

The use of discrete gradient methods for total variation type regularization problems in image processing

V. Grimm*, R. I. McLachlan†, D. I. McLaren‡, G. R. W. Quispel‡, and C.-B. Schönlieb§

Abstract. Discrete gradient methods are well-known methods of Geometric Numerical Integration, which preserve the dissipation of gradient systems. The preservation of the dissipation of a system is an important feature in numerous image processing tasks. We promote the use of discrete gradient methods in image processing by exhibiting experiments with nonlinear total variation (TV) deblurring, denoising, and inpainting.

Key words. Gradient system, gradient flow, discrete gradient, discrete gradient method, geometric numerical integration, total variation deblurring, total variation denoising, total variation inpainting

AMS subject classifications. 94A08, 37N30, 65D18.

1. Introduction. The discrete gradient method is mainly interesting as a computational method that preserves the underlying geometric structure of a gradient flow. Since we always end up with a problem in \mathbb{R}^n after discretization, we discuss the properties of the method in \mathbb{R}^n with some inner product $\langle \cdot, \cdot \rangle$. Given a differentiable functional $V : \mathbb{R}^n \rightarrow \mathbb{R}$, a gradient flow is the solution of the initial value problem

$$(1.1) \quad \dot{x} = -\nabla V(x), \quad x(0) = x_0,$$

where the dot represents differentiation with respect to time. We have the immediate consequence that the functional is nonincreasing along the solution of the evolution equation (1.1). Or, more exactly, we have the decay

$$(1.2) \quad \frac{d}{dt}V(x(t)) = \langle \nabla V(x(t)), \dot{x} \rangle = -\|\nabla V(x(t))\|^2 \leq 0.$$

Gradient systems of this type appear in many areas of image processing, for example, time-marching schemes, nonlinear diffusion filters such as the Perona-Malik model (cf. [33]) and many variants thereof, Sobolev gradient flows, image registration (e.g. [19, 39]) and some applications of active contours, snakes and level sets (e.g. [12, 30, 43]). In all applications, the preservation of the decay and the limit solution are the most important aspects. It is not so important to solve the evolution equation as accurately as possible, but to find the equilibrium as exactly as possible. Therefore, the preservation of the dissipative behavior of the evolution equation is very important. This is certainly not a new observation and has

* Department of Mathematics, Karlsruhe Institute of Technology (KIT), D-76128 Karlsruhe, Germany, email:volker.grimm@kit.edu.

† Institute of Fundamental Sciences, Massey University, Private Bag 11-222, Palmerston North, New Zealand, email:r.mclachlan@massey.ac.nz.

‡Department of Mathematics and Statistics, La Trobe University, Victoria 3086, Australia, email:{D.McLaren, r.quispel}@latrobe.edu.au.

§Department of Applied Mathematics and Theoretical Physics (DAMTP), University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom, email:cbs@cam.ac.uk.

been expressed by several authors. For example, in [42], a thorough discussion of the impact of the preservation of dissipativity in a gradient system with respect to diffusion filtering can be found. If the functional V is non-convex, as in some image processing applications such as sparse ℓ_p regularization [31, 24] and inpainting or data classification with the Ginzburg-Landau energy [3, 7, 4], then a decay guarantee within a gradient flow formulation will at least guarantee monotone convergence to a critical point of V . Structure preservation is the main topic of Geometric Numerical Integration (e.g. [21, 28]), which has been an active research area over the last two decades. Discrete gradient methods turned out to be especially useful to preserve the dissipative structure of a gradient system. We introduce these methods here and show that they have a great potential to be useful for image processing tasks. We also note that discrete gradient methods allow the use of adaptive time steps. Since in most image processing applications the accuracy with respect to the time evolution of the gradient flow is less important than the desire for good descent directions that take the iterates to the equilibrium fast, adaptive time steps seem particularly attractive. This in turn makes it possible to use large time steps initially, followed by smaller time steps once one approaches the equilibrium.

The paper is organized as follows: The discrete gradient method as a numerical method in Geometric Numerical Integration is introduced in Section 2 together with some of their favorable properties in Section 3. In Section 4 several experiments with well-known gradient systems in image processing are conducted that illustrate the importance of the preservation of the dissipativity of a gradient system and that indicate that these methods might be useful in image processing. Finally, a brief conclusion is given in Section 5.

