



Bregman Itoh–Abe Methods for Sparse Optimisation

Martin Benning¹ · Erlend Skaldehaug Riis² · Carola-Bibiane Schönlieb²

Received: 31 May 2019 / Accepted: 17 January 2020 / Published online: 3 February 2020
© The Author(s) 2020

Abstract

In this paper we propose optimisation methods for variational regularisation problems based on discretising the inverse scale space flow with discrete gradient methods. Inverse scale space flow generalises gradient flows by incorporating a generalised Bregman distance as the underlying metric. Its discrete-time counterparts, Bregman iterations and linearised Bregman iterations are popular regularisation schemes for inverse problems that incorporate a priori information without loss of contrast. Discrete gradient methods are tools from geometric numerical integration for preserving energy dissipation of dissipative differential systems. The resultant Bregman discrete gradient methods are unconditionally dissipative and achieve rapid convergence rates by exploiting structures of the problem such as sparsity. Building on previous work on discrete gradients for non-smooth, non-convex optimisation, we prove convergence guarantees for these methods in a Clarke subdifferential framework. Numerical results for convex and non-convex examples are presented.

Keywords Non-convex optimisation · Non-smooth optimisation · Bregman iteration · Inverse scale space · Geometric numerical integration · Discrete gradient methods

Mathematics Subject Classification 49M37 · 49Q15 · 65K10 · 90C26

This article is dedicated to Mila Nikolova whose keen interest and inspiring discussions have encouraged our research on geometric integration for optimisation. MB acknowledges support from the Leverhulme Trust Early Career Fellowship ECF-2016-611 ‘Learning from mistakes: a supervised feedback-loop for imaging applications’. ESR and CBS acknowledge support from CHiPS (Horizon 2020 RISE project grant) and from the Cantab Capital Institute for the Mathematics of Information. CBS acknowledges support from the Leverhulme Trust project on “Breaking the non-convexity barrier” EPSRC grant Nr EP/M00483X/1, and the EPSRC Centre Nr EP/N014588/1. Moreover, CBS acknowledges support from the RISE project NoMADS and the Alan Turing Institute.

✉ Erlend Skaldehaug Riis
e.s.riis@damtp.cam.ac.uk

Martin Benning
m.benning@qmul.ac.uk

Carola-Bibiane Schönlieb
cbs31@cam.ac.uk

¹ School of Mathematical Sciences, Queen Mary University of London, London, UK

² Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, UK

1 Introduction

We consider the constrained optimisation problem

$$\min_{\mathbf{x} \in \mathcal{C}} V(\mathbf{x}), \quad (1.1)$$

for an objective function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ and constraint $\mathcal{C} \subset \mathbb{R}^n$. The function V may be non-convex and non-smooth, as outlined in Assumption 3.1. In this paper, we propose and study optimisation schemes by using tools from geometric numerical integration to solve the *inverse scale space* (ISS) flow.

The ISS flow is a differential system which generalises gradient flows by replacing the Euclidean distance by a *Bregman distance*, defined via a convex *Bregman (distance generating) function* $J : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$. The ISS flow is given by

$$\dot{\mathbf{p}}(t) = -\nabla V(\mathbf{x}(t)), \quad \mathbf{p}(t) \in \partial J(\mathbf{x}(t)). \quad (1.2)$$

Here, ∂J is the convex subdifferential of J , to be defined in Sect. 2, and $\overline{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$. The term *inverse scale space flow* goes back to Scherzer and Groetsch [50]. This flow is typically derived as the continuous-time limit of *Bregman iterations*—methods for solving variational regularisation

problems. Like gradient flows, the ISS flow is a dissipative system, and its dissipative structure is determined by the function J . This allows one to solve (1.1) while incorporating *a priori* information into the optimisation scheme, with the benefits of converging to superior solutions, and doing so faster. For background on the ISS flow and Bregman distances, see Sect. 2.1.

Geometric numerical integration deals with numerical integration methods that preserve geometric structures of the continuous system. Such geometric structures include dissipation or conservation of energy and Lyapunov functions. In recent years, geometric numerical integration—and numerical integration in general—has gained interest within the mathematical optimisation community, as a framework for formulating iterative schemes that are dissipative or amenable to Lyapunov arguments.

We propose to discretise the inverse scale space flow with discrete gradient methods. These are methods from geometric numerical integration that preserve the aforementioned geometric structures in a general setting. In recent papers [21,26,45,46], optimisation schemes based on discretising gradient flows with discrete gradients have been analysed and implemented for various problems. Favourable properties of the discrete gradient methods include unconditional dissipation, i.e. dissipation is ensured for any time step, and the Itoh–Abe discrete gradient method is derivative-free and has convergence guarantees in the non-smooth, non-convex setting [45]. For smooth problems, the theoretical convergence rates of the discrete gradient methods match those of explicit gradient descent and coordinate descent [21]. The drawback of these methods is that the updates are in general implicit. Nevertheless, for many simple variational problems, the updates turn out to be explicit.

In this paper, we study the Itoh–Abe discrete gradient method applied to the ISS flow. We prove that the method is well defined and converges to a set of stationary points for non-smooth, non-convex functions. Furthermore, building on the paper by Miyatake et al. [38] which pointed out the equivalence between the discrete gradient methods for linear systems and successive-over-relaxation (SOR) methods, we point out equivalencies of various approaches to least squares problems.

Bregman iterations, and related methods, are closely tied to inverse problems and regularisation methods, particularly in signal processing. We consider numerical examples in this setting as well.

1.1 Related Literature

Spurred by applications for variational regularisation in image processing and compressed sensing, the ISS flow and the Bregman method have been active areas of research during the last decade. The Bregman iterative method was

originally proposed by Osher et al. [41] in 2005 for total variation-based image denoising, representing an extension of the Bregman proximal algorithm [14,19,32,55] to non-smooth Bregman functions. Subsequently the ISS flow was derived and analysed by Burger et al. [7,9–11]. Since then, researchers have studied the ISS flow with applications to generalised spectral analysis in a nonlinear setting, i.e. by Burger et al. [8], Gilboa et al. [23], and Schmidt et al. [51]. The Bregman method has been studied for ℓ^1 -regularisation and compressed sensing by Goldstein and Osher [24] and Yin et al. [59], and extended to primal-dual algorithms by Zhang et al. [61].

The linearised Bregman method was proposed by Yin et al. [59] for applications to ℓ^1 -regularisation and compressed sensing, and further studied in this setting by Cai et al. [12], and Dong et al. [18]. An extension for non-convex problems was proposed by Benning et al. [3], proving global convergence for functions that satisfy the Kurdyka–Łojasiewicz property. Lorenz et al. [34,52] proposed a sparse variant of the Kaczmarz method for linear problems based on linearised Bregman iterations. These and other methods were unified in a Split Feasibility Problems framework for general convergence results by Lorenz et al. [33]. For further details on Bregman iterations and linearised Bregman methods, we refer to [4].

We review papers on discrete gradient methods for optimisation based on gradient flows. Grimm et al. [26] used discrete gradients to solve variational regularisation problems in image analysis and proved convergence to a set of stationary points for continuously differentiable functions. Ehrhardt et al. [21] provided additional analysis for the methods, including convergence rates for smooth, convex problems and Polyak–Łojasiewicz functions, as well as well-posedness of the implicit equation. Ringholm et al. [46] applied the Itoh–Abe discrete gradient method to non-convex image problems with Euler’s elastica regularisation.

Furthermore, several recent works have looked at discrete gradient methods in more general settings. Riis et al. [45] studied the Itoh–Abe discrete gradient method in the setting of derivative-free optimisation of non-smooth, non-convex objective functions, and proved that the method converges to a set of stationary points in the Clarke subdifferential framework. Celledoni et al. [13] extended the Itoh–Abe discrete gradient method to optimisation problems defined on Riemannian manifolds. Hernández-Solano et al. [29] combined a discrete gradient method with Hopfield networks in order to preserve a Lyapunov function for optimisation problems.

We point out that a central feature of discrete gradient methods is that due to their implicit formulation, no restrictions are required for the time steps. This will also be the case for the analysis in this paper.

Beyond discrete gradients, other methods from numerical integration have led to notable developments in optimisation

in recent years. Recent papers by Su et al. [54] and Wibisono et al. [56] study second-order ODEs based on the continuous-time limit of Nesterov’s accelerated gradient descent [39], in order to gain insight into the acceleration phenomenon. In this setting, Wilson et al. [57] provided a framework for Lyapunov analysis of optimisation schemes and continuous-time dynamics, and Betancourt et al. [5] proposed a framework for symplectic integration for optimisation. On a related note, Scieur et al. [53] showed that several accelerated optimisation schemes can be derived as multi-step integration schemes from numerical analysis. An optimisation scheme based on numerically integrating dissipative Hamiltonian conformal systems was proposed by Maddison et al. [35], with the aim of ensuring linear convergence for a larger group of functions than classical Other numerical integration methods of relevance include implicit Runge–Kutta methods, where energy dissipation is ensured under mild time step restrictions [27], and explicit stabilised methods for solving strongly convex problems [20]. Another example is the study of gradient flows in metric spaces and minimising movement schemes [1], concerning gradient flow trajectories under other measures of distance, such as the Wasserstein metric [49].

1.2 Structure and Contributions

The paper is structured as follows. For the remainder of Sect. 1, we define mathematical notation. In Sect. 2, we provide preliminary material for convex and non-convex analysis, and introduce Bregman distances. In Sect. 3, we propose a Bregman discrete gradient method based on the ISS flow, and prove well-posedness and convergence results in a non-convex, non-smooth framework. In Sect. 4, we discuss particular examples of Bregman discrete gradient methods. In Sect. 5, we present results from numerical experiments. In Sect. 6, we conclude.

1.3 Notation and Preliminaries

Throughout the paper, we denote by $\|\cdot\|$ the ℓ^2 -norm, i.e. $\|\mathbf{x}\|^2 = \sum_{i=1}^n |x_i|^2$. For $p \in [1, +\infty]$, $\|\cdot\|_p$ denotes the usual ℓ^p -norm. For $\varepsilon > 0$ and $\mathbf{x} \in \mathbb{R}^n$, we denote by $B_\varepsilon(\mathbf{x})$ the open ball $\{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{y} - \mathbf{x}\| < \varepsilon\}$. We denote by $\{\mathbf{e}^1, \dots, \mathbf{e}^n\}$ the standard coordinate vectors in \mathbb{R}^n . We define the extended reals as $\overline{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$.

For a sequence $(\mathbf{x}^k)_{k \in \mathbb{N}} \subset \mathbb{R}^n$, we denote by S its limit set, i.e. the set of accumulation points,

$$S = \{\mathbf{x}^* \in \mathbb{R}^n : \exists (\mathbf{x}^{k_j})_{j \in \mathbb{N}} \text{ s.t. } \mathbf{x}^{k_j} \rightarrow \mathbf{x}^*\}.$$

For a matrix $A \in \mathbb{R}^{m \times n}$, we denote by $\mathbf{a}^i \in \mathbb{R}^n$ its i th row, and A^T its transpose. Accordingly, a_j^i refers to the j th element of the i th column of A .

