L(x, y, p) := g(x) + h(y) + \langle p, Ax - y \rangle \quad (10.25)$$

and the *augmented Lagrangian* by

$$\begin{aligned} L_\gamma(x, y, p) &:= g(x) + h(y) + \langle p, Ax - y \rangle + \frac{\gamma}{2} \|Ax - y\|_2^2, \quad \gamma > 0, \\ &= g(x) + h(y) + \frac{\gamma}{2} \|Ax - y + \frac{p}{\gamma}\|_2^2 - \frac{1}{2\gamma} \|p\|_2^2. \end{aligned} \quad (10.26)$$

Based on the Lagrangian (10.25), the primal and dual problem can be written as

$$(P) \quad \underset{x \in \mathbb{R}^d, y \in \mathbb{R}^m}{\text{argmin}} \sup_{p \in \mathbb{R}^m} \{g(x) + h(y) + \langle p, Ax - y \rangle\}, \quad (10.27)$$

$$(D) \quad \underset{p \in \mathbb{R}^m}{\text{argmax}} \inf_{x \in \mathbb{R}^d, y \in \mathbb{R}^m} \{g(x) + h(y) + \langle p, Ax - y \rangle\}. \quad (10.28)$$

Since

$$\min_{y \in \mathbb{R}^m} \{h(y) - \langle p, y \rangle\} = -\max_{y \in \mathbb{R}^m} \{\langle p, y \rangle - h(y)\} = -h^*(p)$$

and in (10.23) further

$$h(Ax) = \max_{p \in \mathbb{R}^m} \{\langle p, Ax \rangle - h^*(p)\},$$

the primal and dual problem can be rewritten as

$$(P) \quad \operatorname{argmin}_{x \in \mathbb{R}^d} \sup_{p \in \mathbb{R}^m} \{g(x) - h^*(p) + \langle p, Ax \rangle\},$$

$$(D) \quad \operatorname{argmax}_{p \in \mathbb{R}^m} \inf_{x \in \mathbb{R}^d} \{g(x) - h^*(p) + \langle p, Ax \rangle\}.$$

If the infimum exists, the dual problem can be seen as Fenchel dual problem

$$(D) \quad \operatorname{argmax}_{p \in \mathbb{R}^m} \{-g^*(-A^T p) - h^*(p)\}. \quad (10.29)$$

Recall that $((\hat{x}, \hat{y}), \hat{p}) \in \mathbb{R}^{dm,m}$ is a *saddle point* of the Lagrangian L in (10.25) if

$$L((\hat{x}, \hat{y}), p) \leq L((\hat{x}, \hat{y}), \hat{p}) \leq L(x, y, \hat{p}) \quad \forall (x, y) \in \mathbb{R}^{dm}, p \in \mathbb{R}^m.$$

If $((\hat{x}, \hat{y}), \hat{p}) \in \mathbb{R}^{dm,m}$ is a saddle point of L , then (\hat{x}, \hat{y}) is a solution of the primal problem (10.27) and \hat{p} is a solution of the dual problem (10.28). The converse is also true. However the existence of a solution of the primal problem $(\hat{x}, \hat{y}) \in \mathbb{R}^{dm}$ does only imply under additional qualification constraint that there exists \hat{p} such that $((\hat{x}, \hat{y}), \hat{p}) \in \mathbb{R}^{dm,m}$ is a saddle point of L .

6.2 Alternating Direction Method of Multipliers

Based on the Lagrangian formulation (10.27) and (10.28), a first idea to solve the optimization problem would be to alternate the minimization of the Lagrangian with respect to (x, y) and to apply a gradient ascent approach with respect to p . This is known as general Uzawa method [5]. More precisely, noting that for differentiable $v(p) := \inf_{x,y} L(x, y, p) = L(\tilde{x}, \tilde{y}, p)$ we have $\nabla v(p) = A\tilde{x} - \tilde{y}$, the algorithm reads

$$(x^{(r+1)}, y^{(r+1)}) \in \operatorname{argmin}_{x \in \mathbb{R}^d, y \in \mathbb{R}^m} L(x, y, p^{(r)}), \quad (10.30)$$

$$p^{(r+1)} = p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)}), \quad \gamma > 0.$$

Linear convergence can be proved under certain conditions (strong convexity of f) [87]. The assumptions on f to ensure convergence of the algorithm can be relaxed by replacing the Lagrangian by the augmented Lagrangian L_γ (10.26) with fixed parameter γ :

$$(x^{(r+1)}, y^{(r+1)}) \in \underset{x \in \mathbb{R}^d, y \in \mathbb{R}^m}{\operatorname{argmin}} L_\gamma(x, y, p^{(r)}), \tag{10.31}$$

$$p^{(r+1)} = p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)}), \quad \gamma > 0.$$

This augmented Lagrangian method is known as *method of multipliers* [101, 141, 147]. It can be shown [35, Theorem 3.4.7], [17] that the sequence $(p^{(r)})_r$ generated by the algorithm coincides with the proximal point algorithm applied to $-v(p)$, i.e.,

$$p^{(r+1)} = \operatorname{prox}_{-\gamma v} \left(p^{(r)} \right).$$

The improved convergence properties came at a cost. While the minimization with respect to x and y can be separately computed in (10.30) using $\langle p^{(r)}, (A| - I) \begin{pmatrix} x \\ y \end{pmatrix} \rangle = \langle \begin{pmatrix} A^T \\ -I \end{pmatrix} p^{(r)}, \begin{pmatrix} x \\ y \end{pmatrix} \rangle$, this is no longer possible for the augmented Lagrangian. A remedy is to alternate the minimization with respect to x and y which leads to

$$x^{(r+1)} \in \underset{x \in \mathbb{R}^d}{\operatorname{argmin}} L_\gamma(x, y^{(r)}, p^{(r)}), \tag{10.32}$$

$$y^{(r+1)} = \underset{y \in \mathbb{R}^m}{\operatorname{argmin}} L_\gamma(x^{(r+1)}, y, p^{(r)}), \tag{10.33}$$

$$\tilde{p}^{(r+1)} = p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)}).$$

This is the *alternating direction method of multipliers (ADMM)* which dates back to [82, 83, 88].

Algorithm 5 Alternating Direction Method of Multipliers (ADMM)

Initialization: $y^{(0)} \in \mathbb{R}^m, p^{(0)} \in \mathbb{R}^m$

Iterations: For $r = 0, 1, \dots$

$$x^{(r+1)} \in \underset{x \in \mathbb{R}^d}{\operatorname{argmin}} \left\{ g(x) + \frac{\gamma}{2} \left\| \frac{1}{\gamma} p^{(r)} + Ax - y^{(r)} \right\|_2^2 \right\}$$

$$y^{(r+1)} = \underset{y \in \mathbb{R}^m}{\operatorname{argmin}} \left\{ h(y) + \frac{\gamma}{2} \left\| \frac{1}{\gamma} p^{(r)} + Ax^{(r+1)} - y \right\|_2^2 \right\} = \operatorname{prox}_{\frac{1}{\gamma} h} \left(\frac{1}{\gamma} p^{(r)} + Ax^{(r+1)} \right)$$

$$p^{(r+1)} = p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)})$$

Setting $b^{(r)} := p^{(r)}/\gamma$ we obtain the following (scaled) ADMM:

A good overview on the ADMM algorithm and its applications is given in [27], where in particular the important issue of choosing the parameter $\gamma > 0$ is addressed. Convergence of the ADMM under various assumptions was proved, e.g., in [83, 96, 116, 170]. The ADMM can be considered for more general problems

$$\operatorname{argmin}_{x_i \in \mathbb{R}^{d_i}} \left\{ \sum_{i=1}^m g_i(x) \quad \text{s.t.} \quad A_i x_i = c \right\}. \tag{10.34}$$

Algorithm 6 Alternating Direction Method of Multipliers (scaled ADMM)

Initialization: $y^{(0)} \in \mathbb{R}^m, b^{(0)} \in \mathbb{R}^m$

Iterations: For $r = 0, 1, \dots$

$$\begin{aligned}
 x^{(r+1)} &\in \operatorname{argmin}_{x \in \mathbb{R}^d} \left\{ g(x) + \frac{\gamma}{2} \|b^{(r)} + Ax - y^{(r)}\|_2^2 \right\} \\
 y^{(r+1)} &= \operatorname{argmin}_{y \in \mathbb{R}^m} \left\{ h(y) + \frac{\gamma}{2} \|b^{(r)} + Ax^{(r+1)} - y\|_2^2 \right\} = \operatorname{prox}_{\frac{1}{\gamma}h}(b^{(r)} + Ax^{(r+1)}) \\
 b^{(r+1)} &= b^{(r)} + Ax^{(r+1)} - y^{(r+1)}
 \end{aligned}$$

Here we refer to [47] and the references therein. We will see that for our problem (10.24) the convergence follows by the relation of the ADMM to the so-called Douglas-Rachford splitting algorithm where convergence can be shown using averaged operators. Few bounds on the global convergence rate of the algorithm can be found in [68] (linear convergence for linear programs depending on a variety of quantities), [102] (linear convergence for sufficiently small step size) and on the local behavior of a specific variation of the ADMM during the course of iteration for quadratic programs in [21]. Further global rates of the ADMM are given in the recent preprints [64, 65].

Theorem 9 (Convergence of ADMM). *Let $g \in \Gamma_0(\mathbb{R}^d)$, $h \in \Gamma_0(\mathbb{R}^m)$ and $A \in \mathbb{R}^{m,d}$. Assume that the Lagrangian (10.25) has a saddle point. Then, for $r \rightarrow \infty$, the sequence $\gamma(b^{(r)})_r$ converges to a solution of the dual problem. If in addition the first step (10.32) in the ADMM algorithm has a unique solution, then $(x^{(r)})_r$ converges to a solution of the primal problem.*

There exist different modifications of the ADMM algorithm presented above:

- *inexact computation* of the first step (10.32) [48, 69] such that it might be handled by an iterative method,
- *variable parameter and metric strategies* [27, 95, 96, 98, 111] where the fixed parameter γ can vary in each step, or the quadratic term $(\gamma/2)\|Ax - y\|_2^2$ within the augmented Lagrangian (10.26) is replaced by the more general proximal operator based on (10.5) such that the ADMM updates (10.32) and (10.33) receive the form

$$\begin{aligned}
 x^{(r+1)} &\in \operatorname{argmin}_{x \in \mathbb{R}^d} \left\{ g(x) + \frac{1}{2} \|b^{(r)} + Ax - y^{(r)}\|_{Q_r}^2 \right\}, \\
 y^{(r+1)} &= \operatorname{argmin}_{y \in \mathbb{R}^m} \left\{ h(y) + \frac{1}{2} \|b^{(r)} + Ax^{(r+1)} - y\|_{Q_r}^2 \right\},
 \end{aligned}$$

respectively, with symmetric, positive definite matrices Q_r . The variable parameter strategies might mitigate the performance dependency on the initial chosen fixed parameter [27, 98, 111, 181] and include monotone conditions [96, 111] or more flexible non-monotone rules [27, 95, 98].

ADMM from the Perspective of Variational Inequalities

The ADMM algorithm presented above from the perspective of Lagrangian functions has been also studied extensively in the area of *variational inequalities (VIs)*, see, e.g., [82, 95, 170]. A VI problem consists of finding for a mapping $F : \mathbb{R}^l \rightarrow \mathbb{R}^l$ a vector $\hat{z} \in \mathbb{R}^l$ such that

$$\langle z - \hat{z}, F(\hat{z}) \rangle \geq 0, \quad \forall z \in \mathbb{R}^l. \tag{10.35}$$

In the following, we consider the minimization problem (10.24), i.e.,

$$\underset{x \in \mathbb{R}^d, y \in \mathbb{R}^m}{\operatorname{argmin}} \{g(x) + h(y) \quad \text{s.t.} \quad Ax = y\},$$

for $g \in \Gamma_0(\mathbb{R}^d)$, $h \in \Gamma_0(\mathbb{R}^m)$. The discussion can be extended to the more general problem (10.34) [95, 170]. Considering the Lagrangian (10.25) and its optimality conditions, solving (10.24) is equivalent to find a triple $\hat{z} = ((\hat{x}, \hat{y}), \hat{p}) \in \mathbb{R}^{d+m,m}$ such that (10.35) holds with

$$z = \begin{pmatrix} x \\ y \\ p \end{pmatrix}, \quad F(z) = \begin{pmatrix} \partial g(x) + A^\top p \\ \partial h(y) - p \\ Ax - y \end{pmatrix},$$

where ∂g and ∂h have to be understood as any element of the corresponding subdifferential for simplicity. Note that ∂g and ∂h are maximal monotone operators [12]. A VI problem of this form can be solved by ADMM as proposed by Gabay [82] and Gabay and Mercier [83]: for a given triple $(x^{(r)}, y^{(r)}, p^{(r)})$ generate new iterates $(x^{(r+1)}, y^{(r+1)}, p^{(r+1)})$ by

- i) find $x^{(r+1)}$ such that

$$\langle x - x^{(r+1)}, \partial g(x^{(r+1)}) + A^\top(p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r)})) \rangle \geq 0, \quad \forall x \in \mathbb{R}^d, \tag{10.36}$$

- ii) find $y^{(r+1)}$ such that

$$\langle y - y^{(r+1)}, \partial h(y^{(r+1)}) - (p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)})) \rangle \geq 0, \quad \forall y \in \mathbb{R}^m, \tag{10.37}$$

- iii) update $p^{(r+1)}$ via

$$p^{(r+1)} = p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)}),$$

where $\gamma > 0$ is a fixed penalty parameter. To corroborate the equivalence of the iteration scheme above to ADMM in Algorithm 5, note that (10.35) reduces to $\langle \hat{z}, F(\hat{z}) \rangle \geq 0$ for $z = 2\hat{z}$. On the other hand, (10.35) is equal to $\langle \hat{z}, F(\hat{z}) \rangle \leq 0$ when $z = -\hat{z}$. The both cases transform (10.35) to find a solution \hat{z} of a system of equations $F(\hat{z}) = 0$. Thus, the VI sub-problems (10.36) and (10.37) can be reduced to find a pair $(x^{(r+1)}, y^{(r+1)})$ with

$$\begin{aligned} \partial g(x^{(r+1)}) + A^\top(p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r)})) &= 0, \\ \partial h(y^{(r+1)}) - (p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)})) &= 0. \end{aligned}$$

The both equations correspond to optimality conditions of the minimization sub-problems (10.32) and (10.33) of the ADMM algorithm, respectively. The theoretical properties of ADMM from the perspective of VI problems were studied extensively and a good reference overview can be found in [95].

Relation to Douglas-Rachford Splitting

Finally we want to point out the relation of the ADMM to the Douglas-Rachford splitting algorithm applied to the dual problem, see [44, 69, 70, 82, 86, 164]. We consider again the problem (10.17), i.e.,

$$\operatorname{argmin}_{x \in \mathbb{R}^d} \{g(x) + h(x)\},$$

where we assume this time only $g, h \in \Gamma_0(\mathbb{R}^d)$ and that g or h is continuous at a point in $\operatorname{dom}g \cap \operatorname{dom}h$. Fermat's rule and subdifferential calculus imply that \hat{x} is a minimizer if and only if

$$0 \in \partial g(\hat{x}) + \partial h(\hat{x}) \Leftrightarrow \exists \hat{\xi} \in \eta \partial g(\hat{x}) \text{ such that } \hat{x} = \operatorname{prox}_{\eta h}(\hat{x} - \hat{\xi}). \quad (10.38)$$

The basic idea for finding such minimizer is to set up a 'nice' operator $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ by

$$T := \operatorname{prox}_{\eta h}(2\operatorname{prox}_{\eta g} - I) - \operatorname{prox}_{\eta g} + I, \quad (10.39)$$

whose fixed points \hat{t} are related to the minimizers as follows: setting $\hat{x} := \operatorname{prox}_{\eta g}(\hat{t})$, i.e., $\hat{t} \in \hat{x} + \eta \partial g(\hat{x})$ and $\hat{\xi} := \hat{t} - \hat{x} \in \eta \partial g(\hat{x})$ we see that

$$\begin{aligned} \hat{t} &= T(\hat{t}) = \operatorname{prox}_{\eta h}(2\hat{x} - \hat{t}) - \hat{x} + \hat{t}, \\ \hat{\xi} + \hat{x} &= \operatorname{prox}_{\eta h}(\hat{x} - \hat{\xi}) + \hat{\xi}, \\ \hat{x} &= \operatorname{prox}_{\eta h}(\hat{x} - \hat{\xi}), \end{aligned}$$

which coincides with (10.38). By the proof of the next theorem, the operator T is firmly nonexpansive such that by Theorem 6 a fixed point of T can be found by Picard iterations. This gives rise to the following *Douglas-Rachford splitting algorithm (DRS)*.