2. Discrete gradient method. For simplicity, we will use \mathbb{R}^n equipped with the standard inner product and its associated norm. The proofs for this case can be generalized to Hilbert spaces, but for our purposes, where we either think of a digital grayscale image taken by a digital camera, \mathbb{R}^n where $n = N_x N_y$ and N_x and N_y correspond to the number of pixels with respect to width and height of the digitized picture, or an *RGB* picture where $n = 3N_x N_y$ and N_x and N_y correspond to the number of pixels with respect to width and height of the digitized picture for the red, green, and blue color channel will be sufficient.

Definition 2.1. *Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be continuously differentiable. The function $\bar{\nabla}V : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a discrete gradient of V iff it is continuous and*

$$\begin{cases} \langle \bar{\nabla}V(x, x'), (x' - x) \rangle = V(x') - V(x), \\ \bar{\nabla}V(x, x) = \nabla V(x), \end{cases} \quad \text{for all } x, x' \in \mathbb{R}^n.$$

Note, that this definition is different from what is often understood to be a discrete gradient in image processing where the term just refers to a discretized gradient. Definition 2.1 asks for a specific condition. Discrete gradients according to this definition have been studied by many researchers in Geometric Numerical Integration (e.g. [13, 14, 16, 18, 20, 27, 29, 35, 40]). Three well-known discrete gradients are the *midpoint discrete gradient* or *Gonzalez discrete gradient* (cf. [18])

$$(2.1) \quad \begin{aligned} \bar{\nabla}_1 V(x, x') &= \nabla V\left(\frac{x'+x}{2}\right) \\ &+ \frac{V(x') - V(x) - \langle \nabla V\left(\frac{x'+x}{2}\right), x' - x \rangle}{\|x - x'\|^2} (x' - x), \quad (x \neq x'), \end{aligned}$$

the mean value discrete gradient

$$\bar{\nabla}_2 V(x, x') = \int_0^1 \nabla V((1-s)x + sx') ds,$$

that is, for example, used in the averaged vector field method (cf. [16]), and the discrete gradient proposed by Itoh & Abe (cf. [25]) that reads

$$(2.2) \quad \bar{\nabla}_3 V(x, x') = \begin{pmatrix} \frac{V(x'_1, x_2, \dots, x_n) - V(x)}{x'_1 - x_1} \\ \frac{V(x'_1, x'_2, x_3, \dots, x_n) - V(x'_1, x_2, \dots, x_n)}{x'_2 - x_2} \\ \vdots \\ \frac{V(x'_1, \dots, x'_{n-1}, x_n) - V(x'_1, \dots, x'_{n-2}, x_{n-1}, x_n)}{\tilde{x}_{n-1} - x_{n-1}} \\ \frac{V(x') - V(x'_1, \dots, x'_{n-1}, x_n)}{x'_n - x_n} \end{pmatrix}, \quad x'_i \neq x_i, \quad i = 1, \dots, n,$$

Note that the Itoh & Abe discrete gradient is derivative-free and hence its computational realization relatively cheap. Besides these discrete gradients, there are many more. For the gradient flow

$$(2.3) \quad \dot{x} = -\nabla V(x), \quad x(0) = x_0$$

every discrete gradient $\bar{\nabla}V$ leads to an associated discrete gradient method

$$(2.4) \quad x_{n+1} - x_n = -\tau_n \bar{\nabla}V(x_n, x_{n+1}),$$

where $\tau_n > 0$ is a time step that might vary from step to step. Due to Definition 2.1 of a discrete gradient, this method preserves the dissipativity of the solution of the gradient system (1.1), that is we have

$$V(x_{n+1}) - V(x_n) = \langle \bar{\nabla}V(x_n, x_{n+1}), (x_{n+1} - x_n) \rangle = -\tau_n \|\bar{\nabla}V(x_n, x_{n+1})\|^2 \leq 0$$

for all steps n and arbitrary $\tau_n > 0$ as a discrete analogue to the decay (1.2) of the continuous solution.

For our numerical illustrations, we will mainly use the Gonzalez and Itoh-Abe discrete gradient. But we would like to stress that the properties just described as well as the theoretical results in the following sections hold for arbitrary discrete gradients - a rich family to pick from.