For $x \in \mathbb{R}$, the *sign operator* $\text{sgn} : \mathbb{R} \rightarrow \{\pm 1, 0\}$ is defined as

$$\text{sgn}(x) = \begin{cases} 1, & \text{if } x > 0, \\ 0, & \text{if } x = 0, \\ -1, & \text{if } x < 0. \end{cases}$$

2 Preliminaries for Convex and Non-convex Analysis

In this section, we provide preliminary material on nonsmooth analysis. In Sect. 2.1, we cover convex analysis and Bregman distances, while in Sect. 2.2 we cover nonconvex, nonsmooth analysis in the Clarke subdifferential framework.

2.1 Convex Analysis

We consider functions $J : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ that are convex, proper, and lower-semicontinuous (see [22,48] for details on this class of functions).

Definition 2.1 (*Effective domain*) The *effective domain* of a function $J : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is defined as $\text{dom}(J) = \{\mathbf{x} \in \mathbb{R}^n : J(\mathbf{x}) < \infty\}$.

Definition 2.2 (*Subgradients and subdifferentials*) The *subdifferential* of a convex function $J : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ at $\mathbf{x} \in \mathbb{R}^n$ is the set of vectors

$$\begin{aligned} \partial J(\mathbf{x}) &= \{\mathbf{p} \in \mathbb{R}^n : J(\mathbf{y}) - J(\mathbf{x}) - \langle \mathbf{p}, \mathbf{y} - \mathbf{x} \rangle \geq 0 \text{ for all } \mathbf{y} \in \mathbb{R}^n\}. \end{aligned}$$

Vectors in $\partial J(\mathbf{x})$ are called *subgradients* of J at \mathbf{x} .

For $i = 1, \dots, n$, we denote by $[\partial J(\mathbf{x})]_i$ the projection of $\partial J(\mathbf{x})$ onto the i th coordinate, i.e. $[\partial J(\mathbf{x})]_i = \{p_i : \mathbf{p} \in \partial J(\mathbf{x})\}$.

For the theory of convex functions and their subdifferentials, see [48].

We consider the *characteristic function* $\chi_{\mathcal{C}}$ of a convex, closed set $\mathcal{C} \subset \mathbb{R}^n$, defined by

$$\chi_{\mathcal{C}}(\mathbf{x}) := \begin{cases} 0, & \text{if } \mathbf{x} \in \mathcal{C}, \\ +\infty, & \text{else.} \end{cases}$$

This is a convex, proper, lower-semicontinuous function, and $\mathbf{0} \in \partial \chi_{\mathcal{C}}(\mathbf{x})$ for all $\mathbf{x} \in \text{dom}(\chi_{\mathcal{C}}) = \mathcal{C}$.

We can now define the generalised Bregman distance [6] of a convex function J .

Definition 2.3 (Bregman distance) Given $\mathbf{p} \in \partial J(\mathbf{x})$, the Bregman distance between \mathbf{x} and \mathbf{y} is given by

$$D_J^p(\mathbf{x}, \mathbf{y}) = J(\mathbf{y}) - J(\mathbf{x}) - \langle \mathbf{p}, \mathbf{y} - \mathbf{x} \rangle.$$

We refer to J as the corresponding Bregman function.

Example 2.4 If $J(\mathbf{x}) = \|\mathbf{x}\|^2/2$, then $D_J^p(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|^2/2$, i.e. the square of the Euclidean distance.

While Bregman distances are non-negative due to convexity of J , they are not metrics as they do not generally satisfy symmetry or a triangle inequality. However, a related object is the symmetric Bregman distance.

Definition 2.5 (Symmetric Bregman distance) Given $\mathbf{p} \in \partial J(\mathbf{x})$ and $\mathbf{q} \in \partial J(\mathbf{y})$, the symmetric Bregman distance between \mathbf{x} and \mathbf{y} is given by

$$D_J^{\text{symm}}(\mathbf{x}, \mathbf{y}) = D_J^q(\mathbf{y}, \mathbf{x}) + D_J^p(\mathbf{x}, \mathbf{y}) = \langle \mathbf{q} - \mathbf{p}, \mathbf{y} - \mathbf{x} \rangle.$$

Definition 2.6 (μ -convexity) A proper, lower-semicontinuous convex function $J : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is μ -convex for $\mu \geq 0$ if either of the following (equivalent) conditions hold.

- (i) The function $J(\cdot) - \frac{\mu}{2} \|\cdot\|^2$ is convex.
- (ii) $J(\mathbf{y}) - J(\mathbf{x}) - \langle \mathbf{p}, \mathbf{y} - \mathbf{x} \rangle \geq \frac{\mu}{2} \|\mathbf{y} - \mathbf{x}\|^2$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, $\mathbf{p} \in \partial J(\mathbf{x})$.

If $\mu > 0$, J is said to be strongly convex.

For strongly convex Bregman functions, the following property of Bregman distances is immediate.

Proposition 2.7 (Bregman distance under μ -convexity) If J is μ -convex, then for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,

$$D_J^p(\mathbf{x}, \mathbf{y}) \geq \frac{\mu}{2} \|\mathbf{y} - \mathbf{x}\|^2, \quad D_J^{\text{symm}}(\mathbf{x}, \mathbf{y}) \geq \mu \|\mathbf{y} - \mathbf{x}\|^2.$$

2.2 Non-Convex Subdifferential Analysis

We summarise the main concepts of the Clarke subdifferential for locally Lipschitz continuous, non-smooth, non-convex functions $V : \mathbb{R}^n \rightarrow \mathbb{R}$, and refer to [16] for further details.

Definition 2.8 (Lipschitz continuity) V is Lipschitz of rank L near \mathbf{x} if there exists $\varepsilon > 0$ such that for all $\mathbf{y}, \mathbf{z} \in B_\varepsilon(\mathbf{x})$, one has

$$|V(\mathbf{y}) - V(\mathbf{z})| \leq L \|\mathbf{y} - \mathbf{z}\|.$$

V is locally Lipschitz continuous if the above property holds for all $\mathbf{x} \in \mathbb{R}^n$.

Definition 2.9 (Clarke directional derivative) For V Lipschitz near \mathbf{x} and for a vector $\mathbf{d} \in \mathbb{R}^n$, the Clarke directional derivative is given by

$$V^o(\mathbf{x}; \mathbf{d}) = \limsup_{\mathbf{y} \rightarrow \mathbf{x}, \lambda \downarrow 0} \frac{V(\mathbf{y} + \lambda \mathbf{d}) - V(\mathbf{y})}{\lambda}.$$

Definition 2.10 (Clarke subdifferential) Let V be locally Lipschitz continuous and $\mathbf{x} \in \mathbb{R}^n$. The Clarke subdifferential of V at \mathbf{x} is given by

$$\partial V(\mathbf{x}) = \{\mathbf{p} \in \mathbb{R}^n : V^o(\mathbf{x}; \mathbf{d}) \geq \langle \mathbf{d}, \mathbf{p} \rangle \text{ for all } \mathbf{d} \in \mathbb{R}^n\}.$$

An element of $\partial V(\mathbf{x})$ is called a Clarke subgradient.

The Clarke subdifferential was introduced by Clarke in [15]. It is well defined for locally Lipschitz functions, coincides with the standard subdifferential for convex functions [16, Proposition 2.2.7], and coincides with the derivative at points of strict differentiability [16, Proposition 2.2.4]. We additionally state two useful results, both of which can be found in Chapter 2 of [16].

Proposition 2.11 (Properties of Clarke subdifferential) Suppose V is locally Lipschitz continuous. Then

- (i) $\partial V(\mathbf{x})$ is nonempty, convex and compact, and if V is Lipschitz of rank L near \mathbf{x} , then $\partial V(\mathbf{x}) \subseteq B_L(\mathbf{0})$.
- (ii) $\partial V(\mathbf{x})$ is outer semicontinuous. That is, for all $\varepsilon > 0$, there exists $\delta > 0$ such that

$$\partial V(\mathbf{y}) \subset \partial V(\mathbf{x}) + B_\varepsilon(\mathbf{0}), \quad \text{for all } \mathbf{y} \in B_\delta(\mathbf{x}).$$

In this paper, we will consider constrained optimisation problems $\arg \min_{\mathbf{x} \in \mathcal{C}} V(\mathbf{x})$ for a closed, convex subset \mathcal{C} . We assume in the convergence analysis that \mathcal{C} is box constraints of the form $\mathcal{C} = \otimes_{i=1}^n [l_i, u_i]$. To have a formal notion of first-order optimality in this setting, we define the tangent cone.

Definition 2.12 (Tangent cone) For a convex, closed set $\mathcal{C} \subset \mathbb{R}^n$, the tangent cone at $\mathbf{x} \in \mathcal{C}$, denoted by $T(\mathcal{C}, \mathbf{x})$, is the closure of the set of vectors $\mathbf{v} \in \mathbb{R}^n$ such that there is $\delta > 0$ such that for all $\varepsilon \in (0, \delta)$, $\mathbf{x} + \varepsilon \mathbf{v} \in \mathcal{C}$.

Definition 2.13 (Clarke stationarity) Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be a locally Lipschitz continuous function and $\mathcal{C} \subset \mathbb{R}^n$ a closed, convex set. A point $\mathbf{x}^* \in \mathbb{R}^n$ is a Clarke stationary point of V restricted to \mathcal{C} if $V^o(\mathbf{x}^*; \mathbf{d}) \geq 0$ for all $\mathbf{d} \in T(\mathcal{C}, \mathbf{x}^*)$.

As the following proposition shows, the above definition generalises minimisers for convex functions, and Clarke stationary points in the unrestricted case.

Proposition 2.14 (Stationary versus minimal) *Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function and $\mathcal{C} \subset \mathbb{R}^n$ be closed and convex. Then $\mathbf{x}^* \in \mathbb{R}^n$ is a Clarke stationary point restricted to \mathcal{C} if and only if $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathcal{C}} V(\mathbf{x})$.*

Proof This follows directly from [31, Theorem 4.19], noting that for convex functions the Clarke directional derivative coincides with the classical directional derivative [16, Proposition 2.2.7]. \square

3 The Discrete Gradient Method for the Inverse Scale Space Flow

In what follows, we discuss the inverse scale space flow, and propose a *Bregman discrete gradient method* by using discrete gradients to solve the differential system.