Algorithm 7 Douglas-Rachford Splitting Algorithm (DRS)

Initialization: $x^{(0)}, t^{(0)} \in \mathbb{R}^d, \eta > 0$

Iterations: For $r = 0, 1, \dots$

$$t^{(r+1)} = \operatorname{prox}_{\eta h}(2x^{(r)} - t^{(r)}) + t^{(r)} - x^{(r)},$$

$$x^{(r+1)} = \operatorname{prox}_{\eta g}(t^{(r+1)}).$$

The following theorem verifies the convergence of the DRS algorithm. For a recent convergence result see also [97].

Theorem 10 (Convergence of Douglas-Rachford Splitting Algorithm). *Let $g, h \in \Gamma_0(\mathbb{R}^d)$ where one of the functions is continuous at a point in $\text{dom}g \cap \text{dom}h$. Assume that a solution of $\text{argmin}_{x \in \mathbb{R}^d} \{g(x) + h(x)\}$ exists. Then, for any initial $t^{(0)}, x^{(0)} \in \mathbb{R}^d$ and any $\eta > 0$, the DRS sequence $(t^{(r)})_r$ converges to a fixed point \hat{t} of T in (10.39) and $(x^{(r)})_r$ to a solution of the minimization problem.*

Proof. It remains to show that T is firmly nonexpansive. We have for $R_{\eta g} := 2\text{prox}_{\eta g} - I$ and $R_{\eta h} := 2\text{prox}_{\eta h} - I$ that

$$\begin{aligned} 2T &= 2\text{prox}_{\eta h}(2\text{prox}_{\eta g} - I) - (2\text{prox}_{\eta g} - I) + I = R_{\eta h} \circ R_{\eta g} + I, \\ T &= \frac{1}{2}I + \frac{1}{2}R_{\eta h} \circ R_{\eta g}. \end{aligned}$$

The operators $R_{\eta g}, R_{\eta h}$ are nonexpansive since $\text{prox}_{\eta g}$ and $\text{prox}_{\eta h}$ are firmly nonexpansive. Hence $R_{\eta h} \circ R_{\eta g}$ is nonexpansive and we are done. \square

For variety of (sharp) convergence rates of DRS we refer to [64, 65].

The relation of the ADMM algorithm and DRS algorithm applied to the Fenchel dual problem (10.29), i.e.,

$$\begin{aligned} t^{(r+1)} &= \text{prox}_{\eta g^* \circ (-A^T)}(2p^{(r)} - t^{(r)}) + t^{(r)} - p^{(r)}, \\ p^{(r+1)} &= \text{prox}_{\eta h^*}(t^{(r+1)}), \end{aligned} \quad (10.40)$$

is given by the following theorem, see [69, 82].

Theorem 11 (Relation Between ADMM and DRS). *The ADMM sequences $(b^{(r)})_r$ and $(y^{(r)})_r$ are related to the sequences (10.40) generated by the DRS algorithm applied to the dual problem by $\eta = \gamma$ and*

$$\begin{aligned} t^{(r)} &= \eta(b^{(r)} + y^{(r)}), \\ p^{(r)} &= \eta b^{(r)}. \end{aligned} \quad (10.41)$$

Proof. First, we show that

$$\hat{p} = \text{argmin}_{p \in \mathbb{R}^m} \left\{ \frac{\eta}{2} \|Ap - q\|_2^2 + g(p) \right\} \Rightarrow \eta(A\hat{p} - q) = \text{prox}_{\eta g^* \circ (-A^T)}(-\eta q) \quad (10.42)$$

holds true. The left-hand side of (10.42) is equivalent to

$$0 \in \eta A^T(A\hat{p} - q) + \partial g(\hat{p}) \Leftrightarrow \hat{p} \in \partial g^*(-\eta A^T(A\hat{p} - q)).$$

Applying $-\eta A$ on both sides and using the chain rule implies

$$-\eta A\hat{p} \in -\eta A\partial g^*(-\eta A^T(A\hat{p} - q)) = \eta \partial(g^* \circ (-A^T))(\eta(A\hat{p} - q)).$$

Adding $-\eta q$ we get

$$-\eta q \in (I + \eta \partial(g^* \circ (-A^T))) (\eta(A\hat{p} - q)),$$

which is equivalent to the right-hand side of (10.42) by the definition of the resolvent (see Remark 1).

Secondly, applying (10.42) to the first ADMM step with $\gamma = \eta$ and $q := y^{(r)} - b^{(r)}$ yields

$$\eta(b^{(r)} + Ax^{(r+1)} - y^{(r)}) = \text{prox}_{\eta g^* \circ (-A^T)}(\eta(b^{(r)} - y^{(r)})).$$

Assume that the ADMM and DRS iterates have the identification (10.41) up to some $r \in \mathbb{N}$. Using this induction hypothesis it follows that

$$\eta(b^{(r)} + Ax^{(r+1)}) = \text{prox}_{\eta g^* \circ (-A^T)}(\underbrace{\eta(b^{(r)} - y^{(r)})}_{2p^{(r)} - t^{(r)}}) + \underbrace{\eta y^{(r)}}_{t^{(r)} - p^{(r)}} \stackrel{(10.40)}{=} t^{(r+1)} \quad (10.43)$$

By definition of $b^{(r+1)}$ we see that $\eta(b^{(r+1)} + y^{(r+1)}) = t^{(r+1)}$. Next we apply (10.42) in the second ADMM step where we replace g by h and A by $-I$ and use $q := -b^{(r)} - Ax^{(r+1)}$. Together with (10.43) this gives

$$\eta(-y^{(r+1)} + b^{(r)} + Ax^{(r+1)}) = \text{prox}_{\eta h^*}(\underbrace{\eta(b^{(r)} + Ax^{(r+1)})}_{t^{(r+1)}}) \stackrel{(10.40)}{=} p^{(r+1)} \quad (10.44)$$

Using again the definition of $b^{(r+1)}$ we obtain $\eta b^{(r+1)} = p^{(r+1)}$ which completes the proof. \square

A recent work [186] shows that the ADMM is in some sense self-dual, i.e., it is equivalent not only to the DRS applied to the dual problem, but also to the primal one.

6.3 Primal Dual Hybrid Gradient Algorithms

The first ADMM step (10.32) requires in general the solution of a linear system of equations. This can be avoided by modifying this step using the Taylor expansion at $x^{(r)}$:

$$\frac{\gamma}{2} \left\| \frac{1}{\gamma} p^{(r)} + Ax - y^{(r)} \right\|_2^2 \approx \text{const} + \gamma \langle A^T(Ax^{(r)} - y^{(r)} + \frac{1}{\gamma} p^{(r)}), x \rangle + \frac{\gamma}{2} (x - x^{(r)})^T A^T A (x - x^{(r)}),$$

approximating $A^T A \approx \frac{1}{\gamma \tau} I$, setting $\gamma := \sigma$ and using $p^{(r)}/\sigma$ instead of $Ax^{(r)} - y^{(r)} + p^{(r)}/\sigma$ we obtain (note that $p^{(r+1)}/\sigma = p^{(r)}/\sigma + Ax^{(r+1)} - y^{(r+1)}$):

$$\begin{aligned}
x^{(r+1)} &= \operatorname{argmin}_{x \in \mathbb{R}^d} \left\{ g(x) + \frac{1}{2\tau} \|x - (x^{(r)} - \tau A^T p^{(r)})\|_2^2 \right\} = \operatorname{prox}_{\tau g} \left(x^{(r)} - \tau A^T p^{(r)} \right), \\
y^{(r+1)} &= \operatorname{argmin}_{y \in \mathbb{R}^m} \left\{ h(y) + \frac{\sigma}{2} \left\| \frac{1}{\sigma} p^{(r)} + Ax^{(r+1)} - y \right\|_2^2 \right\} = \operatorname{prox}_{\frac{1}{\sigma} h} \left(\frac{1}{\sigma} p^{(r)} + Ax^{(r+1)} \right), \\
p^{(r+1)} &= p^{(r)} + \sigma (Ax^{(r+1)} - y^{(r+1)}).
\end{aligned} \tag{10.45}$$

The above algorithm can be deduced in another way by the Arrow-Hurwicz method: we alternate the minimization in the primal and dual problems (10.27) and (10.28) and add quadratic terms. The resulting sequences

$$\begin{aligned}
x^{(r+1)} &= \operatorname{argmin}_{x \in \mathbb{R}^d} \left\{ g(x) + \langle p^{(r)}, Ax \rangle + \frac{1}{2\tau} \|x - x^{(r)}\|_2^2 \right\}, \\
&= \operatorname{prox}_{\tau g} (x^{(r)} - \tau A^T p^{(r)})
\end{aligned} \tag{10.46}$$

$$\begin{aligned}
p^{(r+1)} &= \operatorname{argmin}_{p \in \mathbb{R}^m} \left\{ h^*(p) - \langle p, Ax^{(r+1)} \rangle + \frac{1}{2\sigma} \|p - p^{(r)}\|_2^2 \right\}, \\
&= \operatorname{prox}_{\sigma h^*} (p^{(r)} + \sigma Ax^{(r+1)})
\end{aligned} \tag{10.47}$$

coincide with those in (10.45) which can be seen as follows: For $x^{(r)}$ the relation is straightforward. From the last equation we obtain

$$\begin{aligned}
p^{(r)} + \sigma Ax^{(r+1)} &\in p^{(r+1)} + \sigma \partial h^*(p^{(r+1)}), \\
\frac{1}{\sigma} (p^{(r)} - p^{(r+1)}) + Ax^{(r+1)} &\in \partial h^*(p^{(r+1)}),
\end{aligned}$$

and using that $p \in \partial h(x) \Leftrightarrow x \in \partial h^*(p)$ further

$$p^{(r+1)} \in \partial h \left(\underbrace{\frac{1}{\sigma} (p^{(r)} - p^{(r+1)}) + Ax^{(r+1)}}_{y^{(r+1)}} \right).$$

Setting

$$y^{(r+1)} := \frac{1}{\sigma} (p^{(r)} - p^{(r+1)}) + Ax^{(r+1)},$$

we get

$$p^{(r+1)} = p^{(r)} + \sigma (Ax^{(r+1)} - y^{(r+1)}) \tag{10.48}$$

and $p^{(r+1)} \in \partial h(y^{(r+1)})$ which can be rewritten as

$$\begin{aligned}
y^{(r+1)} + \frac{1}{\sigma} p^{(r+1)} &\in y^{(r+1)} + \frac{1}{\sigma} \partial h(y^{(r+1)}), \\
\frac{1}{\sigma} p^{(r)} + Ax^{(r+1)} &\in y^{(r+1)} + \frac{1}{\sigma} \partial h(y^{(r+1)}), \\
y^{(r+1)} &= \operatorname{prox}_{\frac{1}{\sigma} h} \left(\frac{1}{\sigma} p^{(r)} + Ax^{(r+1)} \right).
\end{aligned}$$

There are several modifications of the basic “linearized” ADMM which improve its convergence properties as

- the predictor corrector proximal multiplier method [48],
- the primal dual hybrid gradient method (PDHG) [191] with convergence proof in [23],
- primal dual hybrid gradient method with extrapolation of the primal or dual variable [43, 140], a preconditioned version [42] and a generalization [58], Douglas-Rachford-type algorithms [25, 26] for solving inclusion equations, see also [54, 177], as well an extension allowing the operator A to be nonlinear [172].

A good overview on primal-dual methods is also given in [110]. Here is the algorithm proposed by Chambolle, Cremers, and Pock [43, 140].

Algorithm 8 Primal Dual Hybrid Gradient Method with Extrapolation of Dual Variable (PDHGMp)

Initialization: $y^{(0)}, b^{(0)} = b^{(-1)} \in \mathbb{R}^m, \tau, \sigma > 0$ with $\tau\sigma < 1/\|A\|_2^2$ and $\theta \in (0, 1]$

Iterations: For $r = 0, 1, \dots$

$$\begin{aligned} x^{(r+1)} &= \operatorname{argmin}_{x \in \mathbb{R}^d} \left\{ g(x) + \frac{1}{2\tau} \|x - (x^{(r)} - \tau\sigma A^\top \bar{b}^{(r)})\|_2^2 \right\} \\ y^{(r+1)} &= \operatorname{argmin}_{y \in \mathbb{R}^m} \left\{ h(y) + \frac{\sigma}{2} \|b^{(r)} + Ax^{(r+1)} - y\|_2^2 \right\} \\ b^{(r+1)} &= b^{(r)} + Ax^{(r+1)} - y^{(r+1)}. \\ \bar{b}^{(r+1)} &= b^{(r+1)} + \theta(b^{(r+1)} - b^{(r)}) \end{aligned}$$

Note that the new first updating step can be also deduced by applying the so-called *inexact Uzawa algorithm* to the first ADMM step (see Section 6.4). Furthermore, it can be directly seen that for A being the identity and $\theta = 1$ and $\gamma = \sigma = \frac{1}{\tau}$, the PDHGMp algorithm corresponds to the ADMM as well as the Douglas-Rachford splitting algorithm as proposed in Section 6.2. The following theorem and convergence proof are based on [43].

Theorem 12. Let $g \in \Gamma_0(\mathbb{R}^d), h \in \Gamma_0(\mathbb{R}^m)$ and $\theta \in (0, 1]$. Let $\tau, \sigma > 0$ fulfill

$$\tau\sigma < 1/\|A\|_2^2. \tag{10.49}$$

Suppose that the Lagrangian $L(x, p) := g(x) - h^*(p) + \langle Ax, p \rangle$ has a saddle point. Then the sequence $\{(x^{(r)}, p^{(r)})\}_r$ produced by PDGHMp converges to a saddle point of the Lagrangian.