3. Some properties of discrete gradient methods. The preservation of the dissipativity by a discrete gradient method leads to useful consequences. Before we can state our first result, we need to recall some definitions.

Definition 3.1. A functional $V : \mathbb{R}^n \rightarrow \mathbb{R}$ is called

- coercive iff

$$V(x_n) \rightarrow \infty \quad \text{for} \quad \|x_n\| \rightarrow \infty.$$

- bounded from below *iff* there exists a constant C such that

$$C \leq V(x), \quad \text{for all } x \in \mathbb{R}^n.$$

- convex *iff* for all $x, y \in \mathbb{R}^n$ and $\lambda \in [0, 1]$

$$V(\lambda x + (1 - \lambda)y) \leq \lambda V(x) + (1 - \lambda)V(y).$$

- strictly convex *iff* for all $x, y \in \mathbb{R}^n$, $x \neq y$, and $\lambda \in (0, 1)$

$$V(\lambda x + (1 - \lambda)y) < \lambda V(x) + (1 - \lambda)V(y).$$

Theorem 3.2. *Let ∇V in (2.3) stem from a functional V which is bounded from below, coercive and continuously differentiable. If $\{x_n\}_{n=0}^{\infty}$ is a sequence generated by the discrete gradient method (2.4) with time steps $0 < c \leq \tau_n \leq M < \infty$, then*

$$\lim_{n \rightarrow \infty} \bar{\nabla} V(x_{n+1}, x_n) = \lim_{n \rightarrow \infty} \nabla V(x_n) = 0.$$

There exists at least one accumulation point of the sequence $\{x_n\}_{n=0}^{\infty}$. And for any accumulation point x_ of the sequence $\{x_n\}_{n=0}^{\infty}$, we have $\nabla V(x_*) = 0$.*

Proof. Since V is bounded from below, say by C , and due to the preservation of the dissipativity, we find

$$C \leq V(x_{n+1}) \leq V(x_n) \leq \dots \leq V(x_0), \quad n = 1, 2, 3, \dots$$

and hence the limit

$$\lim_{n \rightarrow \infty} V(x_n) = V_*$$

exists. From Definition 2.1 and the definition of the discrete gradient method in (2.4), we find

$$\begin{aligned} \tau_n \|\bar{\nabla} V(x_{n+1}, x_n)\|^2 &= -\langle \bar{\nabla} V(x_{n+1}, x_n), x_{n+1} - x_n \rangle = V(x_n) - V(x_{n+1}) \\ &= \frac{1}{\tau_n} \langle -\tau_n \bar{\nabla} V(x_{n+1}, x_n), x_{n+1} - x_n \rangle \\ &= \frac{1}{\tau_n} \|x_{n+1} - x_n\|^2 \geq 0 \end{aligned}$$

for all n . By summing these equations from n to $m - 1$, $m > n$, we obtain

$$\sum_{k=n}^{m-1} \tau_k \|\bar{\nabla} V(x_{k+1}, x_k)\|^2 = \sum_{k=n}^{m-1} \frac{1}{\tau_k} \|x_{k+1} - x_k\|^2 = V(x_n) - V(x_m) \leq V(x_0) - V_*$$

and thus

$$\sum_{k=0}^{\infty} \|\bar{\nabla} V(x_{k+1}, x_k)\|^2 \leq \frac{V(x_0) - V_*}{c} < \infty, \quad \sum_{k=0}^{\infty} \|x_{k+1} - x_k\|^2 \leq M(V(x_0) - V_*) < \infty$$

and therefore

$$\lim_{n \rightarrow \infty} (x_{n+1} - x_n) = \lim_{n \rightarrow \infty} \bar{\nabla} V(x_{n+1}, x_n) = 0.$$