3.1 Inverse Scale Space Flow and Bregman Methods

For a convex function $J : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$, objective function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ and starting points $\mathbf{x}(0) = \mathbf{x}^0 \in \mathbb{R}^n$, $\mathbf{p}(0) \in \partial J(\mathbf{x}^0)$, the ISS flow is the dissipative differential system given by (1.2). If J were twice continuously differentiable and μ -convex, then (1.2) could be rewritten as

$$\dot{\mathbf{x}}(t) = -(\nabla^2 J(\mathbf{x}(t)))^{-1} \nabla V(\mathbf{x}(t)),$$

and the energy $V(\mathbf{x}(t))$ would dissipate over time as

$$\begin{aligned} \frac{d}{dt} V(\mathbf{x}(t)) &= \langle \dot{\mathbf{x}}(t), \nabla V(\mathbf{x}(t)) \rangle \\ &= -\langle \dot{\mathbf{x}}(t), \nabla^2 J(\mathbf{x}(t)) \dot{\mathbf{x}}(t) \rangle \leq -\mu \|\dot{\mathbf{x}}(t)\|^2. \end{aligned}$$

We briefly discuss variants of Bregman methods as discretisations of (1.2). The Bregman method is derived by backward Euler discretisation of (1.2) and is given by

$$\mathbf{p}^{k+1} = \mathbf{p}^k - \tau_k \nabla V(\mathbf{x}^{k+1}), \quad \mathbf{p}^{k+1} \in \partial J(\mathbf{x}^{k+1})$$

which can be rewritten as

$$\mathbf{x}^{k+1} = \arg \min_{\mathbf{x} \in \mathbb{R}^n} V(\mathbf{x}) + \frac{1}{\tau_k} D_J^{\mathbf{p}^k}(\mathbf{x}^k, \mathbf{x}). \tag{3.1}$$

From (3.1), we see that the Bregman method is dissipative, as

$$\begin{aligned} V(\mathbf{x}^{k+1}) - V(\mathbf{x}^k) &\leq -\frac{1}{\tau_k} D_J^{\mathbf{p}^k}(\mathbf{x}^k, \mathbf{x}^{k+1}) \\ &\leq -\frac{\mu}{2\tau_k} \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2. \end{aligned}$$

Similarly, the linearised Bregman method is derived by forward Euler discretisation of (1.2) and is given by

$$\mathbf{p}^{k+1} = \mathbf{p}^k - \tau_k \nabla V(\mathbf{x}^k), \quad \mathbf{p}^{k+1} \in \partial J(\mathbf{x}^{k+1})$$

or equivalently

$$\mathbf{x}^{k+1} = \arg \min_{\mathbf{x} \in \mathbb{R}^n} V(\mathbf{x}^k) + \langle \nabla V(\mathbf{x}^k), \mathbf{x} - \mathbf{x}^k \rangle + \frac{1}{\tau_k} D_J^{\mathbf{p}^k}(\mathbf{x}^k, \mathbf{x}).$$

The ISS flow and Bregman methods are considered for solving ill-conditioned linear systems $A\mathbf{x} = \mathbf{b}$, with the objective function

$$V(\mathbf{x}) = \frac{1}{2} \|A\mathbf{x} - \mathbf{b}\|^2.$$

In this setting, iterates of both the Bregman method and the linearised Bregman method converge [4,33] to a solution of

$$\min_{\mathbf{x} \in \mathbb{R}^n} \{J(\mathbf{x}) \text{ s.t. } A\mathbf{x} = \mathbf{b}\}.$$

Furthermore, applications of the ISS flow include image denoising with reduced contrast-loss and staircasing effects [41], recovering eigenfunctions [51], and identifying sparse or low-rank structures [59].

We make the following assumptions on the objective function V , the constraints \mathcal{C} , as well as the Bregman function J .

- Assumption 3.1** (a) The function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ is locally Lipschitz continuous and bounded below.
- (b) $\mathbf{x}^* \in \mathbb{R}^n$ is a Clarke stationary point of V restricted to \mathcal{C} if and only if for all coordinate vectors \mathbf{e}^i , we have $V^o(\mathbf{x}^*; \mathbf{e}^i), V^o(\mathbf{x}^*; -\mathbf{e}^i) \geq 0$.
- (c) The set $\mathcal{C} \subset \mathbb{R}^n$ consists of coordinate-wise box constraints, i.e. $\mathcal{C} = \otimes_{i=1}^n [l_i, u_i]$.
- (d) The function $J : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is proper, lower-semicontinuous, and μ -convex with $\mu > 0$. Furthermore, $J(\mathbf{x}) = \sum_{i=1}^n j_i(x_i) + \chi_{[l_i, u_i]}(x_i)$, where $j_i : \mathbb{R} \rightarrow \overline{\mathbb{R}}$ are convex and Lipschitz continuous on $[l_i, u_i]$ for each i .

3.2 The Bregman Discrete Gradient Method

In what follows, we define discrete gradients and propose the Bregman discrete gradient method by discretising the ISS flow.

Definition 3.2 (*Discrete gradient*) Let V be a continuously differentiable function. A *discrete gradient* is a continuous map $\overline{\nabla} V : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,

$$\langle \overline{\nabla} V(\mathbf{x}, \mathbf{y}), \mathbf{y} - \mathbf{x} \rangle = V(\mathbf{y}) - V(\mathbf{x}) \quad (\text{Mean value}), \tag{3.2}$$

$$\lim_{\mathbf{y} \rightarrow \mathbf{x}} \bar{\nabla} V(\mathbf{x}, \mathbf{y}) = \nabla V(\mathbf{x}) \quad (\text{Consistency}). \tag{3.3}$$

Discrete gradients are tools from geometric numerical integration. In geometric integration, one studies methods for numerically solving ODEs while preserving certain structures of the continuous system—see [28,36] for an introduction. Discrete gradients are tools for solving first-order ODEs that preserve energy conservation laws, dissipation laws, and Lyapunov functions [25,30,37,44].

The Itoh–Abe discrete gradient [30] (also known as the coordinate increment discrete gradient) is given by

$$\bar{\nabla} V(\mathbf{x}, \mathbf{y}) = \begin{pmatrix} \frac{V(y_1, x_2, \dots, x_n) - V(\mathbf{x})}{y_1 - x_1} \\ \frac{V(y_1, y_2, x_3, \dots, x_n) - V(y_1, x_2, \dots, x_n)}{y_2 - x_2} \\ \vdots \\ \frac{V(\mathbf{y}) - V(y_1, \dots, y_{n-1}, x_n)}{y_n - x_n} \end{pmatrix}, \tag{3.4}$$

where $0/0$ is interpreted as $\partial_i V(\mathbf{x})$.

The Itoh–Abe discrete gradient is derivative-free and is evaluated by successively computing difference quotients. For optimisation problems where the gradient is expensive to compute (e.g. [46]), or unavailable (e.g. [45]), Itoh–Abe discrete gradient methods are useful for ensuring convergence guarantees and competitive convergence rates [21] without requiring derivatives.

We propose the *Bregman discrete gradient method* as follows. For a starting point $\mathbf{x}^0 \in \mathbb{R}^n$, subgradient $\mathbf{p}^0 \in \partial J(\mathbf{x}^0)$, and time steps $(\tau_k)_{k \in \mathbb{N}} > 0$, solve for $k = 0, 1, \dots$,

$$\begin{aligned} \mathbf{p}^{k+1} &= \mathbf{p}^k - \tau_k \bar{\nabla} V(\mathbf{x}^k, \mathbf{x}^{k+1}) \\ \mathbf{p}^{k+1} &\in \partial J(\mathbf{x}^{k+1}). \end{aligned} \tag{3.5}$$

This scheme preserves the dissipative structure of the ISS flow and (linearised) Bregman methods, as we see by applying the mean value property (3.2) and (3.5).

$$\begin{aligned} V(\mathbf{x}^k) - V(\mathbf{x}^{k+1}) &= \langle \mathbf{x}^k - \mathbf{x}^{k+1}, \bar{\nabla} V(\mathbf{x}^k, \mathbf{x}^{k+1}) \rangle \\ &= \frac{1}{\tau_k} \langle \mathbf{x}^k - \mathbf{x}^{k+1}, \mathbf{p}^k - \mathbf{p}^{k+1} \rangle \\ &= \frac{1}{\tau_k} D_J^{\text{symm}}(\mathbf{x}^k, \mathbf{x}^{k+1}) \\ &\geq \frac{\mu}{\tau_k} \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2. \end{aligned} \tag{3.6}$$

Furthermore, if we plug in $J(\mathbf{x}) = \|\mathbf{x}\|^2/2$, we recover the discrete gradient method for gradient flows [21]. Observe that the dissipative structure holds for arbitrary positive time steps $\tau_k > 0$. We thus only require that $\tau_k \in [\tau_{\min}, \tau_{\max}]$ for arbitrary bounds $\tau_{\min}, \tau_{\max} > 0$.

By Assumption 3.1 d), the subdifferential of J is separable in the coordinates, i.e.

$$\partial J(\mathbf{x}) = \prod_{i=1}^n \partial \chi_{[l_i, u_i]}(x_i) + \partial j_i(x_i).$$

It follows that solving the Bregman Itoh–Abe equation (3.5) corresponds to successively solving n scalar inclusions, i.e. given \mathbf{x}^k , and \mathbf{p}^k , we want \mathbf{x}^{k+1} and \mathbf{p}^{k+1} that solve

$$\begin{aligned} p_i^{k+1} &= p_i^k - \tau_{k,i} \frac{V(\mathbf{y}^{k,i}) - V(\mathbf{y}^{k,i-1})}{x_i^{k+1} - x_i^k}, \\ p_i^{k+1} &\in \partial j_i(y_i^{k,i}) + \partial \chi_{[l_i, u_i]}(y_i^{k,i}), \\ \mathbf{y}^{k,i} &= [x_1^{k+1}, \dots, x_i^{k+1}, x_{i+1}^k, \dots, x_n^k], \\ i &= 1, \dots, n. \end{aligned} \tag{3.7}$$

Here $\mathbf{y}^{k,i}$ denotes $[x_1^{k+1}, \dots, x_i^{k+1}, x_{i+1}^k, \dots, x_n^k]^T$. For a choice of $v_i^k \in [\partial V(\mathbf{y}^{k,i-1})]_i$, if $p_i^k - \tau_{k,i} v_i^k \in [\partial J(\mathbf{y}^{k,i-1})]_i$, then we consider $x_i^{k+1} = x_i^k$ and $p_i^{k+1} = p_i^k - \tau_{k,i} v_i^k$ an admissible update.

Remark 3.3 While discrete gradients are not defined for non-smooth functions V , the resultant Bregman discrete gradient method (3.7) is still well defined, provided the function V has a well defined Clarke subdifferential. This is the case, e.g. if V is locally Lipschitz continuous [16]. Furthermore, note that the properties used to derive the dissipative structure (3.6) holds in this setting.

4 Theoretical Results

In this section, we prove that a solution to the Bregman discrete gradient method (3.7) exists, and provide conditions in which the update is unique. Furthermore, we prove that the accumulation points of the iterates $(\mathbf{x}^k)_{k \in \mathbb{N}}$ from the Bregman Itoh–Abe method are Clarke stationary.

4.1 Well-Posedness and Preliminary Results

The Bregman Itoh–Abe scheme (3.7) is in general implicit. We therefore want to ensure that an update exists and is easy to compute. Furthermore, we want to know under what conditions the updates are unique. In what follows, we address these questions. The corresponding analysis generalises that from the previous papers on discrete gradient methods for optimisation [21,45].

We first show that an update (not necessarily unique) always exists.

Lemma 4.1 (Existence) *For any $\tau > 0$, $\mathbf{x}^k \in \mathbb{R}^n$, and $\mathbf{p}^k \in \partial J(\mathbf{x}^k)$, there exists an update $(\mathbf{x}^{k+1}, \mathbf{p}^{k+1})$ that satisfies (3.7).*

Proof As (3.7) consists of successive scalar updates, it is sufficient to consider a scalar problem, $v : \mathbb{R} \rightarrow \mathbb{R}, j : \mathbb{R} \rightarrow \overline{\mathbb{R}}$. For $x \in \mathbb{R}$ and $p \in \partial j(x)$ we either want $y \neq x$ such that

$$p - \tau \frac{v(y) - v(x)}{y - x} \in \partial j(y), \tag{4.1}$$

or $y = x$ and $p - \tau w \in \partial j(x)$, for some $w \in \partial v(x)$.