Proof. We restrict the proof to the case $\theta = 1$. For arbitrary $\bar{x} \in \mathbb{R}^d, \bar{p} \in \mathbb{R}^m$ consider according to (10.46) and (10.47) the iterations

$$\begin{aligned} x^{(r+1)} &= (I + \tau\partial g)^{-1} \left(x^{(r)} - \tau A^\top \bar{p} \right), \\ p^{(r+1)} &= (I + \sigma\partial h^*)^{-1} \left(p^{(r)} + \sigma A \bar{x} \right), \end{aligned}$$

i.e.,

$$\frac{x^{(r)} - x^{(r+1)}}{\tau} - A^T \bar{p} \in \partial g(x^{(r+1)}), \quad \frac{p^{(r)} - p^{(r+1)}}{\sigma} + A\bar{x} \in \partial h^*(p^{(r+1)}).$$

By definition of the subdifferential we obtain for all $x \in \mathbb{R}^d$ and all $p \in \mathbb{R}^m$ that

$$\begin{aligned} g(x) &\geq g(x^{(r+1)}) + \frac{1}{\tau} \langle x^{(r)} - x^{(r+1)}, x - x^{(r+1)} \rangle - \langle A^T \bar{p}, x - x^{(r+1)} \rangle, \\ h^*(p) &\geq h^*(p^{(r+1)}) + \frac{1}{\sigma} \langle p^{(r)} - p^{(r+1)}, p - p^{(r+1)} \rangle + \langle p - p^{(r+1)}, A\bar{x} \rangle \end{aligned}$$

and by adding the equations

$$\begin{aligned} 0 &\geq g(x^{(r+1)}) - h^*(p) - \left(g(x) - h^*(p^{(r+1)}) \right) - \langle A^T \bar{p}, x - x^{(r+1)} \rangle + \langle p - p^{(r+1)}, A\bar{x} \rangle \\ &\quad + \frac{1}{\tau} \langle x^{(r)} - x^{(r+1)}, x - x^{(r+1)} \rangle + \frac{1}{\sigma} \langle p^{(r)} - p^{(r+1)}, p - p^{(r+1)} \rangle. \end{aligned}$$

By

$$\langle x^{(r)} - x^{(r+1)}, x - x^{(r+1)} \rangle = \frac{1}{2} \left(\|x^{(r)} - x^{(r+1)}\|_2^2 + \|x - x^{(r+1)}\|_2^2 - \|x - x^{(r)}\|_2^2 \right)$$

this can be rewritten as

$$\begin{aligned} &\frac{1}{2\tau} \|x - x^{(r)}\|_2^2 + \frac{1}{2\sigma} \|p - p^{(r)}\|_2^2 \\ &\geq \frac{1}{2\tau} \|x^{(r)} - x^{(r+1)}\|_2^2 + \frac{1}{2\tau} \|x - x^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p^{(r)} - p^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p - p^{(r+1)}\|_2^2 \\ &\quad + \left(g(x^{(r+1)}) - h^*(p) + \langle p, Ax^{(r+1)} \rangle \right) - \left(g(x) - h^*(p^{(r+1)}) + \langle p^{(r+1)}, Ax \rangle \right) \\ &\quad - \langle p, Ax^{(r+1)} \rangle + \langle p^{(r+1)}, Ax \rangle - \langle \bar{p}, A(x - x^{(r+1)}) \rangle + \langle p - p^{(r+1)}, A\bar{x} \rangle \\ &= \frac{1}{2\tau} \|x^{(r)} - x^{(r+1)}\|_2^2 + \frac{1}{2\tau} \|x - x^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p^{(r)} - p^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p - p^{(r+1)}\|_2^2 \\ &\quad + \left(g(x^{(r+1)}) - h^*(p) + \langle p, Ax^{(r+1)} \rangle \right) - \left(g(x) - h^*(p^{(r+1)}) + \langle p^{(r+1)}, Ax \rangle \right) \\ &\quad + \langle p^{(r+1)} - p, A(x^{(r+1)} - \bar{x}) \rangle - \langle p^{(r+1)} - \bar{p}, A(x^{(r+1)} - x) \rangle. \end{aligned}$$

For any saddle point (x^*, p^*) we have that $L(x^*, p) \leq L(x^*, p^*) \leq L(x, p^*)$ for all x, p so that in particular $0 \leq L(x^{(r+1)}, p^*) - L(x^*, p^{(r+1)})$. Thus, using $(x, p) := (x^*, p^*)$ in the above inequality, we get

$$\begin{aligned} &\frac{1}{2\tau} \|x^* - x^{(r)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(r)}\|_2^2 \\ &\geq \frac{1}{2\tau} \|x^{(r)} - x^{(r+1)}\|_2^2 + \frac{1}{2\tau} \|x^* - x^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p^{(r)} - p^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(r+1)}\|_2^2 \\ &\quad + \langle p^{(r+1)} - p^*, A(x^{(r+1)} - \bar{x}) \rangle - \langle p^{(r+1)} - \bar{p}, A(x^{(r+1)} - x^*) \rangle. \end{aligned}$$

In the algorithm we use $\bar{x} := x^{(r+1)}$ and $\bar{p} := 2p^{(r)} - p^{(r-1)}$. Note that $\bar{p} = p^{(r+1)}$ would be the better choice, but this is impossible if we want to keep on an explicit algorithm. For these values the above inequality further simplifies to

$$\begin{aligned} & \frac{1}{2\tau} \|x^* - x^{(r)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(r)}\|_2^2 \\ \geq & \frac{1}{2\tau} \|x^{(r)} - x^{(r+1)}\|_2^2 + \frac{1}{2\tau} \|x^* - x^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p^{(r)} - p^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(r+1)}\|_2^2 \\ & + \langle p^{(r+1)} - 2p^{(r)} + p^{(r-1)}, A(x^* - x^{(r+1)}) \rangle. \end{aligned}$$

We estimate the last summand using Cauchy-Schwarz's inequality as follows:

$$\begin{aligned} & \langle p^{(r+1)} - p^{(r)} - (p^{(r)} - p^{(r-1)}), A(x^* - x^{(r+1)}) \rangle \\ = & \langle p^{(r+1)} - p^{(r)}, A(x^* - x^{(r+1)}) \rangle - \langle p^{(r)} - p^{(r-1)}, A(x^* - x^{(r)}) \rangle \\ & - \langle p^{(r)} - p^{(r-1)}, A(x^{(r)} - x^{(r+1)}) \rangle \\ \geq & \langle p^{(r+1)} - p^{(r)}, A(x^* - x^{(r+1)}) \rangle - \langle p^{(r)} - p^{(r-1)}, A(x^* - x^{(r)}) \rangle \\ & - \|A\|_2 \|x^{(r+1)} - x^{(r)}\|_2 \|p^{(r)} - p^{(r-1)}\|_2. \end{aligned}$$

Since

$$2uv \leq \alpha u^2 + \frac{1}{\alpha} v^2, \quad \alpha > 0, \quad (10.50)$$

we obtain

$$\begin{aligned} \|A\|_2 \|x^{(r+1)} - x^{(r)}\|_2 \|p^{(r)} - p^{(r-1)}\|_2 & \leq \frac{\|A\|_2}{2} \left(\alpha \|x^{(r+1)} - x^{(r)}\|_2^2 + \frac{1}{\alpha} \|p^{(r)} - p^{(r-1)}\|_2^2 \right) \\ & = \frac{\|A\|_2 \alpha \tau}{2\tau} \|x^{(r+1)} - x^{(r)}\|_2^2 + \frac{\|A\|_2 \sigma}{2\alpha \sigma} \|p^{(r)} - p^{(r-1)}\|_2^2. \end{aligned}$$

With $\alpha := \sqrt{\sigma/\tau}$ the relation

$$\|A\|_2 \alpha \tau = \frac{\|A\|_2 \sigma}{\alpha} = \|A\|_2 \sqrt{\sigma \tau} < 1$$

holds true. Thus, we get

$$\begin{aligned} & \frac{1}{2\tau} \|x^* - x^{(r)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(r)}\|_2^2 \\ \geq & \frac{1}{2\tau} \|x^* - x^{(r+1)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(r+1)}\|_2^2 \\ & + \frac{1}{2\tau} (1 - \|A\|_2 \sqrt{\sigma \tau}) \|x^{(r+1)} - x^{(r)}\|_2^2 + \frac{1}{2\sigma} \|p^{(r+1)} - p^{(r)}\|_2^2 - \frac{\|A\|_2 \sqrt{\sigma \tau}}{2\sigma} \|p^{(r)} - p^{(r-1)}\|_2^2 \\ & + \langle p^{(r+1)} - p^{(r)}, A(x^* - x^{(r+1)}) \rangle - \langle p^{(r)} - p^{(r-1)}, A(x^* - x^{(r)}) \rangle. \end{aligned} \quad (10.51)$$

Summing up these inequalities from $r = 0$ to $N - 1$ and regarding that $p^{(0)} = p^{(-1)}$, we obtain

$$\begin{aligned}
& \frac{1}{2\tau} \|x^* - x^{(0)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(0)}\|_2^2 \\
\geq & \frac{1}{2\tau} \|x^* - x^{(N)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(N)}\|_2^2 \\
& + (1 - \|A\|_2 \sqrt{\sigma\tau}) \left(\frac{1}{2\tau} \sum_{r=1}^N \|x^{(r)} - x^{(r-1)}\|_2^2 + \frac{1}{2\sigma} \sum_{r=1}^{N-1} \|p^{(r)} - p^{(r-1)}\|_2^2 \right) \\
& + \frac{1}{2\sigma} \|p^{(N)} - p^{(N-1)}\|_2^2 + \langle p^{(N)} - p^{(N-1)}, A(x^* - x^{(N)}) \rangle
\end{aligned}$$

By

$$\langle p^{(N)} - p^{(N-1)}, A(x^{(N)} - x^*) \rangle \leq \frac{1}{2\sigma} \|p^{(N)} - p^{(N-1)}\|_2^2 + \frac{\|A\|_2^2 \sigma\tau}{2\tau} \|x^{(N)} - x^*\|_2^2$$

this can be further estimated as

$$\begin{aligned}
& \frac{1}{2\tau} \|x^* - x^{(0)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(0)}\|_2^2 \\
\geq & \frac{1}{2\tau} (1 - \|A\|_2^2 \sigma\tau) \|x^* - x^{(N)}\|_2^2 + \frac{1}{2\sigma} \|p^* - p^{(N)}\|_2^2 \\
& + (1 - \|A\|_2 \sqrt{\sigma\tau}) \left(\frac{1}{2\tau} \sum_{r=1}^N \|x^{(r)} - x^{(r-1)}\|_2^2 + \frac{1}{2\sigma} \sum_{r=1}^{N-1} \|p^{(r)} - p^{(r-1)}\|_2^2 \right). \quad (10.52)
\end{aligned}$$

By (10.52) we conclude that the sequence $\{(x^{(n)}, p^{(n)})\}_n$ is bounded. Thus, there exists a convergent subsequence $\{(x^{(n_j)}, p^{(n_j)})\}_j$ which converges to some point (\hat{x}, \hat{p}) as $j \rightarrow \infty$. Further, we see by (10.52) that

$$\lim_{n \rightarrow \infty} (x^{(n)} - x^{(n-1)}) = 0, \quad \lim_{n \rightarrow \infty} (p^{(n)} - p^{(n-1)}) = 0.$$

Consequently,

$$\lim_{j \rightarrow \infty} (x^{(n_j-1)} - \hat{x}) = 0, \quad \lim_{j \rightarrow \infty} (p^{(n_j-1)} - \hat{p}) = 0$$

holds true. Let T denote the iteration operator of the PDHGMp cycles, i.e., $T(x^{(r)}, p^{(r)}) = (x^{(r+1)}, p^{(r+1)})$. Since T is the concatenation of affine operators and proximation operators, it is continuous. Now we have that $T(x^{(n_j-1)}, p^{(n_j-1)}) = (x^{(n_j)}, p^{(n_j)})$ and taking the limits for $j \rightarrow \infty$ we see that $T(\hat{x}, \hat{p}) = (\hat{x}, \hat{p})$ so that (\hat{x}, \hat{p}) is a fixed point of the iteration and thus a saddle point of L . Using this particular saddle point in (10.51) and summing up from $r = n_j$ to $N - 1$, $N > n_j$ we obtain

$$\begin{aligned}
 & \frac{1}{2\tau} \|\hat{x} - x^{(n_j)}\|_2^2 + \frac{1}{2\sigma} \|\hat{p} - p^{(n_j)}\|_2^2 \\
 \geq & \frac{1}{2\tau} \|\hat{x} - x^{(N)}\|_2^2 + \frac{1}{2\sigma} \|\hat{p} - p^{(N)}\|_2^2 \\
 & + (1 - \|A\|_2 \sqrt{\sigma\tau}) \left(\frac{1}{2\tau} \sum_{r=n_j}^{N-1} \|x^{(r+1)} - x^{(r)}\|_2^2 + \frac{1}{2\sigma} \sum_{r=n_j+1}^{N-1} \|p^{(r)} - p^{(r-1)}\|_2^2 \right) \\
 & + \frac{1}{2\sigma} \|p^{(N)} - p^{(N-1)}\|_2^2 - \frac{\|A\|_2 \sqrt{\sigma\tau}}{2\sigma} \|p^{(n_j)} - p^{(n_j-1)}\|_2^2 \\
 & + \langle p^{(N)} - p^{(N-1)}, A(\hat{x} - x^{(N)}) \rangle - \langle p^{(n_j)} - p^{(n_j-1)}, A(\hat{x} - x^{(n_j)}) \rangle
 \end{aligned}$$

and further

$$\begin{aligned}
 & \frac{1}{2\tau} \|\hat{x} - x^{(n_j)}\|_2^2 + \frac{1}{2\sigma} \|\hat{p} - p^{(n_j)}\|_2^2 \\
 \geq & \frac{1}{2\tau} \|\hat{x} - x^{(N)}\|_2^2 + \frac{1}{2\sigma} \|\hat{p} - p^{(N)}\|_2^2 \\
 & + \frac{1}{2\sigma} \|p^{(N)} - p^{(N-1)}\|_2^2 - \frac{\|A\|_2 \sqrt{\sigma\tau}}{2\sigma} \|p^{(n_j)} - p^{(n_j-1)}\|_2^2 \\
 & + \langle p^{(N)} - p^{(N-1)}, A(\hat{x} - x^{(N)}) \rangle - \langle p^{(n_j)} - p^{(n_j-1)}, A(\hat{x} - x^{(n_j)}) \rangle
 \end{aligned}$$

For $j \rightarrow \infty$ this implies that $(x^{(N)}, p^{(N)})$ converges also to (\hat{x}, \hat{y}) and we are done. \square

6.4 Proximal ADMM

To avoid the computation of a linear system of equations in the first ADMM step (10.32), we can more generally use the proximal ADMM algorithm [95, 190] that can be interpreted as a preconditioned variant of ADMM. In this variant the minimization step (10.32) is replaced by a proximal-like iteration based on the general proximal operator (10.5),

$$x^{(r+1)} = \operatorname{argmin}_{x \in \mathbb{R}^d} \{L_\gamma(x, y^{(r)}, p^{(r)}) + \frac{1}{2} \|x - x^{(r)}\|_R^2\} \tag{10.53}$$

with a symmetric, positive definite matrix $R \in \mathbb{R}^{d,d}$. The introduction of R provides an additional flexibility to cancel out linear operators which might be difficult to invert. In addition the modified minimization problem is strictly convex inducing a unique minimizer. In the same manner the second ADMM step (10.33) can also be extended by a proximal term $(1/2)\|y - y^{(r)}\|_S^2$ with a symmetric, positive definite matrix $S \in \mathbb{R}^{m,m}$ [190]. The convergence analysis of the proximal ADMM was provided in [190] and the algorithm can be also classified as an inexact Uzawa method. A generalization, where the matrices R and S can non-monotonically vary in each iteration step, was analyzed in [95], additionally allowing an inexact computation of minimization problems.

In case of the PDHGMP algorithm, it was mentioned that the first updating step can be deduced by applying the inexact Uzawa algorithm to the first ADMM step. Using the proximal ADMM, it is straightforward to see that the first updating step of the PDHGMP algorithm with $\theta = 1$ corresponds to (10.53) in case of $R = \frac{1}{\tau}I - \sigma A^T A$ with $0 < \tau < 1/\|\sigma A^T A\|$, see [43, 71]. Further relations of the (proximal) ADMM to primal dual hybrid methods discussed above can be found in [71].

6.5 Bregman Methods

Bregman methods became very popular in image processing by a series papers of Osher and co-workers, see, e.g., [91, 135]. Many of these methods can be interpreted as ADMM methods and its linearized versions. In the following we briefly sketch these relations.

The PPA is a special case of the Bregman PPA. Let $\varphi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be a convex function. Then the *Bregman distance* $D_\varphi^p : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is given by

$$D_\varphi^p(x, y) = \varphi(x) - \varphi(y) - \langle p, x - y \rangle$$

with $p \in \partial\varphi(y)$, $y \in \text{dom}f$. If $\partial\varphi(y)$ contains only one element, we just write D_φ . If φ is smooth, then the Bregman distance can be interpreted as subtracting the first order Taylor expansion of $\varphi(x)$ at y .