The sets V_t defined by

$$V_t := \{x \in \mathbb{R}^n \mid V(x) \leq t\}$$

are empty or compact. Hence the set $V_{V(x_0)}$ is bounded and closed. In particular, $\bar{\nabla}V$ is uniformly continuous on $V_{V(x_0)} \times V_{V(x_0)}$, where we have chosen the usual topology on the product space to coincide with the norm induced by the standard scalar product on \mathbb{R}^{2n} . Therefore, for any $\epsilon > 0$ there exists a δ such that for $\|(x_{n+1}, x_n) - (x_n, x_n)\| = \|x_{n+1} - x_n\| \leq \delta$ we have

$$\|\bar{\nabla}V(x_{n+1}, x_n) - \bar{\nabla}V(x_n, x_n)\| = \|\bar{\nabla}V(x_{n+1}, x_n) - \nabla V(x_n)\| < \epsilon.$$

Since $\lim_{n \rightarrow \infty} \|x_{n+1} - x_n\| = 0$, we find, that for large enough n , we have

$$\|\nabla V(x_n)\| \leq \|\bar{\nabla}V(x_{n+1}, x_n) - \nabla V(x_n)\| + \|\bar{\nabla}V(x_{n+1}, x_n)\| \leq 2\epsilon.$$

Hence, altogether, we conclude

$$\lim_{n \rightarrow \infty} \bar{\nabla}V(x_{n+1}, x_n) = \lim_{n \rightarrow \infty} \nabla V(x_n) = 0.$$

Due to the boundedness of $V_{V(x_0)}$, the sequence $\{x_n\}_{n=0}^{\infty}$ has at least one accumulation point x_* by the Bolzano-Weierstrass theorem. For a subsequence $\{x_{n_l}\}_{l=0}^{\infty}$ with $\lim_{l \rightarrow \infty} x_{n_l} = x_*$, we have

$$0 = \lim_{l \rightarrow \infty} \nabla V(x_{n_l}) = \nabla V(x_*),$$

due to the continuity of ∇V . ■

Theorem 3.2 states that the sequence $\{x_n\}_{n=0}^{\infty}$ generated by any discrete gradient method satisfies $\lim_{n \rightarrow \infty} \nabla V(x_n) = 0$. This property is very important in the minimization of functionals. In image processing, the functionals to be minimized are often convex or even strictly convex. For such functionals any discrete gradient method tends to global minimizers.

Theorem 3.3. *Under the assumptions of Theorem 3.2.*

1. *If V is in addition convex, then a minimizer exists and any accumulation point of the sequence $\{x_n\}_{n=0}^{\infty}$ is a minimizer.*
2. *If V is in addition strictly convex, then*

$$\lim_{n \rightarrow \infty} x_n = x_*, \quad V(x_*) = \min_x V(x),$$

that is, the sequence of the discrete gradient approximations converges to the unique minimizer.

Proof. It is a standard result, that a continuously differentiable function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex, if and only if

$$V(u) + \langle \nabla V(u), w - u \rangle \leq V(w), \quad \text{for all } u, w \in \mathbb{R}^n.$$

For an accumulation point x_* of the sequence $\{x_n\}_{n=0}^{\infty}$ generated by the discrete gradient method, we have $\nabla V(x_*) = 0$ according to Theorem 3.2 and therefore $V(x_*) \leq V(w)$, for all $w \in \mathbb{R}^n$ which means that x_* is a minimizer of the functional V . There is at least one accumulation point of the sequence $\{x_n\}_{n=0}^{\infty}$ according to Theorem 3.2 and therefore a minimizer exists.

Assume that the function V is strictly convex and that x_* and y_* were two different minimizers of V , that is $x_* \neq y_*$ and $V(x_*) = V(y_*) \leq V(w)$ for all $w \in \mathbb{R}^n$. Since V is strictly convex, pick $\lambda \in (0, 1)$ and we obtain

$$V(\lambda x_* + (1 - \lambda)y_*) < \lambda V(x_*) + (1 - \lambda)V(y_*) = \lambda V(x_*) + (1 - \lambda)V(x_*) = V(x_*).$$

Since $\lambda x_* + (1 - \lambda)y_* \in \mathbb{R}^n$, this is a contradiction to x_* (or y_* , respectively) being a minimizer. Hence the minimizer must be unique and therefore all accumulation points of the sequence $\{x_n\}_{n=0}^{\infty}$, which are minimizers, must be identical. Therefore, the sequence converges to the unique minimizer. ■