If such a w exists, we are done. Otherwise, we have $\min\{v^o(x; 1), v^o(x; -1)\} < 0$ and may assume that $v^o(x; 1) < 0$. In this case, we will show that there exists $y > x$ such that (4.1) holds.

Since $p - \tau v^o(x; 1) > p$ and $p \in \partial j(x)$, we deduce that $p - \tau v^o(x; 1) > \partial j(x)$. By the outer semicontinuity of sub-differentials and definition of Clarke directional derivatives, there is $\delta > 0$ such that

$$p - \tau \frac{v(y) - v(x)}{y - x} > \partial j(y) \text{ for all } y \in (x, x + \delta).$$

On the other hand, as v is bounded below,

$$y \mapsto (v(y) - v(x))/(y - x)$$

is bounded below for all $y \in [x + \delta, +\infty)$, while by μ -convexity of j , we have $\partial j(y) \geq \partial j(x) + \mu(y - x)$ for all $y \in [x + \delta, +\infty)$. Hence, there is $r \gg 0$ such that

$$p - \tau \frac{v(y) - v(x)}{y - x} < \partial j(y) \text{ for all } y \geq x + r.$$

By continuity of v , and by outer semicontinuity of subdifferentials, it follows that there exists $y \in (x + \delta, x + r)$ that solves (4.1). This concludes the proof. \square

Next we give conditions in which the update is unique. While this holds for any time step if the objective function is convex, it may fail for nonconvex functions, as discussed in [21, Section 5].

Proposition 4.2 (Uniqueness) *For $\tau > 0, \mathbf{x}^k \in \mathbb{R}^n$, and $\mathbf{p}^k \in \partial J(\mathbf{x}^k)$, the update $(\mathbf{x}^{k+1}, \mathbf{p}^{k+1})$ to (3.7) is unique in the following cases.*

- (i) V is convex.
- (ii) $V(\cdot) + \eta\|\cdot\|^2$ is convex for $\eta > 0$, and $\tau < \mu/\eta$.

Proof *Case (i)* Suppose V is convex. The existence of a solution to (3.7) is guaranteed by Lemma 4.1. To establish uniqueness, we argue as follows. An update $\mathbf{y}^{k,i}$ must satisfy

$$p_i^k - \tau \frac{V(\mathbf{y}^{k,i}) - V(\mathbf{y}^{k,i-1})}{x_i^{k+1} - x_i^k} \in [\partial J(\mathbf{y}^{k,i})]_i. \tag{4.2}$$

The left-hand side is non-increasing with respect to x_i^{k+1} due to the difference quotient term of the convex function V , while the right-hand side is strictly increasing due to the strong convexity of J . Hence there cannot be two distinct solutions for $y_i^{k,i}$ to the scalar equation. This implies uniqueness of the update.

Case (ii): Suppose $V(\cdot) + \eta\|\cdot\|^2$ is convex and $\tau < \mu/\eta$. We rewrite (4.2) as

$$\frac{p_i^k}{\tau} - \frac{V(\mathbf{y}^{k,i}) - V(\mathbf{y}^{k,i-1})}{x_i^{k+1} - x_i^k} - \eta x_i^{k+1} \in \frac{1}{\tau} [\partial J(\mathbf{y}^{k,i})]_i - \eta x_i^{k+1}.$$

Since $\tau\eta < \mu$, the function $J(\cdot) - \tau\eta\|\cdot\|^2/2$ is strongly convex, and therefore the right-hand side is strictly increasing with respect to x_i^{k+1} . On the other side, as $V(\cdot) + \eta\|\cdot\|^2$ is convex, the difference quotient term $\frac{V(\mathbf{y}^{k,i}) - V(\mathbf{y}^{k,i-1})}{x_i^{k+1} - x_i^k} + \eta x_i^{k+1}$ is non-decreasing with respect to x_i^{k+1} . Therefore, by the same reasoning as in the first case, the update to the Bregman Itoh–Abe method must be unique. \square

Remark 4.3 If the update is stationary, i.e. $x_i^{k+1} = x_i^k$, then the subgradient update p_i^{k+1} is unique only up to the choice of subderivative $v_i \in [\partial V(\mathbf{y}^{k,i-1})]_i$.

Note that while the convexity criteria in Proposition 4.2 are global, they are often generalised to local properties such as lower C^2 -smoothness [47] and prox-regularity [43]. However, an analysis of these properties in the context of Itoh–Abe discrete gradient methods is beyond the theoretical and practical scope of this paper. We also add that lack of uniqueness is not an issue for the implementation or computational tractability of the scheme.

The following lemma summarises several useful properties of the iterates $(\mathbf{x}^k, \mathbf{p}^k)_{k \in \mathbb{N}}$, and generalises Lemmas 3.3, 3.4 in [45].

Lemma 4.4 (Properties of scheme) *Let V, J , and \mathcal{C} satisfy Assumption 3.1, and let $(\mathbf{x}^k, \mathbf{p}^k)_{k \in \mathbb{N}}$ be iterates that solve (3.7) for time steps $(\tau_k)_{k \in \mathbb{N}} \subset [\tau_{\min}, \tau_{\max}]$. Then the following properties hold.*

- (i) $V(\mathbf{x}^{k+1}) \leq V(\mathbf{x}^k)$.
- (ii) $\lim_{k \rightarrow \infty} \|\mathbf{x}^{k+1} - \mathbf{x}^k\| = 0$.
- (iii) If V is coercive, then there exists a convergent subsequence of $(\mathbf{x}^k, \mathbf{p}^k)_{k \in \mathbb{N}}$.
- (iv) The set of limit points S is compact, connected, and has empty interior. Furthermore, V is single-valued on S .

Proof Property (i) follows from (3.6).

As V is bounded below and $(V(\mathbf{x}^k))_{k \in \mathbb{N}}$ is decreasing, $V(\mathbf{x}^k) \rightarrow V^*$ for some limit V^* . Therefore, by (3.6),

$$\begin{aligned} V(\mathbf{x}^0) - V^* &= \sum_{k=0}^{\infty} V(\mathbf{x}^k) - V(\mathbf{x}^{k+1}) \\ &\geq \sum_{k=0}^{\infty} \frac{\mu}{\tau_k} \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 \\ &\geq \frac{\mu}{\tau_{\max}} \sum_{k=0}^{\infty} \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2. \end{aligned}$$

This implies property (ii).

Properties (iii) and (iv) follow from (i) and (ii) and are proven, respectively, in [45, Lemma 3.3 and Lemma 3.4]. □

4.2 Amended Scheme

In the next subsection, we prove that accumulation points of the iterates from the Bregman Itoh–Abe method are Clarke stationary. However, we first modify the Bregman Itoh–Abe method to ‘forget’ subgradients induced by the constraints.

We define the modified scheme as follows. For a starting point $\mathbf{x}^0, \mathbf{p}^0 \in \partial J(\mathbf{x}^0)$, and $k \in \mathbb{N}$, update

$$\begin{aligned} p_i^{k+1} + \tilde{q} &= p_i^k - \tau_{k,i} \frac{V(\mathbf{y}^{k,i}) - V(\mathbf{y}^{k,i-1})}{x_i^{k+1} - x_i^k}, \\ p_i^{k+1} &\in \partial j_i(y_i^{k,i}), \quad \tilde{q} \in \partial \chi_{[l_i, u_i]}(y_i^{k,i}), \\ \mathbf{y}^{k,i} &= [x_1^{k+1}, \dots, x_i^{k+1}, x_{i+1}^k, \dots, x_n^k], \\ i &= 1, \dots, n. \end{aligned} \tag{4.3}$$

Observe that since $\mathbf{0} \in \partial \chi_{\mathcal{C}}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{C}$, we still have $\mathbf{p}^{k+1} \in \partial J(\mathbf{x}^{k+1})$. It is also straightforward to verify that the previous results and analysis of dissipative structure in this section also hold for (4.3).

The reason for introducing the modified scheme (4.3) is so that the subgradient iterates p_i^k do not grow arbitrarily large when x_i^k equals u_i or l_i for several updates (recall that when the constraint is active, $\partial \chi_{[l_i, u_i]}$ is unbounded). Otherwise, the iterates might get stuck at a non-stationary point for several iterations, since x_i^k is unable to leave $\{l_i, u_i\}$ until the subgradient update in $\partial \chi_{[l_i, u_i]}$ vanishes, yielding inefficient progress. Furthermore, this leads to pathological, albeit unlikely, examples where the iterates of (3.7) converge to non-stationary accumulation points. For completeness, we give such an example in Section A.

4.3 Main Convergence Theorem

Having introduced a modified Bregman Itoh–Abe scheme in the previous section, we proceed to state and prove the main theorem of this paper, namely that all accumulation points of the scheme (4.3) are non-stationary. We note that this also holds for the original Bregman Itoh–Abe method if the iterates $(\mathbf{x}^k)_{k \in \mathbb{N}}$ converge to a unique limit.

Theorem 4.5 (Stationarity guarantees) *Let V, J , and \mathcal{C} satisfy Assumption 3.1, and suppose the sequence of iterates $(\mathbf{x}^k, \mathbf{p}^k)_{k \in \mathbb{N}}$ solves (4.3) for time steps $(\tau_k)_{k \in \mathbb{N}} \subset [\tau_{\min}, \tau_{\max}]$. Then all accumulation points $\mathbf{x}^* \in S$ are Clarke stationary points restricted to \mathcal{C} .*

Proof Let $\mathbf{x}^* \in S$ and consider a convergent subsequence $(\mathbf{x}^{k_j})_{j \in \mathbb{N}}$. We want to show for each basis vector \mathbf{e}^i that either $V^o(\mathbf{x}^*; \mathbf{e}^i) \geq 0$ or $x_i^* = u_i$, and analogously that either $V^o(\mathbf{x}^*; -\mathbf{e}^i) \geq 0$ or $x_i^* = l_i$. As the arguments are equivalent, we only prove the first case.