Example 5. (Special Bregman Distances)

1. The Bregman distance corresponding to $\varphi(x) := \frac{1}{2}\|x\|_2^2$ is given by $D_\varphi(x, y) = \frac{1}{2}\|x - y\|_2^2$.
2. For the negative Shannon entropy $\varphi(x) := \langle 1_d, x \log x \rangle$, $x > 0$ we obtain the (discrete) I -divergence or generalized Kullback-Leibler entropy $D_\varphi(x, y) = x \log \frac{x}{y} - x + y$.

For $f \in \Gamma_0(\mathbb{R}^d)$ we consider the generalized proximal problem

$$\operatorname{argmin}_{y \in \mathbb{R}^d} \left\{ \frac{1}{\gamma} D_\varphi^p(x, y) + f(y) \right\}.$$

The Bregman Proximal Point Algorithm (BPPA) for solving this problem reads as follows:

Algorithm 9 Bregman Proximal Point Algorithm (BPPA)

Initialization: $x^{(0)} \in \mathbb{R}^d$, $p^{(0)} \in \partial\varphi(x^{(0)})$, $\gamma > 0$

Iterations: For $r = 0, 1, \dots$

$$x^{(r+1)} = \operatorname{argmin}_{y \in \mathbb{R}^d} \left\{ \frac{1}{\gamma} D_\varphi^{p^{(r)}}(y, x^{(r)}) + f(y) \right\}$$

$$p^{(r+1)} \in \partial\varphi(x^{(r+1)})$$

The BPPA converges for any initialization $x^{(0)}$ to a minimizer of f if $f \in \Gamma_0(\mathbb{R}^d)$ attains its minimum and φ is finite, lower semi-continuous and strictly convex. For convergence proofs we refer, e.g., to [107, 108]. We are interested again in the problem (10.24), i.e.,

$$\operatorname{argmin}_{x \in \mathbb{R}^d, y \in \mathbb{R}^m} \{ \Phi(x, y) \quad \text{s.t.} \quad Ax = y \}, \quad \Phi(x, y) := g(x) + h(y).$$

We consider the BPP algorithm for the objective function $f(x, y) := \frac{1}{2} \|Ax - y\|^2$ with the Bregman distance

$$D_{\Phi}^{(p_x^{(r)}, p_y^{(r)})} \left((x, y), (x^{(r)}, y^{(r)}) \right) = \Phi(x, y) - \Phi(x^{(r)}, y^{(r)}) - \langle p_x^{(r)}, x - x^{(r)} \rangle - \langle p_y^{(r)}, y - y^{(r)} \rangle,$$

where $(p_x^{(r)}, p_y^{(r)}) \in \partial \Phi(x^{(r)}, y^{(r)})$. This results in

$$(x^{(r+1)}, y^{(r+1)}) = \operatorname{argmin}_{x \in \mathbb{R}^d, y \in \mathbb{R}^m} \left\{ \frac{1}{\gamma} D_{\Phi}^{(p_x^{(r)}, p_y^{(r)})} \left((x, y), (x^{(r)}, y^{(r)}) \right) + \frac{1}{2} \|Ax - y\|^2 \right\},$$

$$p_x^{(r+1)} = p_x^{(r)} - \gamma A^* (Ax^{(r+1)} - y^{(r+1)}), \tag{10.54}$$

$$p_y^{(r+1)} = p_y^{(r)} + \gamma (Ax^{(r+1)} - y^{(r+1)}), \tag{10.55}$$

where we have used that the first equation implies

$$0 \in \frac{1}{\gamma} \partial \left(\Phi(x^{(r+1)}, y^{(r+1)}) - \langle p_x^{(r)}, p_y^{(r)} \rangle \right) + (A^* (Ax^{(r+1)} - y^{(r+1)}), -(Ax^{(r+1)} - y^{(r+1)})),$$

$$0 \in \partial \Phi(x^{(r+1)}, y^{(r+1)}) - (p_x^{(r+1)}, p_y^{(r+1)}),$$

so that $(p_x^{(r)}, p_y^{(r)}) \in \partial \Phi(x^{(r)}, y^{(r)})$. From (10.54) and (10.55) we see by induction that $p_x^{(r)} = -A^* p_y^{(r)}$. Setting $p^{(r)} = p_y^{(r)}$ and regarding that

$$\begin{aligned} & \frac{1}{\gamma} D_{\Phi}^{p^{(r)}} \left((x, y), (x^{(r)}, y^{(r)}) \right) + \frac{1}{2} \|Ax - y\|_2^2 \\ &= \operatorname{const} + \frac{1}{\gamma} \left(\Phi(x, y) + \langle A^* p^{(r)}, x \rangle - \langle p^{(r)}, y \rangle \right) + \frac{1}{2} \|Ax - y\|_2^2 \\ &= \operatorname{const} + \frac{1}{\gamma} \left(\Phi(x, y) + \frac{\gamma}{2} \left\| \frac{p^{(r)}}{\gamma} + Ax - y \right\|_2^2 \right) \end{aligned}$$

we obtain the following split Bregman method, see [91]:

Obviously, this is exactly the form of the augmented Lagrangian method in (10.31). Concerning the elastic net example 2 we refer to [38, 114, 187] for convergence results of the so-called linearized Bregman algorithm.

Algorithm 10 Split Bregman Algorithm**Initialization:** $x^{(0)} \in \mathbb{R}^d, p^{(0)}, \gamma > 0$ **Iterations:** For $r = 0, 1, \dots$

$$\begin{aligned} (x^{(r+1)}, y^{(r+1)}) &= \operatorname{argmin}_{x \in \mathbb{R}^d, y \in \mathbb{R}^m} \left\{ \Phi(x, y) + \frac{\gamma}{2} \left\| \frac{p^{(r)}}{\gamma} + Ax - y \right\|_2^2 \right\} \\ p^{(r+1)} &= p^{(r)} + \gamma(Ax^{(r+1)} - y^{(r+1)}) \end{aligned}$$

7 Iterative Regularization for Ill-Posed Problems

So far we have discussed the use of splitting methods for the numerical solution of well-posed variational problems, which arise in a discrete setting and in particular for the standard approach of Tikhonov-type regularization in inverse problems in imaging. The latter is based on minimizing a weighted sum of a data fidelity and a regularization functional, and can be more generally analyzed in Banach spaces, cf. [163]. However, such approaches have several disadvantages, in particular it has been shown that they lead to unnecessary bias in solutions, e.g., a contrast loss in the case of total variation regularization, cf. [136, 31]. A successful alternative to overcome this issue is iterative regularization, which directly applies iterative methods to solve the constrained variational problem

$$\operatorname{argmin}_{x \in \mathcal{X}} \{g(x) \quad \text{s.t.} \quad Ax = f\}. \quad (10.56)$$

Here $A : \mathcal{X} \rightarrow \mathcal{Y}$ is a bounded linear operator between Banach spaces (also nonlinear versions can be considered, cf. [9, 105]) and f are given data. In the well-posed case, (10.56) can be rephrased as the saddle-point problem

$$\min_{x \in \mathcal{X}} \sup_q (g(x) - \langle q, Ax - f \rangle) \quad (10.57)$$

The major issue compared to the discrete setting is that for many prominent examples the operator A does not have a closed range (and hence a discontinuous pseudo-inverse), which makes (10.56) ill-posed. From the optimization point of view, this raises two major issues:

- *Emptiness of the constraint set:* In the practically relevant case of noisy measurements one has to expect that f is not in the range of A , i.e., the constraint cannot be satisfied exactly. Reformulated in the constrained optimization view, the standard paradigm in iterative regularization is to construct an iteration slowly increasing the functional g while decreasing the error in the constraint.
- *Nonexistence of saddle points:* Even if the data or an idealized version Ax^* to be approximated are in the range of A , the existence of a saddle point (x^*, q^*) of (10.57) is not guaranteed. The optimality condition for the latter would yield

$$A^*q^* \in \partial g(x^*), \quad (10.58)$$

which is indeed an abstract smoothness condition on the subdifferential of g at x^* if A and consequently A^* are smoothing operators, it is known as source condition in the field of inverse problems, cf. [31].

Due to the above reasons the use of iterative methods for solving respectively approximating (10.56) has a different flavor than iterative methods for well-posed problems. The key idea is to employ the algorithm as an iterative regularization method, cf. [105], where appropriate stopping in dependence on the noise, i.e. a distance between Ax^* and f , needs to be performed in order to obtain a suitable approximation. The notion to be used is called semiconvergence, i.e., if $\delta > 0$ denotes a measure for the data error (noise level) and $\hat{r}(\delta)$ is the stopping index of the iteration in dependence on δ , then we look for convergence

$$x^{(\hat{r}(\delta))} \rightarrow x^* \quad \text{as } \delta \rightarrow 0, \quad (10.59)$$

in a suitable topology. The minimal ingredient in the convergence analysis is the convergence $x^{(r)} \rightarrow x^*$, which already needs different approaches as discussed above. For iterative methods working on primal variables one can at least use the existence of (10.56) in this case, while real primal-dual iterations still suffer from the potential nonexistence of solutions of the saddle point problem (10.57).

A well-understood iterative method is the Bregman iteration

$$x^{(r+1)} \in \operatorname{argmin}_{x \in \mathcal{X}} \left(\frac{\mu}{2} \|Ax - f\|^2 + D_g^{p^{(r)}}(x, x^{(r)}) \right), \quad (10.60)$$

with $p^{(r)} \in \partial g(x^{(r)})$, which has been analyzed as an iterative method in [136], respectively for nonlinear A in [9]. Note that again with $p^{(r)} = A^*q^{(r)}$ the Bregman iteration is equivalent to the augmented Lagrangian method for the saddle-point problem (10.57). With such iterative regularization methods superior results compared to standard variational methods can be computed for inverse and imaging problems, in particular bias can be eliminated, cf. [31].

The key properties are the *decrease of the data fidelity*

$$\|Ax^{(r+1)} - f\|^2 \leq \|Ax^{(r)} - f\|^2, \quad (10.61)$$

for all r and the *decrease of the Bregman distance to the clean solution*

$$D_g^{p^{(r+1)}}(x^*, x^{(r+1)}) \leq D_g^{p^{(r)}}(x^*, x^{(r)}) \quad (10.62)$$

for those r such that

$$\|Ax^{(r)} - f\| \geq \|Ax^* - f\| = \delta.$$

Together with a more detailed analysis of the difference between consecutive Bregman distances, this can be used to prove semiconvergence results in appropriate weak topologies, cf. [136, 163]. In [9] further variants approximating the quadratic terms, such as the linearized Bregman iteration, are analyzed, however with further restrictions on g . For all other iterative methods discussed above, a convergence

analysis in the case of ill-posed problems is completely open and appears to be a valuable task for future research. Note that in the realization of the Bregman iteration, a well-posed but complex variational problem needs to be solved in each step. By additional operator splitting in an iterative regularization method one could dramatically reduce the computational effort.

If the source condition is satisfied, i.e. if there exists a saddle-point (x^*, q^*) , one can further exploit the decrease of *dual distances*

$$\|q^{(r+1)} - q^*\| \leq \|q^{(r)} - q^*\| \quad (10.63)$$

to obtain a quantitative estimate on the convergence speed, we refer to [30, 32] for a further discussion.

8 Applications

So far we have focused on technical aspects of first order algorithms whose (further) development has been heavily forced by practical applications. In this section we give a rough overview of the use of first order algorithms in practice. We start with applications from classical imaging tasks such as computer vision and image analysis and proceed to applications in natural and life sciences. From the area of biomedical imaging, we will present the *Positron Emission Tomography (PET)* and *Spectral X-ray CT* in more detail and show some results reconstructed with first order algorithms.

At the beginning it is worth to emphasize that many algorithms based on proximal operators such as proximal point algorithm, proximal forward-backward splitting, ADMM, or Douglas-Rachford splitting have been introduced in the 1970s, cf. [83, 116, 146, 147]. However these algorithms have found a broad application in the last two decades, mainly caused by the technological progress. Due to the ability of distributed convex optimization with ADMM related algorithms, these algorithms seem to be qualified for ‘big data’ analysis and large-scale applications in applied statistics and machine learning, e.g., in areas as artificial intelligence, internet applications, computational biology and medicine, finance, network analysis, or logistics [27, 139]. Another boost for the popularity of first order splitting algorithms was the increasing use of sparsity-promoting regularizers based on ℓ_1 - or L_1 -type penalties [91, 151], in particular in combination with inverse problems considering nonlinear image formation models [9, 172] and/or statistically derived (inverse) problems [31]. The latter problems lead to non-quadratic fidelity terms which result from the non-Gaussian structure of the noise model. The overview given in the following mainly concentrates on the latter mentioned applications.

The most classical application of first order splitting algorithms is image analysis such as denoising, where these methods were originally pushed by the Rudin, Osher, and Fatemi (ROF) model [151]. This model and its variants are frequently used as prototypes for total variation methods in imaging to illustrate the applicability of

proposed algorithms in case of non-smooth cost functions, cf. [43, 70, 71, 91, 136, 164, 191]. Since the standard L_2 fidelity term is not appropriate for non-Gaussian noise models, modifications of the ROF problem have been considered in the past and were solved using splitting algorithms to denoise images perturbed by non-Gaussian noise, cf. [19, 43, 76, 153, 168]. Due to the close relation of total variation techniques to image segmentation [31, 140], first order algorithms have been also applied in this field of applications (cf. [42, 43, 90, 140]). Other image analysis tasks where proximal based algorithms have been applied successfully are deblurring and zooming (cf. [23, 43, 71, 166]), inpainting [43], stereo and motion estimation [46, 43, 188], and segmentation [10, 115, 140, 106].

Due to increasing noise level in modern biomedical applications, the requirement on statistical image reconstruction methods has been risen recently and the proximal methods have found access to many applied areas of biomedical imaging. Among the enormous amount of applications from the last two decades, we only give the following selection and further links to the literature:

- *X-ray CT*: Recently statistical image reconstruction methods have received increasing attention in X-ray CT due to increasing noise level encountered in modern CT applications such as sparse/limited-view CT and low-dose imaging, cf., e.g., [173, 178, 179], or K-edge imaging where the concentrations of K-edge materials are inherently low, see, e.g., [159, 158, 160]. In particular, first order splitting methods have received strong attention due to the ability to handle non-standard noise models and sparsity-promoting regularizers efficiently. Beside the classical fan-beam and cone-beam X-ray CT (see, e.g., [4, 45, 51, 104, 130, 145, 167, 173]), the algorithms have also found applications in emerging techniques such as spectral CT, see Section 8.2 and [85, 155, 185] or phase contrast CT [59, 131, 184].
- *Magnetic resonance imaging (MRI)*: Image reconstruction in MRI is mainly achieved by inverting the Fourier transform which can be performed efficiently and robustly if a sufficient number of Fourier coefficients is measured. However, this is not the case in special applications such as fast MRI protocols, cf., e.g., [117, 142], where the Fourier space is undersampled so that the Nyquist criterion is violated and Fourier reconstructions exhibit aliasing artifacts. Thus, compressed sensing theory have found the way into MRI by exploiting sparsity-promoting variational approaches, see, e.g., [15, 103, 118, 144]. Furthermore, in advanced MRI applications such as velocity-encoded MRI or diffusion MRI, the measurements can be modeled more accurately by nonlinear operators and splitting algorithms provide the ability to handle the increased reconstruction complexity efficiently [172].
- *Emission tomography*: Emission tomography techniques used in nuclear medicine such as *positron emission tomography (PET)* and *single photon emission computed tomography (SPECT)* [183] are classical examples for inverse problems in biomedical imaging where statistical modeling of the reconstruction problem is essential due to Poisson statistics of the data. In addition, in cases where short time or low tracer dose measurements are available (e.g., using cardiac and/or respiratory gating [34]) or tracer with a short radioactive half-life are used (e.g., radioactive water $H_2^{15}O$ [157]), the measurements

suffer from inherently high noise level and thus a variety of first order splitting algorithms has been utilized in emission tomography, see, e.g., [4, 16, 122, 123, 143, 153].