Suppose for contradiction that $V^o(\mathbf{x}^*; \mathbf{e}^i) < -\eta$ for some $\eta > 0$, and that $x_i^* < u_i$. By the definition of the Clarke directional derivative, Definition 2.9, there are $\varepsilon, \delta > 0$ such that for all $\mathbf{x} \in B_\varepsilon(\mathbf{x}^*)$ and $\lambda \in (0, \delta)$, we have

$$\frac{V(\mathbf{x} + \lambda \mathbf{e}^i) - V(\mathbf{x})}{\lambda} \leq -\frac{\eta}{2}. \tag{4.4}$$

Since $\mathbf{x}^{k_j} \rightarrow \mathbf{x}^*$ and $\|\mathbf{x}^{k_j+1} - \mathbf{x}^{k_j}\| \rightarrow 0$, for each $N \in \mathbb{N}$ there exists K such that for all $j \geq K$, we have $\mathbf{x}^k \in B_\varepsilon(\mathbf{x}^*)$ and $\|\mathbf{x}^k - \mathbf{x}^{k+1}\| < \delta$ for $k = k_j, k_j + 1, \dots, k_j + N$. By making $\varepsilon > 0$ sufficiently small, we have $B_\varepsilon(x_i^*) < u_i$. Furthermore, since $x_i^{k+1} \geq x_i^k$ for $k = k_j, \dots, k_j + N - 1$, we deduce that the constraint component \tilde{q} is zero. By (4.4), it follows that

$$\begin{aligned} p_i^{k_j} - p_i^{k_j+N} &= \sum_{k=k_j}^{N-1} p_i^k - p_i^{k+1} \\ &= \sum_{k=k_j}^{N-1} \tau_i^k \frac{V(\mathbf{y}^{k,i}) - V(\mathbf{y}^{k,i-1})}{x_i^{k+1} - x_i^k} \\ &\leq -\tau_{\min} \sum_{k=k_j}^{N-1} \frac{\eta}{2} = -N \tau_{\min} \frac{\eta}{2}. \end{aligned} \tag{4.5}$$

By Assumption 3.1, ∂j_i is bounded on $U = B_\varepsilon(\mathbf{x}^*) \cap [l_i, u_i]$. Since $p_i^{k_j}, \dots, p_i^{k_j+N} \in \partial j_i(U)$, we can choose N such that $N \tau_{\min} \frac{\eta}{2} > \max \partial j_i(U) - \min \partial j_i(U)$ and arrive at a contradiction. Thus, \mathbf{x}^* is a Clarke stationary point restricted to \mathcal{C} . □

5 Examples of Bregman Itoh–Abe Discrete Gradient Schemes

In this section, we describe several schemes based on the Bregman Itoh–Abe discrete gradient scheme (3.7). We will primarily consider objective functions of the form

$$V(\mathbf{x}) = \frac{1}{2} \langle \mathbf{x}, A\mathbf{x} \rangle - \langle \mathbf{b}, \mathbf{x} \rangle, \tag{5.1}$$

where $A \in \mathbb{R}^{n \times n}$ is a symmetric, positive semi-definite matrix, with strictly positive entries $a_i^i > 0$ on the diagonal.

We are particularly interested in problems with underlying sparsity and/or constraints, with applications in image analysis. Throughout this section, we use a time step vector τ^k coordinate-wise scaled by the diagonal of A , i.e. $\tau^k = \tau/\text{diag}(A) = [\tau/a_1^1, \dots, \tau/a_n^n]$ for all $k \in \mathbb{N}$, and some $\tau > 0$.

We first introduce some well-known coordinate descent schemes for solving linear systems, which Miyatake et al. [38] showed were equivalent to the Itoh–Abe discrete gradient method. The successive-over-relaxation (SOR) method [60] updates each coordinate sequentially according to the rule

$$\begin{aligned} \mathbf{y}^{k,0} &= \mathbf{x}^k \\ \mathbf{y}^{k,i} &= \mathbf{y}^{k,i-1} - \frac{\omega}{a_i^i} (\langle \mathbf{a}^i, \mathbf{y}^{k,i-1} \rangle - b_i) \mathbf{e}^i, \\ \mathbf{x}^{k+1} &= \mathbf{y}^{k,n}, \end{aligned} \tag{5.2}$$

where $\omega \in (0, 2)$. For $\omega = 1$, this is the Gauss–Seidel method [60]. The SOR method is equivalent to the Itoh–Abe discrete gradient method

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \tau \bar{\nabla} V(\mathbf{x}^k, \mathbf{x}^{k+1}),$$

with V given by (5.1) with the time steps $\tau_i = 2\omega / ((2 - \omega)a_i^i)$.

5.1 Sparse SOR Method

We consider underdetermined linear systems and want to find sparse solutions \mathbf{x}^* . Hence we seek to apply the Bregman Itoh–Abe method (3.7) with objective function V given by (5.1), and

$$J(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|^2 + \gamma \|\mathbf{x}\|_1, \tag{5.3}$$

for $\gamma > 0$. We term this the *Bregman SOR (BSOR) method*.

By Proposition 4.2, the updates of this method are well defined and unique. One can verify that the updates are given as follows. Denote by \tilde{x}_i^{k+1} the standard SOR coordinate update from x_i^k , (5.2). Furthermore, for $\mathbf{p}^k \in \partial J(\mathbf{x}^k)$, we write $\mathbf{p}^k = \mathbf{x}^k + \gamma \mathbf{r}^k$, where $\mathbf{r}^k \in \partial \|\mathbf{x}^k\|_1$. Then (x_i^{k+1}, r_i^{k+1}) are given in closed form as

$$\begin{aligned} x_i^{k+1} &= S \left(\tilde{x}_i^{k+1} + \frac{2\gamma}{2 + \tau} r_i^k, \frac{2\gamma}{2 + \tau} \right), \\ r_i^{k+1} &= r_i^k + \frac{\tau}{\gamma a_i^i} \left(b_i - \langle \mathbf{a}^i, \mathbf{x}^k \rangle - \frac{a_i^i(2 + \tau)}{2\tau} (y_i - x_i) \right). \end{aligned} \tag{5.4}$$

Here $S : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ denotes the shrinkage operator $S(\mathbf{x}, \lambda) = \text{sgn}(\mathbf{x}) \max\{|\mathbf{x}| - \lambda, 0\}$, applied elementwise to a vector \mathbf{x} .

5.2 Sparse, Regularised SOR

If $\mathbf{b} = \mathbf{A}\mathbf{x}^{\text{true}} + \mathbf{n}$, where \mathbf{x}^{true} is the sparse ground truth and \mathbf{n} is noise, then it may be necessary to regularise the objective function as well. Hence we consider the objective function

$$V(\mathbf{x}) = \frac{1}{2} \langle \mathbf{x}, \mathbf{A}\mathbf{x} \rangle - \langle \mathbf{b}, \mathbf{x} \rangle + \lambda \|\mathbf{x}\|_1, \tag{5.5}$$

for some regularisation parameter $\lambda > 0$. The non-smoothness induced by $\|\cdot\|_1$ satisfies Assumption 3.1, so Theorem 4.5 implies Bregman Itoh–Abe discrete gradient methods will still converge to stationary points of this problem.

For both $J(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|^2 + \gamma \|\mathbf{x}\|_1$ and $J(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|^2$, the scheme (3.7) can be expressed in closed form for (5.5), on a case-by-case basis. However, for purposes of brevity, we leave including this in ‘‘Appendix’’.

6 Equivalence of Iterative Methods for Linear Systems

In what follows, we discuss and demonstrate equivalencies for different iterative methods for solving linear systems. We recall from the previous section that the SOR method (5.2) is equivalent to the Itoh–Abe discrete gradient method [38].

The explicit coordinate descent method [2,58] for minimising V is given by

$$\begin{aligned} \mathbf{y}^{k,0} &= \mathbf{x}^k \\ \mathbf{y}^{k,i} &= \mathbf{y}^{k,i-1} - \alpha_i [\nabla V(\mathbf{y}^{k,i-1})]_i \mathbf{e}^i, \\ \mathbf{x}^{k+1} &= \mathbf{y}^{k,n}, \end{aligned} \tag{6.1}$$

where $\alpha_i > 0$ is the time step. As mentioned in [58], the SOR method is also equivalent to the coordinate descent method with V given by (5.1) and the time step $\alpha_i = \omega/a_i^i$. Hence, in this setting, the Itoh–Abe discrete gradient method is equivalent not only to SOR methods, but to explicit coordinate descent.

It is not surprising that these iterative coordinate methods turn out to be the same, given that the gradient V in (5.1) is linear. Furthermore, these equivalencies extend to discretisations of the inverse scale space flow with J given by (5.3). The resultant Bregman Itoh–Abe scheme for (5.1) is described in Sect. 5.1. We may compare this to a *Bregman linearised coordinate descent* scheme,

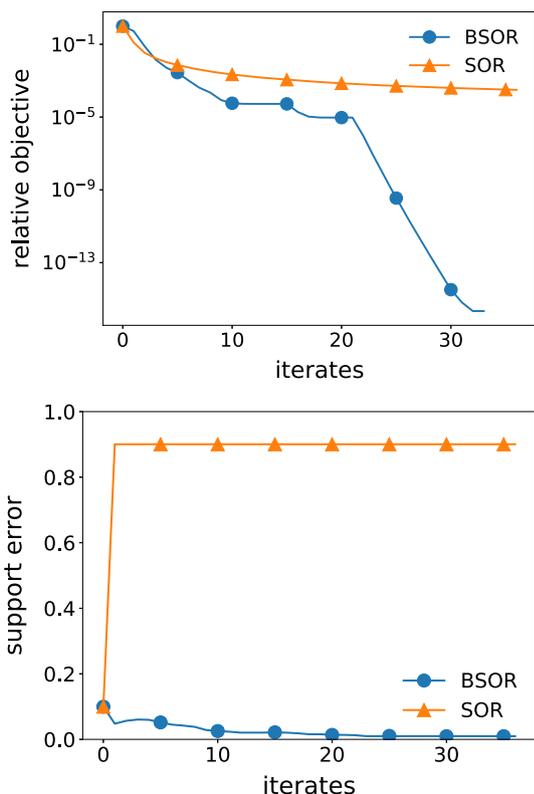


Fig. 1 Comparison of SOR and sparse SOR methods, for Gaussian linear system without noise. Top: Convergence rate for relative objective, i.e. $[V(\mathbf{x}^k) - V^*]/[V(\mathbf{x}^0) - V^*]$. Bottom: Support error with respect to iterates, i.e. proportion of indices i s.t. $\text{sgn}(x_i^k) = \text{sgn}(x_i^*)$

$$\begin{aligned} \mathbf{y}^{k,0} &= \mathbf{x}^k, \quad \mathbf{p}^k \in \partial J(\mathbf{x}^k) \\ z_i &= \arg \min_y [\nabla V(\mathbf{y}^{k,i-1})]_i \cdot y \\ &\quad + \frac{a_i^i}{\alpha_i} D_J^{\mathbf{p}^k}(\mathbf{y}^{k,i-1}, \mathbf{y}^{k,i-1} + y\mathbf{e}^i), \\ \mathbf{y}^{k,i} &= \mathbf{y}^{k,i-1} + z_i \mathbf{e}^i, \\ \mathbf{x}^{k+1} &= \mathbf{y}^{k,n}. \end{aligned}$$

One can verify that this scheme is equivalent to (6.1) for the parameters

$$\tau_i = \frac{2\alpha}{(2-\alpha)a_i^i}, \quad \lambda^* = \frac{\lambda}{1 + \frac{\alpha}{2-\alpha}}.$$

7 Numerical Examples

In this section, we present numerical results for the schemes described in Sect. 5.

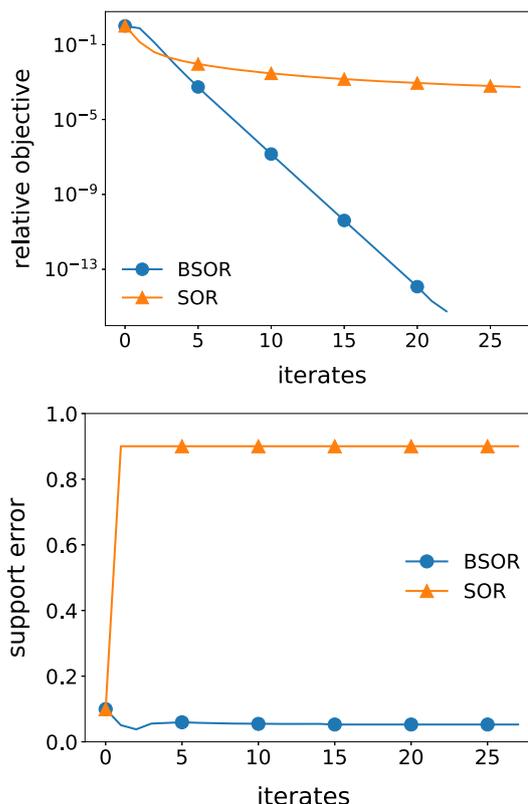


Fig. 2 Comparison of SOR and sparse SOR methods, for Gaussian linear system without noise, and binary ground truth. Top: Convergence rate for relative objective. Bottom: Support error with respect to iterates

7.1 Sparse SOR

We construct a matrix $A \in \mathbb{R}^{1024 \times 1024}$ from independent standard (zero mean, unit variance) Gaussian draws, and construct the sparse ground truth \mathbf{x}^{true} by choosing 10% of the indices at random determined by uniform draws on the unit interval. We then solve the problem

$$\arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{Ax} - \mathbf{b}\|^2,$$

where $\mathbf{b} = \mathbf{Ax}^{\text{true}}$. We compare the SOR method ($J(\mathbf{x}) = \|\mathbf{x}\|^2/2$) and the BSOR method ($J(\mathbf{x}) = \|\mathbf{x}\|^2/2 + \gamma \|\mathbf{x}\|_1$), where $\gamma = 1$. We set time steps to $\tau = 2/\text{diag}(A)$, corresponding to the Gauss–Seidel method. See Fig. 1 for the results.

For the same test problem, but where the ground truth is binary, i.e. only takes values 1 or 0, see Fig. 2.

7.2 Sparse, Regularised SOR

We construct $A \in \mathbb{R}^{1024 \times 1024}$ and \mathbf{x}^{true} as in the previous subsection. However, we add noise to the data, i.e. $\mathbf{b} = \mathbf{Ax}^{\text{true}} + \mathbf{n}$, where \mathbf{n} is independent Gaussian noise with a

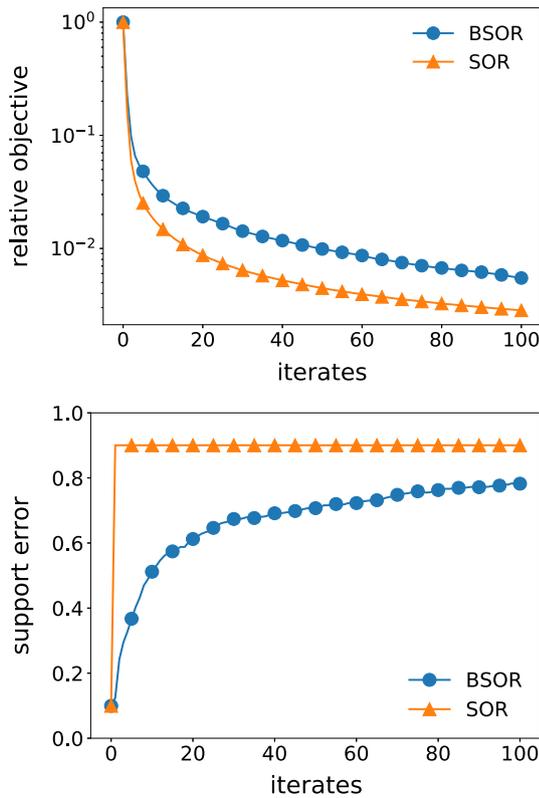


Fig. 3 Comparison of SOR and sparse SOR methods, for Gaussian linear system with noise. Top: Convergence rate for relative objective. Bottom: Support error with respect to iterates

standard deviation of $0.1\|A\mathbf{x}^{\text{true}}\|_{\infty}$. Since the added noise destroys the sparsity structure of $A^{-1}\mathbf{b}$, the sparse SOR method fails to improve the convergence rate. The results for $V(\mathbf{x}) = \|\mathbf{Ax} - \mathbf{b}\|^2/2$ are in Fig. 3.

We therefore include regularisation in the objective function of the form

$$V(\mathbf{x}) = \frac{1}{2}\|\mathbf{Ax} - \mathbf{b}\|^2/2 + \lambda\|\mathbf{x}\|_1,$$

where $\lambda = 100$, and with initialisation \mathbf{x}^0 constructed by random, independent Gaussian draws. The results are visualised in Fig. 4.

7.3 Student-t Regularised Image Denoising

We consider a nonconvex image denoising model, previously presented in [40], given by

$$F : \mathbb{R}^n \rightarrow \mathbb{R}, \quad F(\mathbf{x}) := \sum_{i=1}^N \varphi_i \Phi(K_i \mathbf{x}) + \|\mathbf{x} - \mathbf{x}^\delta\|_1. \quad (7.1)$$

Here $\{K_i\}_{i=1}^N$ is a collection of linear filters, $(\varphi_i)_{i=1}^N \subset [0, \infty)$ are coefficients, $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is the nonconvex function based on the student-t distribution, defined as

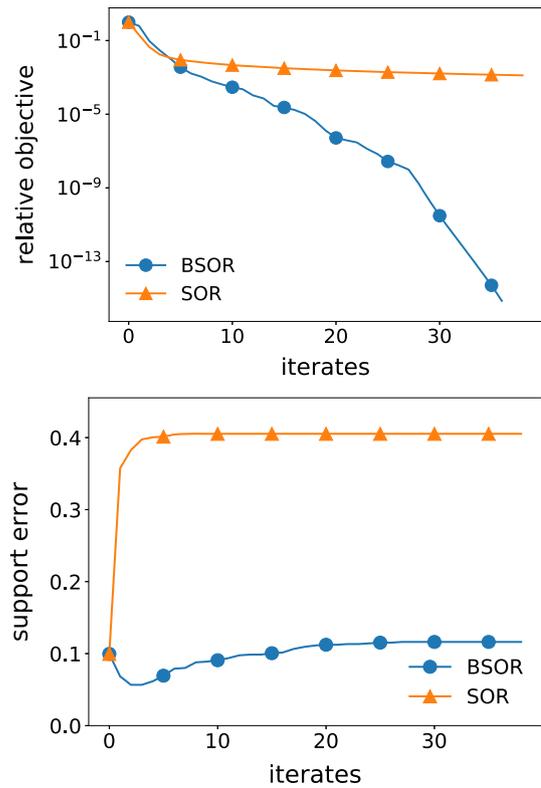


Fig. 4 Comparison of SOR and sparse SOR methods, for ℓ^1 -regularised linear system with noise. Top: Convergence rate for relative objective. Bottom: Support error with respect to iterates

$$\Phi(x) := \sum_{j=1}^n \psi(x_j), \quad \psi(x) := \log(1 + x^2),$$

and \mathbf{x}^δ is an image corrupted by impulse noise (salt & pepper noise).

As impulse noise only affects a small subset of pixels, we use the data fidelity term $\mathbf{x} \mapsto \|\mathbf{x} - \mathbf{x}^\delta\|_1$ to promote sparsity of $\mathbf{x}^* - \mathbf{x}^\delta$ for $\mathbf{x}^* \in \arg \min F(\mathbf{x})$. As linear filters, we consider the simple case of finite difference approximations to first-order derivatives of \mathbf{x} . We note that by applying a gradient flow to this regularisation function, we observe a similarity to Perona–Malik diffusion [42].

We consider a Bregman Itoh–Abe method (abbreviated to BIA) for solving $\min_{\mathbf{x} \in \mathbb{R}^n} F(\mathbf{x})$ with the Bregman function

$$J(\mathbf{x}) := \frac{1}{2}\|\mathbf{x}\|^2 + \gamma\|\mathbf{x} - \mathbf{x}^\delta\|_1,$$

to account for the sparsity of the residual $\mathbf{x}^* - \mathbf{x}^\delta$, and compare the method to the regular Itoh–Abe discrete gradient method (abbreviated to IA).

We set the starting point $\mathbf{x}^0 = \mathbf{x}^\delta$, and the parameters to $\tau_k = 1$ for all k , $\gamma = 0.5$, and $\varphi_i = 2$, $i = 1, 2$. For the impulse noise, we use a noise density of 10%. In the case

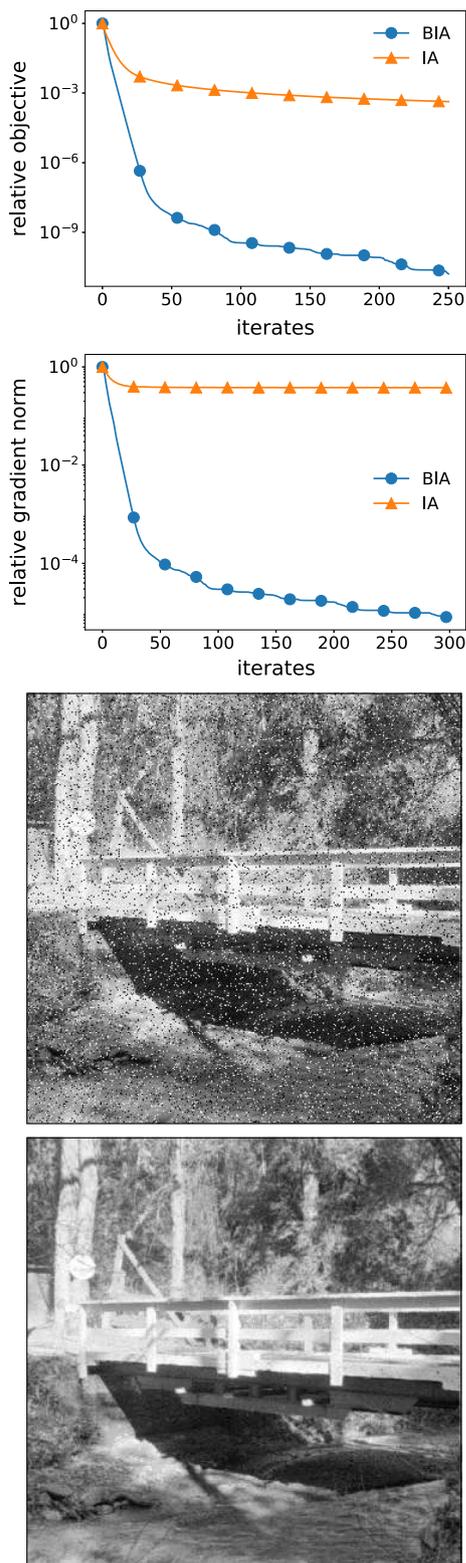


Fig. 5 Comparison of BIA and IA methods, for student-t regularised image denoising. First: Convergence rate for relative objective. Second: Convergence rate for relative gradient norm. Third: Input data. Fourth: Reconstruction

where x_i^{k+1} is not set to x_i^δ , we use the scalar root solver `scipy.optimize.brenth` on Python. Otherwise, the updates are in closed form.

See Fig. 5 for numerical results. By gradient norm, we mean $\text{dist}(\partial F(\mathbf{x}^k), 0)$. We observe that, as in the convex case, the Bregman Itoh–Abe method converges significantly faster than the Itoh–Abe method, as the former utilises the sparse structure of the problem.

8 Conclusion

In this paper, we propose to discretise the ISS flow with the Itoh–Abe discrete gradient. The resultant schemes exhibit a dissipative structure (3.6) related to the symmetrised Bregman distance of a function J . This generalises the discrete gradient method for gradient flows and can be viewed as a discrete gradient analogue to Bregman iterations. Building on previous studies of the Itoh–Abe discrete gradient method in the non-smooth, non-convex setting, we prove convergence guarantees of the Bregman Itoh–Abe discrete gradient method in such a setting.