- *Optical microscopy*: In modern light microscopy techniques such as *stimulated emission depletion (STED)* or *4Pi-confocal fluorescence microscopy* [99, 100] resolutions beyond the diffraction barrier can be achieved allowing imaging at nanoscales. However, by reaching the diffraction limit of light, measurements suffer from blurring effects and Poisson noise with low photon count rates [66, 162], in particular in live imaging and in high resolution imaging at nanoscopic scales. Thus, regularized (blind) deconvolution addressing appropriate Poisson noise is quite beneficial and proximal algorithms have been applied to achieve this goal, cf., e.g., [30, 81, 153, 166].
- *Other modalities*: It is quite natural that first order splitting algorithms have found a broad usage in biomedical imaging, in particular in such applications where the measurements are highly perturbed by noise and thus regularization with probably a proper statistical modeling are essential as, e.g., in optical tomography [1, 80], medical ultrasound imaging [154], hybrid photo-/optoacoustic tomography [84, 180], or electron tomography [92].

8.1 Positron Emission Tomography (PET)

PET is a biomedical imaging technique visualizing biochemical and physiological processes such as glucose metabolism, blood flow, or receptor concentrations, see, e.g., [183]. This modality is mainly applied in nuclear medicine and the data acquisition is based on weak radioactively marked pharmaceuticals (so-called tracers), which are injected into the blood circulation. Then bindings dependent on the choice of the tracer to the molecules are studied. Since the used markers are radioisotopes, they decay by emitting a positron which annihilates almost immediately with an electron. The resulting emission of two photons is detected and, due to the radioactive decay, the measured data can be modeled as an inhomogeneous Poisson process with a mean given by the X-ray transform of the spatial tracer distribution (cf., e.g., [124, 174]). Note that, up to notation, the X-ray transform coincides with the more popular Radon transform in the two dimensional case [124]. Thus, the underlying reconstruction problem can be modeled as

$$\sum_{m=1}^M ((Ku)_m - f_m \log((Ku)_m)) + \alpha R(u) \rightarrow \min_{u \geq 0}, \quad \alpha > 0, \quad (10.64)$$

where M is the number of measurements, f are the given data, and K is the system matrix which describes the full properties of the PET data acquisition.

To solve (10.64), algorithms discussed above can be applied and several of them have been already studied for PET recently. In the following, we will give a (certainly incomplete) performance discussion of different first order splitting algorithms on synthetic PET data and highlight the strengths and weaknesses of them

which could be carried over to many other imaging applications. For the study below, the total variation (TV) was applied as regularization energy R in (10.64) and the following algorithms and parameter settings were used for the performance evaluation:

- **FB-EM-TV:** The FB-EM-TV algorithm [153] represents an instance of the proximal forward-backward (FB) splitting algorithm discussed in Section 5.2 using a variable metric strategy (10.22). The preconditioned matrices $Q^{(r)}$ in (10.22) are chosen in a way that the gradient descent step corresponds to an expectation-maximization (EM) reconstruction step. The EM algorithm is a classically applied (iterative) reconstruction method in emission tomography [124, 174]. The TV proximal problem was solved by an adopted variant of the modified Arrow-Hurwicz method proposed in [43] since it was shown to be the most efficient method for TV penalized weighted least-squares denoising problems in [152]. Furthermore, a warm starting strategy was used to initialize the dual variables within the TV proximal problem and the inner iteration sequence was stopped if the relative error of primal and dual optimality conditions was below an error tolerance δ , i.e., using the notations from [43], if

$$\max\{d^{(r)}, p^{(r)}\} \leq \delta \quad (10.65)$$

with

$$\begin{aligned} d^{(r)} &= \|(y^{(r)} - y^{(r-1)})/\sigma_{r-1} + \nabla(x^{(r)} - x^{(r-1)})\| / \|\nabla x^{(r)}\|, \\ p^{(r)} &= \|x^{(r)} - x^{(r-1)}\| / \|x^{(r)}\|. \end{aligned}$$

The damping parameter $\eta^{(r)}$ in (10.22) was set to $\eta^{(r)} = 1$ as indicated in [153].

- **FB-EM-TV-Nes83:** A modified version of FB-EM-TV described above using the acceleration strategy proposed by Nesterov in [128]. This modification can be seen as a variant of FISTA [13] with a variable metric strategy (10.22). Here, $\eta^{(r)}$ in (10.22) was chosen fixed (i.e. $\eta^{(r)} = \eta$) but has to be adopted to the predefined inner accuracy threshold δ (10.65) to guarantee the convergence of the algorithm and it was to be done manually.
- **CP-E:** The fully explicit variant of the Chambolle-Pock's primal-dual algorithm [43] (cf. Section 6.3) studied for PET reconstruction problems in [3] (see *CP2TV* in [3]). The dual step size σ was set manually and the primal one corresponding to [43] as $\tau\sigma(\|\nabla\|^2 + \|K\|^2) = 1$, where $\|K\|$ was pre-estimated using the Power method.
- **Precond-CP-E:** The CP-E algorithm described above but using the diagonal preconditioning strategy proposed in [42] with $\alpha = 1$ in [42, Lemma 2].
- **CP-SI:** The semi-implicit variant of the Chambolle-Pock's primal-dual algorithm [43] (cf. Section 6.3) studied for PET reconstruction problems in [3] (see *CP1TV* in [3]). The difference to CP-E is that a TV proximal problem has to be solved in each iteration step. This was performed as in case of FB-EM-TV method. Furthermore, the dual step size σ was set manually and the primal one corresponding to [43] as $\tau\sigma\|K\|^2 = 1$, where $\|K\|$ was pre-estimated using the Power method.

- **Precond-CP-SI:** The CP-SI algorithm described above but using the diagonal preconditioning strategy proposed in [42] with $\alpha = 1$ in [42, Lemma 2].
- **PIDSplit+:** An ADMM based algorithm (cf. Section 6.2) that has been discussed for Poisson deblurring problems of the form (10.64) in [166]. However, in case of PET reconstruction problems, the solution of a linear system of equations of the form

$$(I + K^T K + \nabla^T \nabla) u^{(r+1)} = z^{(r)} \quad (10.66)$$

has to be computed in a different way in contrast to deblurring problems. This was done by running two preconditioned conjugate gradient (PCG) iterations with warm starting and cone filter preconditioning whose effectiveness has been validated in [145] for X-ray CT reconstruction problems. The cone filter was constructed as described in [73, 74] and diagonalized by the discrete cosine transform (DCT-II) supposing Neumann boundary conditions. The PIDSplit+ algorithm described above can be accomplished by a strategy of adaptive augmented Lagrangian parameters γ in (10.26) as proposed for the PIDSplit+ algorithm in [27, 169]. The motivation behind this strategy is to mitigate the performance dependency on the initial chosen fixed parameter that may strongly influence the speed of convergence of ADMM based algorithms.

All algorithms were implemented in MATLAB and executed on a machine with 4 CPU cores, each 2.83 GHz, and 7.73 GB physical memory, running a 64 bit Linux system and MATLAB 2013b. The built-in multi-threading in MATLAB was disabled such that all computations were limited to a single thread. The algorithms were evaluated on a simple object (image size 256×256) and the synthetic 2D PET measurements were obtained via a Monte-Carlo simulation with 257 radial and 256 angular samples, using one million simulated events (see Figure 10.2). Due to sustainable image and measurement dimensions, the system matrix K was pre-computed for all reconstruction runs. To evaluate the performance of algorithms described above, the following procedure was applied. First, since K is injective and thus an unique minimizer of (10.64) is guaranteed [153], we can run a well-performing method for a very long time to compute a “ground truth” solution u_α^* for a fixed α . To this end, we have run the Precond-CP-E algorithm for 100,000 iterations for the following reasons: (1) all iteration steps can be solved exactly such that the solution cannot be influenced by inexact computations (see discussion below); (2) due to preconditioning strategy, no free parameters are available those may influence the speed of convergence negatively such that u_α^* is expected to be of high accuracy after 100,000 iterations. Having u_α^* , each algorithm was applied until the relative error

$$\|u_\alpha^{(r)} - u_\alpha^*\| / \|u_\alpha^*\| \quad (10.67)$$

was below a predefined threshold ε (or a maximum number of iterations adopted for each algorithm individually was reached). The “ground truth” solutions for three different values of α are shown in Figure 10.3.




Fig. 10.2: Synthetic 2D PET data. *Left*: Exact object. *Middle*: Exact Radon data. *Right*: Simulated PET measurements via a Monte-Carlo simulation using one million events.




Fig. 10.3: The “ground truth” solutions for regularization parameter values $\alpha = 0.04$ (left), $\alpha = 0.08$ (middle), and $\alpha = 0.20$ (right).

Figures 10.4–10.8 show the performance evaluation of algorithms plotting the propagation of the relative error (10.67) in dependency on the number of iterations and CPU time in seconds. Since all algorithms have a specific set of unspecified parameters, different values of them are plotted to give a representative overall impression. The reason for showing the performance both in dependency on the number of iterations and CPU time is twofold: (1) in the presented case where the PET system matrix K is pre-computed and thus available explicitly, the evaluation of forward and backward projections is nearly negligible and TV relevant computations have the most contribution to the run time such that the CPU time will be a good indicator for algorithm’s performance; (2) in practically relevant cases where the forward and backward projections have to be computed in each iteration step implicitly and in general are computationally consuming, the number of iterations and thus the number of projection evaluations will be the crucial factor for algorithm’s efficiency. In the following, we individually discuss the behavior of algorithms observed for the regularization parameter $\alpha = 0.08$ (10.64) with the “ground truth” solution shown in Figure 10.3:

- **FB-EM-TV(-Nes83)**: The evaluation of FB-EM-TV based algorithms is shown in Figure 10.4. The major observation for any δ in (10.65) is that the inexact computations of TV proximal problems lead to a restrictive approximation of the “ground truth” solution where the approximation accuracy stagnates after a specific number of iterations depending on δ . In addition, it can also be observed that the relative error (10.67) becomes better with more accurate TV

proximal solutions (i.e., smaller δ) indicating that a decreasing sequence $\delta^{(r)}$ should be used to converge against the solution of (10.64) (see, e.g., [161, 175] for convergence analysis of inexact proximal gradient algorithms). However, as indicated in [119] and is shown in Figure 10.4, the choice of δ provides a trade-off between the approximation accuracy and computational cost such that the convergence rates proved in [161, 175] might be computationally not optimal. Another observation concerns the accelerated modification FB-EM-TV-Nes83. In Figure 10.4 we can observe that the performance of FB-EM-TV can actually be improved by FB-EM-TV-Nes83 regarding the number of iterations but only for smaller values of δ . One reason might be that using FB-EM-TV-Nes83 we have seen in our experiments that the gradient descent parameter $0 < \eta \leq 1$ [153] in (10.22) has to be chosen smaller with increased TV proximal accuracy (i.e., smaller δ). Since in such cases the effective regularization parameter value in each TV proximal problem is $\eta\alpha$, a decreasing η will result in poorer denoising properties increasing the inexactness of TV proximal operator. Recently, an (accelerated) inexact variable metric proximal gradient method was analyzed in [52] providing a theoretical view on such a type of methods.






Fig. 10.4: Performance of FB-EM-TV (dashed lines) and FB-EM-TV-Nes83 (solid lines) for different accuracy thresholds δ (10.65) within the TV proximal step. Evaluation of relative error (10.67) is shown as a function of the number of iterations (left) and CPU time in seconds (right).

- **(Precond-)CP-E:** In Figure 10.5, the algorithms CP-E and Precond-CP-E are evaluated. In contrast to FB-EM-TV-(Nes83), the approximated solution cannot be influenced by inexact computations such that a decaying behavior of relative error can be observed. The single parameter that affects the convergence rate is the dual steplength σ and we observe in Figure 10.5(a) that some values yield a fast initial convergence (see, e.g., $\sigma = 0.05$ and $\sigma = 0.1$), but are less suited to achieve fast asymptotic convergence and vice versa (see, e.g., $\sigma = 0.3$ and $\sigma = 0.5$). However, the plots in Figure 10.5(b) indicate that $\sigma \in [0.2, 0.3]$ may provide an acceptable trade-off between initial and asymptotic convergence in terms of the number of iterations and CPU time. Regarding the latter mentioned aspect we note that in case of CP-E the more natural setting of σ would be $\sigma = \sqrt{\|\nabla\|^2 + \|K\|^2}$ what is approximately 0.29 in our experiments providing acceptable trade-off between initial and asymptotic convergence. Finally, no

acceleration was observed in case of Precond-CP-E algorithm due to the regular structure of linear operators ∇ and K in our experiments such that the performance is comparable to CP-E with $\sigma = 0.5$ (see Figure 10.5(a)).



(a) Evaluation of relative error for fixed dual step sizes σ .



(b) Performance to get the relative error below the threshold ϵ as a function of dual step size σ .

Fig. 10.5: Performance of Precond-CP-E (dashed lines in ((a))) and CP-E (solid lines) for different dual step sizes σ . ((a)) Evaluation of relative error as a function of the number of iterations (left) and CPU time in seconds (right). ((b)) Required number of iterations (left) and CPU time (right) to get the relative error below a predefined threshold ϵ as a function of σ .

- **(Precond-)CP-SI:** In Figures 10.6 and 10.7, the evaluation of CP-SI and Precond-CP-SI is presented. Since a TV proximal operator has to be approximated in each iteration step, the same observations can be made as in case of FB-EM-TV that depending on δ the relative error stagnates after a specific number of iterations and that the choice of δ provides a trade-off between approximation accuracy and computational time (see Figure 10.6 for Precond-CP-SI). In addition, since the performance of CP-SI not only depends on δ but also on the dual steplength σ , the evaluation of CP-SI for different values of σ and two stopping values δ is shown in Figure 10.7. The main observation is that for smaller σ a better initial convergence can be achieved in terms of the number of iterations but results in less efficient performance regarding the CPU time. The reason is that the effective regularization parameter within the TV proximal problem is $\tau\alpha$ (see (10.22)) with $\tau = (\sigma\|K\|^2)^{-1}$ and a decreasing σ leads to an increasing TV denoising effort. Thus, in practically relevant cases, σ should be chosen optimally in a way balancing the required number of iterations and TV proximal computation.
- **PIDSplit+:** In Figure 10.8 the performance of PIDSplit+ is shown. It is well evaluated that the convergence of ADMM based algorithms is strongly depen-






Fig. 10.6: Performance of Precond-CP-SI for different accuracy thresholds δ (10.65) within the TV proximal step. Evaluation of relative error as a function of number of iterations (left) and CPU time in seconds (right).



(a) Performance with TV proximal accuracy $\delta = 0.05$ (65).



(b) Performance with TV proximal accuracy $\delta = 0.005$ (65).

Fig. 10.7: Performance of Precond-CP-SI (dashed lines) and CP-SI (solid lines) for different dual step sizes σ . Evaluation of relative error as a function of number of iterations (left) and CPU time in seconds (right) for accuracy thresholds $\delta = 0.05$ ((a)) and $\delta = 0.005$ ((b)) within the TV proximal problem.

dent on the augmented Lagrangian parameter γ (10.26) and that some values yield a fast initial convergence but are less suited to achieve a fast asymptotic convergence and vice versa. This behavior can also be observed in Figure 10.8 (see $\gamma = 30$ in upper row).