We consider numerical examples motivated by linear systems and searching for sparse solutions, as well as a non-convex image denoising example. These results indicate that for sparse reconstructions, popular iterative solvers such as the SOR method can be significantly sped up by incorporating a Bregman step.

Future work is dedicated to proving convergence rates for the Bregman Itoh–Abe methods, and to compare the scheme to related methods such as the sparse Kaczmarz method [33]. Furthermore, we will study corresponding inverse scale space schemes using other discrete gradients, such as the mean value discrete gradient.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Counterexample of Theorem 4.5 for (3.7)

In the following example, we describe an example of a constrained optimisation problem for which the iterates of the

unmodified Bregman Itoh–Abe method (3.7) fails to converge to a limit set of stationary points.

Example A.1 As a starting point, we consider a C^2 -smooth objective function $W : \mathbb{R}^2 \rightarrow \mathbb{R}$ as described by Curry [17], for which the trajectory of a Euclidean gradient flow spirals along a gully, asymptotically trending towards the unit circle $S^1 = \{[x, y]^T \in \mathbb{R}^2 : x^2 + y^2 = 1\}$.

Denote by $[\tilde{x}^k, \tilde{y}^k]^T$ the iterates from the (standard) Itoh–Abe method for W with (arbitrary) time steps $\tau \equiv 1$ and starting point $[\tilde{x}^0, \tilde{y}^0]^T$ in the gully. As each update of the method consists of moving to a local point of descent, the iterates of this method will also remain in this gully and spiral towards S^1 .

We assume that there exists compact, disjoint sets $A, B \subset \mathbb{R}^2$ with nonempty interior, such that $B \cap S^1 \neq \emptyset$, and such that most of the iterates $([\tilde{x}^k, \tilde{y}^k]^T)_{k \in \mathbb{N}}$ are in A . That is, for all $k \in \mathbb{N}$, the set $\{i \leq k : [\tilde{x}^i, \tilde{y}^i]^T \in A\}$ is strictly greater than the set $\{i \leq k : [\tilde{x}^i, \tilde{y}^i]^T \notin A\}$. Such a set A exists as long as the iterates are spiralling around the unit circle at approximately a steadily decreasing rate.

Next we define an objective function $V : \mathbb{R}^3 \rightarrow \mathbb{R}$ as follows. Let $f_A, f_B : \mathbb{R}^2 \rightarrow \mathbb{R}$ be smooth support functions of A and B , respectively, with disjoint support sets. That is, $f_A(x, y) \in [0, 1]$, $f_B(x, y) \in [0, 1]$, and $f_A(x, y) \cdot f_B(x, y) = 0$ for all $[x, y]^T \in \mathbb{R}^2$, and furthermore, $f_A(A) \equiv 1$ and $f_B(B) \equiv 1$. We then consider the optimisation problem

$$\min_{z \geq 0} V(x, y, z) := W(x, y) + z(f_A(x, y) - f_B(x, y)). \tag{A.1}$$

We implement the unmodified Bregman Itoh–Abe method (3.7) with $\tau \equiv 1$, $\mathbf{x}^0 = [\tilde{x}^0, \tilde{y}^0, 0]^T$, and $J(\mathbf{x}) = \|\mathbf{x}\|^2/2 + \chi_{[z \geq 0]}(z)$, where we use the shorthand notation $\mathbf{x} = [x, y, z]^T$.

If $[x, y]^T \in B \cap S^1$, then $\partial_z V(x, y, 0) = -1$, meaning that $[x, y, 0]^T$ is a non-stationary point of V restricted to $\{z \geq 0\}$. Since the limit set of $([\tilde{x}^k, \tilde{y}^k])_{k \in \mathbb{N}}$ is S^1 , it follows that the iterates $([\tilde{x}^k, \tilde{y}^k, 0])_{k \in \mathbb{N}}$ admit non-stationary accumulation points for V . It therefore remains to show that for all $k \in \mathbb{N}$, $\mathbf{x}^k = [\tilde{x}^k, \tilde{y}^k, 0]$ are admissible updates to the Bregman Itoh–Abe method (3.7) for (A.1).

To verify that $\mathbf{x}^k = [\tilde{x}^k, \tilde{y}^k, 0]$ for all k , we can argue by induction. If $\mathbf{x}^k = [\tilde{x}^k, \tilde{y}^k, 0]$, then $V(\tilde{x}^k, \tilde{y}^k, 0) = W(\tilde{x}^k, \tilde{y}^k)$ so \tilde{x}^{k+1} and \tilde{y}^{k+1} are admissible updates for the x - and y -coordinates.

It remains to verify the same for the z -coordinate as well. If $[\tilde{x}^{k+1}, \tilde{y}^{k+1}]^T \in A$, then $\partial_z V(\tilde{x}^{k+1}, \tilde{y}^{k+1}, 0) = 1$, so we have the update $z^{k+1} = 0$ and $p_3^{k+1} = p_3^k - 1 \in \chi_{[z \geq 0]}(z^{k+1})$. On the other hand, if $[\tilde{x}^{k+1}, \tilde{y}^{k+1}]^T \notin A$, then $\partial_z V(\tilde{x}^{k+1}, \tilde{y}^{k+1}, 0) \in [-1, 1]$, so if in addition $p_3^k < -1$, then we have the update $z^{k+1} = 0$, $p_3^{k+1} = p_3^k -$

$\partial_z V(\tilde{x}^{k+1}, \tilde{y}^{k+1}, 0)$. Denoting by M and N the cardinalities of $\{i \leq k : [\tilde{x}^i, \tilde{y}^i]^T \in A\}$ and $\{i \leq k : [\tilde{x}^i, \tilde{y}^i]^T \notin A\}$, respectively, we know by construction of A that $M - N \geq 1$. Furthermore, we observe that $p_3^k \leq N - M \leq -1$, which concludes the argument.

Appendix B ℓ^1 -regularised sparse SOR method

In what follows, we describe the update rule for the Bregman Itoh–Abe method with V given in (5.5) and J given in (5.3).

Denote by \tilde{x}_i^{k+1} the standard SOR update (5.2) for the i th coordinate. Then the ℓ^1 -regularised sparse SOR method can be expressed as follows.

1. If $x_i^k = 0$ and $|\tilde{x}_i^{k+1} - \gamma r_i^k| \leq \gamma + \lambda \tau / a_i^i$, then

$$x_i^{k+1} = 0, \quad r_i^{k+1} = \frac{\gamma r_i^k - \tilde{x}_i^{k+1}}{\gamma + \lambda \tau / a_i^i}.$$

2. Else if

$$|(\tau/2 + 1)\tilde{x}_i^{k+1} + \gamma r_i^k| \geq \gamma + \tau \lambda / a_i^i,$$

then

$$x_i^{k+1} = \tilde{x}_i^{k+1} + \frac{\gamma r_i^k - (\gamma + \tau \lambda / a_i^i) \operatorname{sgn}(\tilde{x}_i^{k+1} + \frac{\gamma}{\tau/2+1} r_i^k)}{\tau/2 + 1}$$

$$r_i^{k+1} = \operatorname{sgn}\left(\tilde{x}_i^{k+1} + \frac{\gamma r_i^k}{(\tau/2 + 1)}\right).$$

3. Else if $x_i^k \neq 0$ and

$$\left| (\tau/2 + 1)\tilde{x}_i^{k+1} + \gamma r_i^k - (\lambda \tau / a_i^i) \operatorname{sgn}(x_i^k) \right| \leq \gamma,$$

then set

$$x_i^{k+1} = 0,$$

$$r_i^{k+1} = r_i^k + \frac{1}{\gamma} ((\tau/2 + 1)\tilde{x}_i^{k+1} - (\tau \lambda / a_i^i) \operatorname{sgn}(x_i^k)).$$

4. Else if $x_i \neq 0$ and

$$\left| 2 \left(\frac{a_i^i}{2} + \frac{a_i^i}{\tau} \right) \tilde{x}_i^{k+1} + \left(\frac{2a_i^i \gamma}{\tau} + \lambda \right) \operatorname{sgn}(x_i^k) \right|^2$$

$$\leq (b_i - \langle \mathbf{a}^i, \mathbf{y}^{k,i-1} \rangle) + \left(\frac{2a_i^i \gamma}{\tau} + \lambda \right) \operatorname{sgn}(x_i^k))^2 \dots$$

$$+ 8\lambda \left(\frac{a_i^i}{2} + \frac{a_i^i}{\tau} \right) |x_i^k|,$$

then set

$$x_i^{k+1} = \tilde{x}_i^{k+1} + \frac{\operatorname{sgn}(x_i^k)}{2\left(\frac{a_i^i}{2} + \frac{a_i^i}{\tau}\right)} \left(\frac{2a_i^i\gamma}{\tau} + \lambda \dots \right. \\ \left. - \sqrt{(b_i - \langle a^i, x^k \rangle + \left(\frac{2a_i^i\gamma}{\tau} + \lambda\right) \operatorname{sgn}(x_i^k))^2 + 8\lambda \left(\frac{a_i^i}{2} + \frac{a_i^i}{\tau}\right) |x_i^k|} \right), \\ r_i^{k+1} = -r_i^k.$$

References

- Ambrosio, L., Gigli, N., Savare, G.: Gradient flows: In Metric Spaces and in the Space of Probability Measures, 2nd edn. Lectures in Mathematics. ETH Zürich. Birkhäuser Basel (2008)
- Beck, A., Tetruashvili, L.: On the convergence of block coordinate descent type methods. *SIAM J. Optim.* **23**(4), 2037–2060 (2013)
- Benning, M., Betcke, M.M., Ehrhardt, M.J., Schönlieb, C.B.: Choose your path wisely: gradient descent in a Bregman distance framework. *ArXiv e-prints* (2017). <http://arxiv.org/abs/1712.04045>
- Benning, M., Burger, M.: Modern regularization methods for inverse problems. *Acta Numerica* **27**, 1–111 (2018)
- Betancourt, M., Jordan, M.I., Wilson, A.C.: On symplectic optimization. *arXiv e-prints* (2018). <http://arxiv.org/abs/1802.03653>
- Bregman, L.M.: The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. Math. Phys.* **7**(3), 200–217 (1967)
- Burger, M.: Bregman distances in inverse problems and partial differential equations. In: *Advances in Mathematical Modeling, Optimization and Optimal Control*, pp. 3–33. Springer (2016)
- Burger, M., Gilboa, G., Moeller, M., Eckardt, L., Cremers, D.: Spectral decompositions using one-homogeneous functionals. *SIAM J. Imag. Sci.* **9**(3), 1374–1408 (2016)
- Burger, M., Gilboa, G., Osher, S., Xu, J.: Nonlinear inverse scale space methods. *Commun. Math. Sci.* **4**(1), 179–212 (2006)
- Burger, M., Moeller, M., Benning, M., Osher, S.: An adaptive inverse scale space method for compressed sensing. *Math. Comput.* **82**(281), 269–299 (2013)
- Burger, M., Resmerita, E., He, L.: Error estimation for Bregman iterations and inverse scale space methods in image restoration. *Computing* **81**(2–3), 109–135 (2007)
- Cai, J.F., Osher, S., Shen, Z.: Linearized Bregman iterations for compressed sensing. *Math. Comput.* **78**(267), 1515–1536 (2009)
- Celledoni, E., Eidnes, S., Owren, B., Ringholm, T.: Dissipative numerical schemes on Riemannian manifolds with applications to gradient flows. *SIAM J. Sci. Comput.* **40**(6), A3789–A3806 (2018)
- Censor, Y., Zenios, S.A.: Proximal minimization algorithm with d -functions. *J. Optim. Theory. Appl.* **73**(3), 451–464 (1992)
- Clarke, F.H.: Necessary conditions for nonsmooth problems in optimal control and the calculus of variations. Ph.D. thesis, University of Washington (1973)
- Clarke, F.H.: *Optimization and Nonsmooth Analysis*. Classics in Applied Mathematics, 1st edn. SIAM, Philadelphia (1990)
- Curry, H.B.: The method of steepest descent for non-linear minimization problems. *Q. Appl. Math.* **2**(3), 258–261 (1944)
- Dong, B., Mao, Y., Osher, S., Yin, W.: Fast linearized Bregman iteration for compressive sensing and sparse denoising. *Commun. Math. Sci.* **8**(1), 93–111 (2010)
- Eckstein, J.: Nonlinear proximal point algorithms using Bregman functions, with applications to convex programming. *Math. Oper. Res.* **18**(1), 202–226 (1993)
- Eftekhari, A., Vandereycken, B., Vilmart, G., Zygalakis, K.C.: Explicit Stabilised Gradient Descent for Faster Strongly Convex Optimisation. *ArXiv e-prints* (2018). <http://arxiv.org/abs/1805.07199>
- Ehrhardt, M.J., Riis, E.S., Ringholm, T., Schönlieb, C.B.: A geometric integration approach to smooth optimisation: Foundations of the discrete gradient method. *ArXiv e-prints* (2018). <http://arxiv.org/abs/1805.06444>
- Ekeland, I., Témam, R.: *Convex Analysis and Variational Problems*, 1st edn. SIAM, Philadelphia (1999)
- Gilboa, G., Moeller, M., Burger, M.: Nonlinear spectral analysis via one-homogeneous functionals: overview and future prospects. *J. Math. Imaging Vis.* **56**(2), 300–319 (2016)
- Goldstein, T., Osher, S.: The split Bregman method for 11-regularized problems. *SIAM J. Imaging Sci.* **2**(2), 323–343 (2009)
- Gonzalez, O.: Time integration and discrete Hamiltonian systems. *J. Nonlinear Sci.* **6**(5), 449–467 (1996)
- Grimm, V., McLachlan, R.I., McLaren, D.I., Quispel, G.R.W., Schönlieb, C.B.: Discrete gradient methods for solving variational image regularisation models. *J. Phys. A* **50**(29), 295201 (2017)
- Hairer, E., Lubich, C.: Energy-diminishing integration of gradient systems. *IMA J. Numer. Anal.* **34**(2), 452–461 (2013)
- Hairer, E., Lubich, C., Wanner, G.: *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. vol. 31, 2nd edn. Springer, Berlin (2006)
- Hernández-Solano, Y., Atencia, M., Joya, G., Sandoval, F.: A discrete gradient method to enhance the numerical behaviour of Hopfield networks. *Neurocomputing* **164**, 45–55 (2015)
- Itoh, T., Abe, K.: Hamiltonian-conserving discrete canonical equations based on variational difference quotients. *J. Comput. Phys.* **76**(1), 85–102 (1988)
- Jahn, J.: *Introduction to the Theory of Nonlinear Optimization*, 3rd edn. Springer, Berlin (2007)
- Kiwiel, K.C.: Proximal minimization methods with generalized Bregman functions. *SIAM J. Control Optim.* **35**(4), 1142–1168 (1997)
- Lorenz, D.A., Schöpfer, F., Wenger, S.: The linearized Bregman method via split feasibility problems: analysis and generalizations. *SIAM J. Imaging Sci.* **7**(2), 1237–1262 (2014)
- Lorenz, D.A., Wenger, S., Schöpfer, F., Magnor, M.: A sparse Kaczmarz solver and a linearized Bregman method for online compressed sensing. *arXiv e-prints* (2014). <http://arxiv.org/abs/1403.7543>
- Maddison, C.J., Paulin, D., Teh, Y.W., O’Donoghue, B., Doucet, A.: Hamiltonian descent methods. *arXiv e-prints* (2018). <http://arxiv.org/abs/1809.05042>
- McLachlan, R.I., Quispel, G.R.W.: Six lectures on the geometric integration of ODEs, p. 155–210. London Mathematical Society Lecture Note Series. Cambridge University Press, Cambridge (2001)
- McLachlan, R.I., Quispel, G.R.W., Robidoux, N.: Geometric integration using discrete gradients. *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **357**(1754), 1021–1045 (1999)
- Miyatake, Y., Sogabe, T., Zhang, S.L.: On the equivalence between SOR-type methods for linear systems and the discrete gradient methods for gradient systems. *J. Comput. Appl. Math.* **342**, 58–69 (2018)
- Nesterov, Y.: A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Sov. Math. Doklady* **27**, 372–376 (1983)
- Ochs, P., Chen, Y., Brox, T., Pock, T.: iPiano: inertial proximal algorithm for nonconvex optimization. *SIAM J. Imaging Sci.* **7**(2), 1388–1419 (2014)

41. Osher, S., Burger, M., Goldfarb, D., Xu, J., Yin, W.: An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.* **4**(2), 460–489 (2005)
42. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**(7), 629–639 (1990)
43. Poliquin, R., Rockafellar, R.: Prox-regular functions in variational analysis. *Trans. Am. Math. Soc.* **348**(5), 1805–1838 (1996)
44. Quispel, G.R.W., Turner, G.S.: Discrete gradient methods for solving ODEs numerically while preserving a first integral. *J. Phys. A* **29**(13), 341–349 (1996)
45. Riis, E.S., Ehrhardt, M.J., Quispel, G.R.W., Schönlieb, C.B.: A geometric integration approach to nonsmooth, nonconvex optimisation. *ArXiv e-prints* (2018). <http://arxiv.org/abs/1807.07554>
46. Ringholm, T., Lazić, J., Schönlieb, C.B.: Variational image regularization with Euler’s elastica using a discrete gradient scheme. *SIAM J. Imaging Sci.* **11**(4), 2665–2691 (2018)
47. Rockafellar, R.: Maximal monotone relations and the second derivatives of nonsmooth functions. In: *Annales de l’Institut Henri Poincaré (C) Non Linear Analysis*, vol. 2, pp. 167–184. Elsevier (1985)
48. Rockafellar, R.T.: *Convex Analysis*. Princeton Landmarks in Mathematics and Physics, 1st edn. Princeton University Press, Princeton (2015)
49. Santambrogio, F.: {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bull. Math. Sci.* **7**(1), 87–154 (2017)
50. Scherzer, O., Groetsch, C.: Inverse scale space theory for inverse problems. In: *International Conference on Scale-Space Theories in Computer Vision*, pp. 317–325. Springer (2001)
51. Schmidt, M.F., Benning, M., Schönlieb, C.B.: Inverse scale space decomposition. *Inverse Probl.* **34**(4), 179–212 (2018)
52. Schöpfer, F., Lorenz, D.A.: Linear convergence of the randomized sparse Kaczmarz method. *Math. Program.* **173**(1), 509–536 (2019)
53. Scieur, D., Roulet, V., Bach, F., d’Aspremont, A.: Integration methods and optimization algorithms. In: *Advances in Neural Information Processing Systems*, pp. 1109–1118 (2017)
54. Su, W., Boyd, S., Candes, E.J.: A differential equation for modeling Nesterov’s accelerated gradient method: theory and insights. *J. Mach. Learn. Res.* **17**(153), 1–43 (2016)
55. Teboulle, M.: Entropic proximal mappings with applications to nonlinear programming. *Math. Oper. Res.* **17**(3), 670–690 (1992)
56. Wibisono, A., Wilson, A.C., Jordan, M.I.: A variational perspective on accelerated methods in optimization. *Proc. Natl. Acad. Sci.* **113**(47), E7351–E7358 (2016)
57. Wilson, A.C., Recht, B., Jordan, M.I.: A Lyapunov Analysis of Momentum Methods in Optimization. *arXiv e-prints* (2016). <http://arxiv.org/abs/1611.02635>
58. Wright, S.J.: Coordinate descent algorithms. *Math. Program.* **1**(151), 3–34 (2015)
59. Yin, W., Osher, S., Goldfarb, D., Darbon, J.: Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing. *SIAM J. Imaging Sci.* **1**(1), 143–168 (2008)
60. Young, D.M.: *Iterative Solution of Large Linear Systems*. Computer Science and Applied Mathematics, 1st edn. Academic Press Inc., Orlando (1971)
61. Zhang, X., Burger, M., Osher, S.: A unified primal–dual algorithm framework based on Bregman iteration. *J. Sci. Comput.* **46**(1), 20–46 (2011)



Martin Benning received both his Master’s degree in 2008 and the Ph.D. degree in 2011 at the University of Münster, Germany. He held postdoctoral positions at the University of Münster and the University of Cambridge, where he was also a Leverhulme Trust Early Career Fellow supported by the Isaac Newton Trust from 2016 to 2018 and member of the Cambridge Image Analysis group. Since 2018 he has been a Lecturer in Optimisation and Machine Learning at the School of Mathematical Sciences, Queen Mary University of London, UK, where he is a member of the Complex Systems and Networks group. His current research focuses on inverse problem theory, non-smooth variational regularisation methods, convex- and non-convex optimisation techniques and machine learning. His main applications of interest are magnetic resonance imaging (MRI), electron tomography (both from incomplete measurements) and positron emission tomography (PET).



Erlend Skaldehaug Riis received the MMath degree in 2015 from the University of Warwick and completed the Ph.D. degree in 2019 at the University of Cambridge, as a member of the Cambridge Image Analysis research group and the Cambridge Centre for Analysis. In 2020 he will be a London Mathematical Society Early Career Fellow, based at the University of Edinburgh, UK, and Escuela Politécnica Nacional, Quito, Ecuador. His current research focuses on methods from geometric numerical integration for optimisation, non-smooth optimisation, and bilevel parameter optimisation problems.



Carola-Bibiane Schönlieb received the B.S. degree from the Institute for Mathematics, University of Salzburg, Salzburg, Austria, in 2004, and the Ph.D. degree from the University of Cambridge, Cambridge, U.K., in 2009. From 2004 to 2005, she held a teaching position in Salzburg. She held a postdoctoral position at the University of Göttingen, Göttingen, Germany, for 1 year. She became a Lecturer with the Department of Applied Mathematics and Theoretical Physics (DAMTP), University of Cambridge, in 2010, where she was promoted to a Reader

in 2015 and a Professor in 2018. Since 2011, she has been a fellow of Jesus College, University of Cambridge, where she is currently a Professor of applied mathematics with DAMTP, the Head of the Cambridge Image Analysis Group, the Director of the Cantab Capital Institute for Mathematics of Information, and the Director of the Engineering and Physical Sciences Research Council Center for Mathematical and Statistical Analysis of Multimodal Clinical Imaging. Her research interests include variational methods, partial differential equations, and machine learning for image analysis, image processing, and inverse imaging problems.