Finally, to get a feeling how the algorithms perform against each other, the required CPU time and number of projection evaluations to get the relative error (10.67)




Fig. 10.8: Performance of PIDSplit+ for fixed augmented Lagrangian penalty parameters γ (10.26). Evaluation of relative error as a function of number of iterations (left) and CPU time in seconds (right).

below a predefined threshold are shown in Table 10.1 for two different values of ε . The following observations can be made:

- The FB-EM-TV based algorithms are competitive in terms of required number of projection evaluations but have a higher CPU time due to the computation of TV proximal operators, in particular the CPU time strongly grows with decreasing ε since TV proximal problems have to be approximated with increased accuracy. However, in our experiments, a fixed δ was used in each TV denoising step and thus the performance can be improved utilizing the fact that a rough accuracy is sufficient at the beginning of the iteration sequence without influencing the performance regarding the number of projector evaluations negatively (cf. Figure 10.4). Thus, a proper strategy to iteratively decrease δ in (10.65) can strongly improve the performance of FB-EM-TV based algorithms.
- The CP-E algorithm is optimal in our experiments in terms of CPU time since the TV regularization is computed by the shrinkage operator and thus is simply to evaluate. However, this algorithm needs almost the highest number of projection evaluations that will result in a slow algorithm in practically relevant cases where the projector evaluations are highly computationally expensive.
- The PIDSplit+ algorithm is slightly poorer in terms of CPU time than CP-E but required a smaller number of projector evaluations. However, we remind that this performance may probably be improved since two PCG iterations were used in our experiments and thus two forward and backward projector evaluations are required in each iteration step of PIDSplit+ method. Thus, if only one PCG step is used, the CPU time and number of projector evaluations can be decreased leading to a better performing algorithm. However, in the latter case, the total number of iteration steps might be increased since a poorer approximation of (10.66) will be performed if only one PCG step is used. Another opportunity to improve the performance of PIDSplit+ algorithm is to use the proximal ADMM strategy described in Section 6.4, namely, to remove $K^T K$ from (10.66). That will result in only a single evaluation of forward and backward projectors in each iteration step but may lead to an increased number of total number of algorithm iterations.

Table 10.1: Performance evaluation of algorithms described above for $\alpha = 0.08$ (see u_α^* in Figure 10.3 (middle)). The table displays the CPU time in seconds and required number of forward and backward projector evaluations (K/K^T) to get the relative error (10.67) below the error tolerance ε . For each algorithm the best performance regarding the CPU time and K/K^T evaluations are shown where \parallel means that the value coincides with the value directly above.

		$\varepsilon = 0.05$		$\varepsilon = 0.005$	
		K/K^T	CPU	K/K^T	CPU
FB-EM-TV	(best K/K^T)	20	40.55	168	4999.71
\parallel	(best CPU)	\parallel	\parallel	230	3415.74
FB-EM-TV-Nes83	(best K/K^T)	15	14.68	231	308.74
\parallel	(best CPU)	\parallel	\parallel	\parallel	\parallel
CP-E	(best K/K^T)	48	4.79	696	69.86
\parallel	(best CPU)	\parallel	\parallel	\parallel	\parallel
CP-SI	(best K/K^T)	22	198.07	456	1427.71
\parallel	(best CPU)	25	23.73	780	1284.56
PIDSplit+	(best K/K^T)	30	7.51	698	179.77
\parallel	(best CPU)	\parallel	\parallel	\parallel	\parallel

Finally, to study the algorithm's stability regarding the choice of regularization parameter α , we have run the algorithms for two additional values of α using the parameters shown the best performance in Table 10.1. The additional penalty parameters include a slightly under-smoothed and over-smoothed result respectively as shown in Figure 10.3 and the evaluation results are shown in Tables 10.2 and 10.3. In the following we describe the major observations:

- The FB-EM-TV method has the best efficiency in terms of projector evaluations, independently from the penalty parameter α , but has the disadvantage of solving a TV proximal problem in each iteration step which get harder to solve with increasing smoothing level (i.e., larger α) leading to a negative computational time. The latter observation holds also for the CP-SI algorithm. In case of a rough approximation accuracy (see Table 10.2), the FB-EM-TV-Nes83 scheme is able to improve the overall performance, respectively at least the computational time for higher accuracy in Table 10.3, but here the damping parameter η in (10.22) has to be chosen carefully to ensure the convergence (cf. Table 10.2 and 10.3 in case of $\alpha = 0.2$). Additionally based on Table 10.1, a proper choice of η is dependent not only on α but also on the inner accuracy of TV proximal problems.
- In contrast to FB-EM-TV and CP-SI, the remaining algorithms provide a superior computational time due to the solution of TV related steps by the shrinkage formula but show a strongly increased requirements on projector evaluations across all penalty parameters α . In addition, the performance of these algorithms is strongly dependent on the proper setting of free parameters (σ in case

of CP-E and γ in PIDSplit+) which unfortunately are able to achieve only a fast initial convergence or a fast asymptotic convergence. Thus different parameter settings of σ and γ were used in Tables 10.2 and 10.3.

Table 10.2: Performance evaluation for different values of α (see Figure 10.3). The table displays the CPU time in seconds and required number of forward and backward projector evaluations (K/K^T) to get the relative error (10.67) below the error tolerance $\varepsilon = 0.05$. For each α , the algorithms were run using the following parameters: FB-EM-TV ($\delta = 0.1$), FB-EM-TV-Nes83 ($\delta = 0.1, \eta = 0.5$), CP-E ($\sigma = 0.07$), CP-SI ($\delta = 0.1, \sigma = 0.05$), PIDSplit+ ($\gamma = 10$), which were chosen based on the “best” performance regarding K/K^T for $\varepsilon = 0.05$ in Table 10.1.

	$\alpha = 0.04$		$\alpha = 0.08$		$\alpha = 0.2$	
	K/K^T	CPU	K/K^T	CPU	K/K^T	CPU
FB-EM-TV	28	16.53	20	40.55	19	105.37
FB-EM-TV-Nes83	17	5.26	15	14.68	-	-
CP-E	61	6.02	48	4.79	51	5.09
CP-SI	-	-	25	23.73	21	133.86
PIDSplit+	32	8.08	30	7.51	38	9.7

Table 10.3: Performance evaluation for different values of α (see Figure 10.3) as in Table 10.2 but for $\varepsilon = 0.005$ and using the following parameters: FB-EM-TV ($\delta = 0.005$), FB-EM-TV-Nes83 ($\delta = 0.005, \eta = 0.05$), CP-E ($\sigma = 0.2$), CP-SI ($\delta = 0.005, \sigma = 0.3$), PIDSplit+ ($\gamma = 3$).

	$\alpha = 0.04$		$\alpha = 0.08$		$\alpha = 0.2$	
	K/K^T	CPU	K/K^T	CPU	K/K^T	CPU
FB-EM-TV	276	2452.14	168	4999.71	175	12612.7
FB-EM-TV-Nes83	512	222.98	231	308.74	-	-
CP-E	962	94.57	696	69.86	658	65.42
CP-SI	565	1117.12	456	1427.71	561	7470.18
PIDSplit+	932	239.35	698	179.77	610	158.94

8.2 Spectral X-Ray CT

Conventional X-ray CT is based on recording changes in the X-ray intensity due to attenuation of X-ray beams traversing the scanned object and has been applied in clinical practice for decades. However, the transmitted X-rays carry more information than just intensity changes since the attenuation of an X-ray depends strongly

on its energy [2, 109]. It is well understood that the transmitted energy spectrum contains valuable information about the structure and material composition of the imaged object and can be utilized to better distinguish different types of absorbing material, such as varying tissue types or contrast agents. But the detectors employing in traditional CT systems provide an integral measure of absorption over the transmitted energy spectrum and thus eliminate spectral information [39, 158]. Even so, spectral information can be obtained by using different input spectra [77, 193] or using the concept of dual-layer (integrating) detectors [40]. This has limited the practical usefulness of energy-resolving imaging, also referred to as spectral CT, to dual energy systems. Recent advances in detector technology towards binned photon-counting detectors have enabled a new generation of detectors that can measure and analyze incident photons individually [39] providing the availability of more than two spectral measurements. This development has led to a new imaging method named K-edge imaging [113] that can be used to selectively and quantitatively image contrast agents loaded with K-edge materials [75, 137]. For a compact overview on technical and practical aspects of spectral CT we refer to [39, 158].

Two strategies have been proposed to reconstruct material specific images from spectral CT projection data and we refer to [158] for a compact overview. Either of them is a projection-based material decomposition with a subsequent image reconstruction. This means that in the first step, estimates of material-decomposed sinograms are computed from the energy-resolved measurements, and in the second step, material images are reconstructed from the decomposed material sinograms. A possible decomposition method to estimate the material sinograms f_l , $l = 1, \dots, L$, from the acquired data is a maximum-likelihood estimator assuming a Poisson noise distribution [150], where L is the number of materials considered. An accepted noise model for line integrals f_l is a multivariate Gaussian distribution [158, 159] leading to a penalized weighted least squares (PWLS) estimator to reconstruct material images u_l :

$$\frac{1}{2} \|f - (I_L \otimes K)u\|_{\Sigma^{-1}}^2 + \alpha R(u) \rightarrow \min_u, \quad \alpha > 0, \quad (10.68)$$

where $f = (f_1^T, \dots, f_L^T)^T$, $u = (u_1^T, \dots, u_L^T)^T$, I_L denotes the $L \times L$ identity matrix, \otimes represents the Kronecker product, and K is the forward projection operator. The given block matrix Σ is the covariance matrix representing the (multivariate) Gaussian distribution, where the off-diagonal block elements describe the inter-sinogram correlations, and can be estimated, e.g., using the inverse of the Fisher information matrix [149, 159]. Since f is computed from a common set of measurements, the correlation of the decomposed data is very high and thus a significant improvement can in fact be expected intuitively by exploiting the fully populated covariance matrix Σ in (10.68). In the following, we exemplarily show reconstruction results on spectral CT data where (10.68) was solved by a proximal ADMM algorithm with a material independent total variation penalty function R as discussed in [155]. For a discussion why ADMM based methods are more preferable for PWLS problems in X-ray CT than, e.g., gradient descent based techniques, we refer to [145].

Figures 10.10 and 10.11 show an example for a statistical image reconstruction method applied to K-edge imaging. A numerical phantom as shown in Figure 10.9 was employed in a spectral CT simulation study assuming a photon-counting detector. Using an analytical spectral attenuation model, spectral measurements were computed. The assumed X-ray source spectrum and detector response function of the photon-counting detector were identical to those employed in a prototype spectral CT scanner described in [160]. The scan parameters were set to tube voltage 130 kVp, anode current $200 \mu\text{A}$, detector width/height 946.38/1.14 mm, number of columns 1024, source-to-isocenter/-detector distance 570/1040 mm, views per turn 1160, time per turn 1 s, and energy thresholds to 25, 46, 61, 64, 76, and 91 keV. The spectral data were then decomposed into ‘photo-electric absorption’, ‘Compton effect’, and ‘ytterbium’ by performing a maximum-likelihood estimation [150]. By computing the covariance matrix Σ of the material decomposed sinograms via the Fisher information matrix [149, 159] and treating the sinograms as the mean and Σ as the variance of a Gaussian random vector, noisy material sinograms were computed. Figures 10.10 and 10.11 show material images that were then reconstructed using the traditional filtered backprojection (upper row) and proximal ADMM algorithm as described in [155] (middle and lower row). In the latter case, two strategies were performed: (1) keeping only the diagonal block elements of Σ in (10.68) and thus neglecting cross-correlations and decoupling the reconstruction of material images (middle row); (2) using the fully populated covariance matrix Σ in (10.68) such that all material images have to be reconstructed jointly (lower row). The results suggest, best visible in the K-edge images in Figure 10.11, that the iterative reconstruction method, which exploits knowledge of the inter-sinogram correlations, produces images that possess a better reconstruction quality. For comparison of iterative reconstruction strategies, the regularization parameters were chosen manually so the reconstructed images possessed approximately the same variance within the region indicated by the dotted circle in Figure 10.9. Further (preliminary) results that demonstrate advantages of exploiting inter-sinogram correlations on computer-simulated and experimental data in spectral CT can be found in [155, 189].




Fig. 10.9: Software thorax phantom comprising sternum, ribs, lungs, vertebrae, and one circle and six ellipsoids containing different concentrations of K-edge material ytterbium [137]. The phantom was used to simulate spectral CT measurements with a six-bin photon-counting detector.




Fig. 10.10: Reconstructions based on the thorax phantom (see Figure 10.9) using the traditional filtered backprojection with Shepp-Logan filter (upper row) and a proximal ADMM algorithm as described in [155] (middle and lower row). The middle row shows results based on (10.68) neglecting cross-correlations between the material decomposed sinograms and lower row using the fully populated covariance matrix Σ . The material images show the results for the ‘Compton effect’ (left column) and ‘photo-electric absorption’ (right column). The K-edge material ‘ytterbium’ is shown in Figure 10.11.

Acknowledgements

The authors thank Frank Wübbeling (University of Münster, Germany) for providing the Monte-Carlo simulation for the synthetic 2D PET data. The authors also thank Thomas Koehler (Philips Technologie GmbH, Innovative Technologies, Hamburg, Germany) and Simon Setzer (G-Research, London) for comments that improved the manuscript and Eva-Maria Brinkmann (University of Münster, Germany) for the careful proofreading. The work on spectral CT was performed when A. Sawatzky was with the Computational Bioimaging Laboratory at Washington University in St. Louis, USA, and was supported in part by NIH award EB009715 and funded from Philips Research North America. G. Steidl’s work was partially supported within the BMBF Project 05M13UKA.




Fig. 10.11: Reconstructions of the K-edge material ‘ytterbium’ using the thorax phantom shown in Figure 10.9. For details see Figure 10.10. To recognize the differences, the maximal intensity value of original reconstructed images shown in left column was set down in the right column.

References

1. J. F. P.-J. Abascal, J. Chamorro-Servent, J. Aguirre, S. Arridge, T. Correia, J. Ripoll, J. J. Vaquero, and M. Desco. Fluorescence diffuse optical tomography using the split Bregman method. *Med. Phys.*, 38:6275, 2011.
2. R. E. Alvarez and A. Macovski. Energy-selective reconstructions in X-ray computerized tomography. *Phys. Med. Biol.*, 21(5):733–744, 1976.
3. S. Anthoine, J.-F. Aujol, Y. Boursier, and C. Mélot. On the efficiency of proximal methods in CBCT and PET. In *Proc. IEEE Int. Conf. Image Proc. (ICIP)*, 2011.
4. S. Anthoine, J.-F. Aujol, Y. Boursier, and C. Mélot. Some proximal methods for CBCT and PET. In *Proc. SPIE (Wavelets and Sparsity XIV)*, volume 8138, 2011.
5. K. J. Arrow, L. Hurwitz, and H. Uzawa. *Studies in Linear and Nonlinear Programming*. Stanford University Press, 1958.
6. H. Attouch and J. Bolte. On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Math. Program.*, 116(1–2):5–16, 2009.
7. H. Attouch, J. Bolte, and B. F. Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods. *Math. Program. Series A*, 137(1–2):91–129, 2013.
8. J.-P. Aubin. *Optima and Equilibria: An Introduction to Nonlinear Analysis*. Springer, Berlin, Heidelberg, New York, 2nd edition, 2003.

9. M. Bachmayr and M. Burger. Iterative total variation schemes for nonlinear inverse problems. *Inverse Problems*, 25(10):105004, 2009.
10. E. Bae, J. Yuan, and X.-C. Tai. Global minimization for continuous multiphase partitioning problems using a dual approach. *International Journal of Computer Vision*, 92(1):112–129, 2011.
11. J. Barzilai and J. M. Borwein. Two-point step size gradient methods. *IMA J. Numer. Anal.*, 8(1):141–148, 1988.
12. H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, New York, 2011.
13. A. Beck and M. Teboulle. Fast gradient-based algorithms for constrained total variation image denoising and deblurring. *SIAM J. Imag. Sci.*, 2:183–202, 2009.
14. S. Becker, J. Bobin, and E. J. Candès. NESTA: a fast and accurate first-order method for sparse recovery. *SIAM J. Imag. Sci.*, 4(1):1–39, 2011.
15. M. Benning, L. Gladden, D. Holland, C.-B. Schönlieb, and T. Valkonen. Phase reconstruction from velocity-encoded MRI measurement - A survey of sparsity-promoting variational approaches. *J. Magn. Reson.*, 238:26–43, 2014.
16. M. Benning, P. Heins, and M. Burger. A solver for dynamic PET reconstructions based on forward-backward-splitting. In *AIP Conf. Proc.*, volume 1281, pages 1967–1970, 2010.
17. D. P. Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods*. Academic Press, New York, 1982.
18. D. P. Bertsekas. Incremental proximal methods for large scale convex optimization. *Math. Program., Ser. B*, 129(2):163–195, 2011.
19. J. M. Bioucas-Dias and M. A. T. Figueiredo. Multiplicative noise removal using variable splitting and constrained optimization. *IEEE Trans. Image Process.*, 19(7):1720–1730, 2010.
20. A. Björck. *Least Squares Problems*. SIAM, Philadelphia, 1996.
21. D. Boley. Local linear convergence of the alternating direction method of multipliers on quadratic or linear programs. *SIAM J. Optim.*, 2014.
22. J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for non-convex and nonsmooth problems. *Math. Program., Series A*, 2013.
23. S. Bonettini and V. Ruggiero. On the convergence of primal-dual hybrid gradient algorithms for total variation image restoration. *J. Math. Imaging Vis.*, 44:236–253, 2012.
24. J. F. Bonnans, J. C. Gilbert, C. Lemaréchal, and C. A. Sagastizábal. A family of variable metric proximal methods. *Mathematical Programming*, 68:15–47, 1995.
25. R. I. Boş and C. Hendrich. A Douglas-Rachford type primal-dual method for solving inclusions with mixtures of composite and parallel-sum type monotone operators. *SIAM Journal on Optimization*, 23(4):2541–2565, 2013.
26. R. I. Boş and C. Hendrich. Convergence analysis for a primal-dual monotone + skew splitting algorithm with applications to total variation minimization. *Journal of Mathematical Imaging and Vision*, 49(3):551–568, 2014.
27. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):101–122, 2011.
28. K. Bredies, K. Kunisch, and T. Pock. Total generalized variation. *SIAM Journal on Imaging Sciences*, 3(3):492–526, 2010.
29. F. E. Browder and W. V. Petryshyn. The solution by iteration of nonlinear functional equations in Banach spaces. *Bulletin of the American Mathematical Society*, 72:571–575, 1966.
30. C. Brune, A. Sawatzky, and M. Burger. Primal and dual Bregman methods with application to optical nanoscopy. *International Journal of Computer Vision*, 92(2):211–229, 2010.
31. M. Burger and S. Osher. A guide to the TV zoo. In *Level Set and PDE Based Reconstruction Methods in Imaging*, pages 1–70. Springer, 2013.
32. M. Burger, E. Resmerita, and L. He. Error estimation for Bregman iterations and inverse scale space methods in image restoration. *Computing*, 81(2–3):109–135, 2007.
33. J. V. Burke and M. Qian. A variable metric proximal point algorithm for monotone operators. *SIAM Journal on Control and Optimization*, 37:353–375, 1999.

34. F. Büther, M. Dawood, L. Stegger, F. Wübbeling, M. Schäfers, O. Schober, and K. P. Schäfers. List mode-driven cardiac and respiratory gating in PET. *J. Nucl. Med.*, 50(5):674–681, 2009.
35. D. Butnariu and A. N. Iusem. *Totally convex functions for fixed points computation and infinite dimensional optimization*, volume 40 of *Applied Optimization*. Kluwer, Dordrecht, 2000.
36. C. Byrne. A unified treatment of some iterative algorithms in signal processing and image reconstruction. *Inverse Problems*, 20:103–120, 2004.
37. J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.
38. J.-F. Cai, S. Osher, and Z. Shen. Convergence of the linearized Bregman iteration for ℓ_1 -norm minimization. *Mathematics of Computation*, 78(268):2127–2136, 2009.
39. J. Cammin, J. S. Iwanczyk, and K. Taguchi. *Emerging Imaging Technologies in Medicine*, chapter Spectral/Photo-Counting Computed Tomography, pages 23–39. CRC Press, 2012.
40. R. Carmi, G. Naveh, and A. Altman. Material separation with dual-layer CT. In *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, pages 1876–1878, 2005.
41. A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock. An introduction to total variation for image analysis. In *Theoretical Foundations and Numerical Methods for Sparse Recovery*, volume 9 of *Radon Series Compl. Appl. Math.*, pages 263–340. Walter de Gruyter, 2010.
42. A. Chambolle and T. Pock. Diagonal preconditioning for first order primal-dual algorithms in convex optimization. *ICCV*, pages 1762–1769, 2011.
43. A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
44. T. F. Chan and R. Glowinski. Finite element approximation and iterative solution of a class of mildly non-linear elliptic equations. Technical report, STAN-CS-78-674, Stanford University, 1978.
45. R. Chartrand, E. Y. Sidky, and X. Pan. Nonconvex compressive sensing for X-ray CT: an algorithm comparison. In *Asilomar Conference on Signals, Systems, and Computers*, 2013.
46. C. Chaux, M. El-Gheche, J. Farah, J. Pesquet, and B. Popescu. A parallel proximal splitting method for disparity estimation from multicomponent images under illumination variation. *Journal of Mathematical Imaging and Vision*, 47(3):1–12, 2012.
47. C. H. Chen, B. S. He, Y. Y. Ye, and X. M. Yuan. The direct extension of ADMM for multi-block convex minimization problems is not necessarily convergent. *Mathematical Programming*, to appear.
48. G. Chen and M. Teboulle. A proximal-based decomposition method for convex minimization problems. *Mathematical Programming*, 64:81–101, 1994.
49. G. H.-G. Chen and R. T. Rockafellar. Convergence rates in forward-backward splitting. *SIAM Journal on Optimization*, 7:421–444, 1997.
50. G. Chierchia, N. Pustelnik, J.-C. Pesquet, and B. Pesquet-Popescu. Epigraphical projection and proximal tools for solving constrained convex optimization problems: Part I. arXiv preprint arXiv:1210.5844 (2014).
51. K. Choi, J. Wang, L. Zhu, T.-S. Suh, S. Boyd, and L. Xing. Compressed sensing based cone-beam computed tomography reconstruction with a first-order method. *Med. Phys.*, 37(9):5113–5125, 2010.
52. E. Chouzenoux, J.-C. Pesquet, and A. Repetti. Variable metric forward-backward algorithm for minimizing the sum of a differentiable function and a convex function. *J. Optim. Theory Appl.*, 2013.
53. P. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 185–212. Springer, 2011.
54. P. Combettes and J.-C. Pesquet. Primal-dual splitting algorithm for solving inclusions with mixture of composite, Lipschitzian, and parallel-sum type monotone operators. *Set-Valued and Variational Analysis*, 20(2):307–330, 2012.

55. P. L. Combettes. Solving monotone inclusions via compositions of nonexpansive averaged operators. *Optimization*, 53(5–6):475–504, 2004.
56. P. L. Combettes and J.-C. Pesquet. Proximal thresholding algorithm for minimization over orthonormal bases. *SIAM Journal on Optimization*, 18(4):1351–1376, 2007.
57. P. L. Combettes and B. C. Vu. Variable metric forward-backward splitting with applications to monotone inclusions in duality. *Optimization*, pages 1–30, 2012.
58. L. Condat. A primal-dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms. *J. Optim. Theory Appl.*, 158(2):460–479, 2013.
59. W. Cong, J. Yang, and G. Wang. Differential phase-contrast interior tomography. *Phys. Med. Biol.*, 57:2905–2914, 2012.
60. J. Dahl, P. J. Hansen, S. H. Jensen, and T. L. Jensen. Algorithms and software for total variation image reconstruction via first order methods. *Numerical Algorithms*, 53:67, 2010.
61. I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 51:1413–1541, 2004.
62. I. Daubechies, M. Fornasier, and I. Loris. Accelerated projected gradient methods for linear inverse problems with sparsity constraints. *The Journal of Fourier Analysis and Applications*, 14(5–6):764–792, 2008.
63. C. Davis. All convex invariant functions of Hermitian matrices. *Archive in Mathematics*, 8:276–278, 1957.
64. D. Davis and W. Yin. Convergence rate analysis of several splitting schemes. In: R. Glowinski, S. Osher, W. Yin (eds.) *Splitting Methods in Communication and Imaging*, Science and Engineering. Springer, 2016.
65. D. Davis and W. Yin. Faster convergence rates of relaxed Peaceman-Rachford and ADMM under regularity assumptions. *ArXiv Preprint 1407.5210*, 2014.
66. N. Dey, L. Blanc-Feraud, C. Zimmer, P. Roux, Z. Kam, J.-C. Olivo-Marin, and J. Zerubia. Richardson-Lucy algorithm with total variation regularization for 3D confocal microscope deconvolution. *Microsc. Res. Tech.*, 69:260–266, 2006.
67. J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the l_1 -ball for learning in high dimensions. In *ICML '08 Proceedings of the 25th International Conference on Machine Learning*, ACM New York, 2008.
68. J. Eckstein and D. P. Bertsekas. An alternating direction method for linear programming. *Tech. Report MIT Lab. for Info. and Dec. Sys.*, 1990.
69. J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55:293–318, 1992.
70. E. Esser. Applications of Lagrangian-based alternating direction methods and connections to split Bregman. Technical report, UCLA Computational and Applied Mathematics, March 2009.
71. E. Esser, X. Zhang, and T. F. Chan. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM J Imag Sci*, 3(4):1015–1046, 2010.
72. F. Facchinei and J.-S. Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*, volume II. Springer, New York, 2003.
73. J. A. Fessler. Conjugate-gradient preconditioning methods: Numerical results. Technical Report 303, Commun. Signal Process. Lab., Dept. Elect. Eng. Comput. Sci., Univ. Michigan, Ann Arbor, MI, Jan. 1997. available from <http://web.eecs.umich.edu/~fessler/>.
74. J. A. Fessler and S. D. Booth. Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction. *IEEE Trans. Image Process.*, 8(5):688–699, 1999.
75. S. Feuerlein, E. Roessl, R. Proksa, G. Martens, O. Klass, M. Jeltsch, V. Rasche, H.-J. Brambs, M. H. K. Hoffmann, and J.-P. Schlomka. Multienergy photon-counting K-edge imaging: Potential for improved luminal depiction in vascular imaging. *Radiology*, 249(3):1010–1016, 2008.

76. M. Figueiredo and J. Bioucas-Dias. Deconvolution of Poissonian images using variable splitting and augmented Lagrangian optimization. In *IEEE Workshop on Statistical Signal Processing*, Cardiff, 2009.
77. T. G. Flohr, C. H. McCollough, H. Bruder, M. Petersilka, K. Gruber, C. Süß, M. Grasruck, K. Stierstorfer, B. Krauss, R. Raupach, A. N. Primak, A. Küttner, S. Achenbach, C. Becker, A. Kopp, and B. M. Ohnesorge. First performance evaluation of a dual-source CT (DSCT) system. *Eur. Radiol.*, 16:256–268, 2006.
78. M. Fornasier. *Theoretical Foundations and Numerical Methods for Sparse Recovery*, volume 9. Walter de Gruyter, 2010.
79. G. Frassoldati, L. Zanni, and G. Zanghirati. New adaptive stepsize selections in gradient methods. *Journal of Industrial and Management Optimization*, 4(2):299–312, 2008.
80. M. Freiberger, C. Clason, and H. Scharfetter. Total variation regularization for nonlinear fluorescence tomography with an augmented Lagrangian splitting approach. *Appl. Opt.*, 49(19):3741–3747, 2010.
81. K. Frick, P. Marnitz, and A. Munk. Statistical multiresolution estimation for variational imaging: With an application in Poisson-biophotonics. *J. Math. Imaging Vis.*, 46:370–387, 2013.
82. D. Gabay. Applications of the method of multipliers to variational inequalities. In M. Fortin and R. Glowinski, editors, *Augmented Lagrangian Methods: Applications to the Solution of Boundary Value Problems*, chapter IX, pages 299–340. North-Holland, Amsterdam, 1983.
83. D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximations. *Computer and Mathematics with Applications*, 2:17–40, 1976.
84. H. Gao, S. Osher, and H. Zhao. *Mathematical Modeling in Biomedical Imaging II: Optical, Ultrasound, and Opto-Acoustic Tomographies*, chapter Quantitative Photoacoustic Tomography, pages 131–158. Springer, 2012.
85. H. Gao, H. Yu, S. Osher, and G. Wang. Multi-energy CT based on a prior rank, intensity and sparsity model (PRISM). *Inverse Problems*, 27(11):115012, 2011.
86. R. Glowinski. On alternating direction methods of multipliers: a historical perspective. In *Modeling, Simulation and Optimization for Science and Technology*, pages 59–82. Springer, 2014.
87. R. Glowinski and P. Le Tallec. *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*, volume 9 of *SIAM Studies in Applied and Numerical Mathematics*. SIAM, Philadelphia, 1989.
88. R. Glowinski and A. Marroco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de Dirichlet non linéaires. *Revue française d’automatique, informatique, recherche opérationnelle. Analyse numérique*, 9(2):41–76, 1975.
89. D. Goldfarb and K. Scheinberg. Fast first-order methods for composite convex optimization with line search. *SIAM Journal on Imaging Sciences*, 2011.
90. T. Goldstein, X. Bresson, and S. Osher. Geometric applications of the split Bregman method: Segmentation and surface reconstruction. *J. Sci. Comput.*, 45:272–293, 2010.
91. T. Goldstein and S. Osher. The split Bregman method for L1-regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009.
92. B. Goris, W. Van den Broek, K. J. Batenburg, H. H. Mezerji, and S. Bals. Electron tomography based on a total variation minimization reconstruction technique. *Ultramicroscopy*, 113:120–130, 2012.
93. O. Güler. New proximal point algorithms for convex minimization. *SIAM J. Optim.*, 2(4):649–664, 1992.
94. S. Harizanov, J.-C. Pesquet, and G. Steidl. Epigraphical projection for solving least squares Anscombe transformed constrained optimization problems. In A. K. et al., editor, *Scale-Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science, SSVM 2013, LNCS 7893*, pages 125–136, Berlin, 2013. Springer.
95. B. He, L.-Z. Liao, D. Han, and H. Yang. A new inexact alternating directions method for monotone variational inequalities. *Math. Program., Ser. A*, 92(1):103–118, 2002.

96. B. He and H. Yang. Some convergence properties of the method of multipliers for linearly constrained monotone variational operators. *Operation Research Letters*, 23:151–161, 1998.
97. B. He and X. Yuan. On the $\mathcal{O}(1/n)$ convergence rate of the Douglas-Rachford alternating direction method. *SIAM Journal on Numerical Analysis*, 2:700–709, 2012.
98. B. S. He, H. Yang, and S. L. Wang. Alternating direction method with self-adaptive penalty parameters for monotone variational inequalities. *J. Optimiz. Theory App.*, 106(2):337–356, 2000.
99. S. W. Hell. Toward fluorescence nanoscopy. *Nat. Biotechnol.*, 21(11):1347–1355, 2003.
100. S. W. Hell. Far-field optical nanoscopy. *Science*, 316(5828):1153–1158, 2007.
101. M. R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4:303–320, 1969.
102. M. Hong and Z. Q. Luo. On linear convergence of the alternating direction method of multipliers. *Arxiv preprint 1208.3922*, 2012.
103. J. Huang, S. Zhang, and D. Metaxas. Efficient MR image reconstruction for compressed MR imaging. *Med. Image Anal.*, 15:670–679, 2011.
104. X. Jia, Y. Lou, R. Li, W. Y. Song, and S. B. Jiang. GPU-based fast cone beam CT reconstruction from undersampled and noisy projection data via total variation. *Med. Phys.*, 37(4):1757–1760, 2010.
105. B. Kaltenbacher, A. Neubauer, and O. Scherzer. *Iterative regularization methods for nonlinear ill-posed problems*, volume 6. Walter de Gruyter, 2008.
106. S. H. Kang, B. Shafei, and G. Steidl. Supervised and transductive multi-class segmentation using p -Laplacians and RKHS methods. *J. Visual Communication and Image Representation*, 25(5):1136–1148, 2014.
107. K. C. Kiwiel. Free-steering relaxation methods for problems with strictly convex costs and linear constraints. *Mathematics of Operations Research*, 22(2):326–349, 1997.
108. K. C. Kiwiel. Proximal minimization methods with generalized Bregman functions. *SIAM Journal on Control and Optimization*, 35(4):1142–1168, 1997.
109. G. F. Knoll. *Radiation Detection and Measurement*. Wiley, 3rd edition, 2000.
110. N. Komodakis and J.-C. Pesquet. Playing with duality: an overview of recent primal-dual approaches for solving large-scale optimization problems. *ArXiv Preprint arXiv:1406.5429*, 2014.
111. S. Kontogiorgis and R. R. Meyer. A variable-penalty alternating directions method for convex optimization. *Math. Program.*, 83(1–3):29–53, 1998.
112. M. A. Krasnoselskii. Two observations about the method of successive approximations. *Uspekhi Matematicheskikh Nauk*, 10:123–127, 1955. In Russian.
113. R. A. Kruger, S. J. Riederer, and C. A. Mistretta. Relative properties of tomography, K-edge imaging, and K-edge tomography. *Med. Phys.*, 4(3):244–249, 1977.
114. M.-J. Lai and W. Yin. Augmented ℓ_1 and nuclear-norm models with a globally linearly convergent algorithm. *SIAM Journal on Imaging Sciences*, 6(2):1059–1091, 2013.
115. J. Lellmann, J. Kappes, J. Yuan, F. Becker, and C. Schnörr. Convex multi-class image labeling with simplex-constrained total variation. In X.-C. Tai, K. Morken, M. Lysaker, and K.-A. Lie, editors, *Scale Space and Variational Methods, volume 5567 of LNCS*, volume 5567 of *Lecture Notes in Computer Science*, pages 150–162. Springer, 2009.
116. P. L. Lions and B. Mercier. Splitting algorithms for the sum of two linear operators. *SIAM Journal on Numerical Analysis*, 16:964–976, 1979.
117. M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.*, 58:1182–1195, 2007.
118. S. Ma, W. Y. Y. Zhang, and A. Chakraborty. An efficient algorithm for compressed MR imaging using total variation and wavelets. In *Proc. IEEE Comput. Vision Pattern Recognit.*, 2008.
119. P. Machart, S. Anthoine, and L. Baldassarre. Optimal computational trade-off of inexact proximal methods. *arXiv preprint arXiv:1210.5034*, 2012.
120. W. R. Mann. Mean value methods in iteration. *Proceedings of the American Mathematical Society*, 16(4):506–510, 1953.

121. B. Martinet. Régularisation d'inéquations variationnelles par approximations successives. *Revue Française d'Informatique et de Recherche Operationelle*, 4(3):154–158, 1970.
122. A. Mehranian, A. Rahmim, M. R. Ay, F. Kotasidis, and H. Zaidi. An ordered-subsets proximal preconditioned gradient algorithm for edge-preserving PET image reconstruction. *Med. Phys.*, 40(5):052503, 2013.
123. J. Müller, C. Brune, A. Sawatzky, T. Kösters, F. Wübbeling, K. Schäfers, and M. Burger. Reconstruction of short time PET scans using Bregman iterations. In *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, 2011.
124. F. Natterer and F. Wübbeling. *Mathematical Methods in Image Reconstruction*. SIAM, 2001.
125. A. S. Nemirovsky and D. B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. J. Wiley & Sons, Ltd., 1983.
126. Y. Nesterov. *Introductory Lectures on Convex Optimization - A Basic Course*, volume 87 of *Applied Optimization*. Springer US, 2004.
127. Y. Nesterov. Gradient methods for minimizing composite functions. *Math. Program., Series B*, 140(1):125–161, 2013.
128. Y. E. Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376, 1983.
129. Y. E. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103:127–152, 2005.
130. H. Nien and J. A. Fessler. Fast X-ray CT image reconstruction using the linearized augmented Lagrangian method with ordered subsets. arXiv preprint arXiv:1402.4381, 2014.
131. M. Nilchian, C. Vonesch, P. Modregger, M. Stampanoni, and M. Unser. Fast iterative reconstruction of differential phase contrast X-ray tomograms. *Optics Express*, 21(5):5511–5528, 2013.
132. P. Ochs, Y. Chen, T. Brox, and T. Pock. iPiano: Inertial proximal algorithm for nonconvex optimization. *SIAM J Imaging Sci*, 7(2):1388–1419, 2014.
133. Z. Opial. Weak convergence of a sequence of successive approximations for nonexpansive mappings. *Bulletin of the American Mathematical Society*, 73:591–597, 1967.
134. J. M. Ortega and W. C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. SIAM, New York, 1970.
135. S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An iterative regularization method for the total variation based image restoration. *Multiscale Modeling and Simulation*, 4:460–489, 2005.
136. S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An iterative regularization method for total variation-based image restoration. *Multiscale Modeling & Simulation*, 4(2):460–489, 2005.
137. D. Pan, C. O. Schirra, A. Senpan, A. H. Schmieder, A. J. Stacy, E. Roessl, A. Thran, S. A. Wickline, R. Proksa, and G. M. Lanza. An early investigation of ytterbium nanocolloids for selective and quantitative “multicolor” spectral CT imaging. *ACS Nano*, 6(4):3364–3370, 2012.
138. L. A. Parente, P. A. Lotito, and M. V. Solodov. A class of inexact variable metric proximal point algorithms. *SIAM J. Optim.*, 19(1):240–260, 2008.
139. N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.
140. T. Pock, A. Chambolle, D. Cremers, and H. Bischof. A convex relaxation approach for computing minimal partitions. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 810–817, 2009.
141. M. J. D. Powell. A method for nonlinear constraints in minimization problems. *Optimization*, 1972.
142. K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: Sensitivity encoding for fast MRI. *Magn. Reson. Med.*, 42:952–962, 1999.
143. N. Pustelnik, C. Chauv, J.-C. Pesquet, and C. Comtat. Parallel algorithm and hybrid regularization for dynamic PET reconstruction. In *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, 2010.

144. S. Ramani and J. A. Fessler. Parallel MR image reconstruction using augmented Lagrangian methods. *IEEE Trans. Med. Imag.*, 30(3):694–706, 2011.
145. S. Ramani and J. A. Fessler. A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction. *IEEE Trans. Med. Imag.*, 31(3):677–688, 2012.
146. R. T. Rockafellar. Augmented Lagrangians and applications of the proximal point algorithm in convex programming. *Math. Oper. Res.*, 1(2):97–116, 1976.
147. R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14:877–898, 1976.
148. R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 10 edition, 1997.
149. E. Roessl and C. Herrmann. Cramér-Rao lower bound of basis image noise in multiple-energy x-ray imaging. *Phys. Med. Biol.*, 54(5):1307–1318, 2009.
150. E. Roessl and R. Proksa. K-edge imaging in x-ray computed tomography using multi-bin photon counting detectors. *Phys. Med. Biol.*, 52(15):4679–4696, 2007.
151. L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
152. A. Sawatzky. Performance of first-order algorithms for TV penalized weighted least-squares denoising problem. In *Image and Signal Processing*, volume 8509 of *Lecture Notes in Computer Science*, pages 340–349. Springer International Publishing, 2014.
153. A. Sawatzky, C. Brune, T. Kösters, F. Wübbeling, and M. Burger. EM-TV methods for inverse problems with Poisson noise. In *Level Set and PDE Based Reconstruction Methods in Imaging*, pages 71–142. Springer, 2013.
154. A. Sawatzky, D. Tenbrinck, X. Jiang, and M. Burger. A variational framework for region-based segmentation incorporating physical noise models. *J. Math. Imaging Vis.*, 47(3):179–209, 2013.
155. A. Sawatzky, Q. Xu, C. O. Schirra, and M. A. Anastasio. Proximal ADMM for multi-channel image reconstruction in spectral X-ray CT. *IEEE Trans. Med. Imag.*, 33(8):1657–1668, 2014.
156. H. Schäfer. Über die Methode sukzessiver Approximationen. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 59:131–140, 1957.
157. K. P. Schäfers, T. J. Spinks, P. G. Camici, P. M. Bloomfield, C. G. Rhodes, M. P. Law, C. S. R. Baker, and O. Rimoldi. Absolute quantification of myocardial blood flow with $H_2^{15}O$ and 3-dimensional PET: An experimental validation. *J. Nucl. Med.*, 43:1031–1040, 2002.
158. C. O. Schirra, B. Brendel, M. A. Anastasio, and E. Roessl. Spectral CT: a technology primer for contrast agent development. *Contrast Media Mol. Imaging*, 9(1):62–70, 2014.
159. C. O. Schirra, E. Roessl, T. Koehler, B. Brendel, A. Thran, D. Pan, M. A. Anastasio, and R. Proksa. Statistical reconstruction of material decomposed data in spectral CT. *IEEE Trans. Med. Imag.*, 32(7):1249–1257, 2013.
160. J. P. Schlomka, E. Roessl, R. Dorscheid, S. Dill, G. Martens, T. Istel, C. Bäumer, C. Herrmann, R. Steadmann, G. Zeitler, A. Livne, and R. Proksa. Experimental feasibility of multi-energy photon-counting K-edge imaging in pre-clinical computed tomography. *Phys. Med. Biol.*, 53(15):4031–4047, 2008.
161. M. Schmidt, N. Le Roux, and F. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. Technical report, arXiv e-print, 2011. <http://arxiv.org/abs/1109.2415>.
162. M. Schrader, S. W. Hell, and H. T. M. van der Voort. Three-dimensional super-resolution with a 4Pi-confocal microscope using image restoration. *J. Appl. Phys.*, 84(8):4033–4042, 1998.
163. T. Schuster, B. Kaltenbacher, B. Hofmann, and K. S. Kazimierski. *Regularization methods in Banach spaces*, volume 10. Walter de Gruyter, 2012.
164. S. Setzer. Operator splittings, Bregman methods and frame shrinkage in image processing. *International Journal of Computer Vision*, 92(3):265–280, 2011.
165. S. Setzer, G. Steidl, and J. Morgenthaler. A cyclic projected gradient method. *Computational Optimization and Applications*, 54(2):417–440, 2013.
166. S. Setzer, G. Steidl, and T. Teuber. Deblurring Poissonian images by split Bregman techniques. *J. Vis. Commun. Image R.*, 21:193–199, 2010.

167. E. Y. Sidky, J. H. Jørgensen, and X. Pan. Convex optimization problem prototyping for image reconstruction in computed tomography with the Chambolle-Pock algorithm. *Phys. Med. Biol.*, 57(10):3065–3091, 2012.
168. G. Steidl and T. Teuber. Removing multiplicative noise by Douglas-Rachford splitting methods. *J. Math. Imaging Vis.*, 36:168–184, 2010.
169. T. Teuber. *Anisotropic Smoothing and Image Restoration Facing Non-Gaussian Noise*. PhD thesis, Technische Universität Kaiserslautern, Apr. 2012. available from <https://kluedo.uni-kl.de/frontdoor/index/index/docId/3219>.
170. P. Tseng. Applications of a splitting algorithm to decomposition in convex programming and variational inequalities. *SIAM Journal on Control and Optimization*, 29:119–138, 1991.
171. P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. Technical report, Dept. of Mathematics, University of Washington, Seattle, 2008.
172. T. Valkonen. A primal-dual hybrid gradient method for nonlinear operators with applications to MRI. *Inverse Problems*, 30(5):055012, 2014.
173. B. Vandeghinste, B. Goossens, J. D. Beenhouwer, A. Pizurica, W. Philips, S. Vandenberghe, and S. Staelens. Split-Bregman-based sparse-view CT reconstruction. *Proc. Int. Meeting Fully 3D Image Recon. Rad. Nucl. Med.*, pages 431–434, 2011.
174. Y. Vardi, L. A. Shepp, and L. Kaufman. A statistical model for positron emission tomography. *J. Am. Stat. Assoc.*, 80(389):8–20, 1985.
175. S. Villa, S. Salzo, L. Baldassarre, and A. Verri. Accelerated and inexact forward-backward algorithms. *SIAM J. Optim.*, 23(3):1607–1633, 2013.
176. J. von Neumann. Some matrix inequalities and metrization of matrix-space. In *Collected Works, Pergamon, Oxford, 1962, Volume IV, 205–218*, pages 286–300. Tomsk University Review, 1937.
177. B. C. Vũ. A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Computational Mathematics*, 38(3):667–681, 2013.
178. G. Wang, H. Yu, and B. D. Man. An outlook on x-ray CT research and development. *Med. Phys.*, 35(3):1051–1064, 2008.
179. J. Wang, T. Li, H. Lu, and Z. Liang. Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose X-ray computed tomography. *IEEE Trans. Med. Imag.*, 25(10):1272–1283, 2006.
180. K. Wang, R. Su, A. A. Oraevsky, and M. A. Anastasio. Investigation of iterative image reconstruction in three-dimensional optoacoustic tomography. *Phys. Med. Biol.*, 57:5399–5423, 2012.
181. S. L. Wang and L. Z. Liao. Decomposition method with a variable parameter for a class of monotone variational inequality problems. *J. Optimiz. Theory App.*, 109(2):415–429, 2001.
182. G. A. Watson. Characterization of the subdifferential of some matrix norms. *Linear Algebra and its Applications*, 170:33–45, 1992.
183. M. N. Wernick and J. N. Aarsvold, editors. *Emission Tomography: The Fundamentals of PET and SPECT*. Elsevier Academic Press, 2004.
184. Q. Xu, A. Sawatzky, and M. A. Anastasio. A multi-channel image reconstruction method for grating-based X-ray phase-contrast computed tomography. In *Proc. SPIE 9033, Medical Imaging 2014: Physics of Medical Imaging*, 2014.
185. Q. Xu, A. Sawatzky, M. A. Anastasio, and C. O. Schirra. Sparsity-regularized image reconstruction of decomposed K-edge data in spectral CT. *Phys. Med. Biol.*, 59(10):N65, 2014.
186. M. Yan and W. Yin. Self equivalence of the alternating direction method of multipliers. In: R. Glowinski, S. Osher, W. Yin (eds.) *Splitting Methods in Communication and Imaging*, Science and Engineering. Springer, 2016.
187. W. Yin. Analysis and generalizations of the linearized Bregman method. *SIAM Journal on Imaging Sciences*, 3(4):856–877, 2010.
188. J. Yuan, C. Schnörr, and G. Steidl. Simultaneous higher order optical flow estimation and decomposition. *SIAM Journal on Scientific Computing*, 29(6):2283–2304, 2007.
189. R. Zhang, J.-B. Thibault, C. A. Bouman, and K. D. S. J. Hsieh. A model-based iterative algorithm for dual-energy X-ray CT reconstruction. In *Proc. Int. Conf. Image Form. in X-ray CT*, pages 439–443, 2012.

190. X. Zhang, M. Burger, and S. Osher. A unified primal-dual algorithm framework based on Bregman iteration. *J. Sci. Comput.*, 46(1):20–46, 2011.
191. M. Zhu and T. F. Chan. An efficient primal-dual hybrid gradient algorithm for total variation image restoration. *UCLA CAM Report 08-34*, 2008.
192. H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of Royal Statistical Society: Series B (Statistical Methodology)*, 67:301–320, 2005.
193. Y. Zou and M. D. Silver. Analysis of fast kV-switching in dual energy CT using a pre-reconstruction decomposition technique. In *Proc. SPIE (Medical Imaging 2008)*, volume 6913, page 691313, 2008.