4. Nonlinear Examples. In this section, we illustrate the positive effect of the preservation of dissipativity by a series of numerical experiments on standard models in image processing, that involve a gradient flow. In Subsection 4.1, we study the TV denoising (also TV cartooning) functional, whose discretized version is strictly convex. The theory of Section 3 is applicable and we illustrate numerically with the Gonzalez discrete gradient that the discrete gradient method shows the predicted behavior. As a simple method that does not possess the preservation of decay property, the explicit Euler method is used for comparison. The introduction of a blurring kernel in the functional in Subsection 4.1, which leads to a deblurring example, shows the same good effects of the preservation of the decay of the functional. After these basic examples, we provide three more experiments that generalize the application of discrete gradient methods in different ways. In Subsection 4.3 we apply the discrete gradient method to solve TV image inpainting. We discuss its performance using the Itoh & Abe gradient with a simple adaptive step size rule and compare it with the so-called lagged diffusivity method [1], which for convex functionals V shares the dissipation property of the discrete gradient approach [10]. A question that is always important to answer is whether newly proposed methods are useful in actual applications. We therefore study a real-world color denoising example in Subsection 4.4. Finally, in Subsection 4.5 we study non-convex TV denoising [24] computed with the Itoh & Abe discrete gradient method. We include this example to demonstrate the flexibility of the discrete gradient, and its robust structure-preserving properties, which guarantee monotonic decay of V even in the non-convex case.

In what follows, we denote by T_α the continuous functional and by V_α its discretization for which the discrete gradient is computed. The parameter α indicates the dependence of the functional on a positive parameter α that weights the TV regularization against a fitting term to the given image data.

4.1. Grayscale image denoising with TV regularization. For the denoising of a grayscale image, the following gradient descent method has been proposed in [37]. The gradient system is based on the functional

$$(4.1) \quad T_\alpha(u) = \frac{1}{2} \int_{\Omega} (u(x, y) - u_0(x, y))^2 d(x, y) + \alpha TV(u),$$

where we use the smoothed TV functional

$$(4.2) \quad TV(u) = \int_{\Omega} \sqrt{\left(\frac{\partial u}{\partial x}\right)^2 + \left(\frac{\partial u}{\partial y}\right)^2} + \beta d(x, y),$$

with parameter $0 < \beta \ll 1$ as suggested in [1]. Under the smoothness assumption $u \in W^{1,1}$, this leads to the gradient system

$$u_t = \alpha \nabla \cdot \left[\frac{\nabla u}{|\nabla u|_\beta} \right] - (u - u_0) \quad \frac{\partial u}{\partial \mathbf{n}} \Big|_{\partial \Omega} = 0$$

where $|\nabla u|_\beta = \sqrt{|\nabla u|^2 + \beta}$. For the computation, discretization of the continuous functional is necessary. As in [37], we use finite differences, where the homogeneous Neumann boundary conditions are discretized by duplicating the boundary rows and columns of the original picture array. The discretized functional reads

$$(4.3) \quad V_\alpha(u) = \frac{1}{2} \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} (u_{i,j} - (u_0)_{i,j})^2 + \alpha J(u),$$

with

$$(4.4) \quad J(u) = \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \psi \left((D_{ij}^x u)^2 + (D_{ij}^y u)^2 \right),$$

where $\psi(t) = \sqrt{t + \beta}$ and

$$D_{ij}^x u = \frac{u_{i,j} - u_{i-1,j}}{\Delta x}, \quad D_{ij}^y u = \frac{u_{i,j} - u_{i,j-1}}{\Delta y}.$$

Here $u \in \mathbb{R}^{N_x} \times \mathbb{R}^{N_y}$ is the discretized picture. As usual, we identify the matrix $u \in \mathbb{R}^{N_x} \times \mathbb{R}^{N_y}$ with the vector $u \in \mathbb{R}^{N_x N_y}$ by running successively through the columns of $u \in \mathbb{R}^{N_x} \times \mathbb{R}^{N_y}$. The gradient system in $\mathbb{R}^{N_x N_y}$ then follows analogously to the continuous system by computing the gradient:

$$\dot{u} = -\nabla V_\alpha(u), \quad u(0) = u_0.$$

This gradient system satisfies the assumption of our theorems in Section 3.

Lemma 4.1. *The discretized functional V_α in (4.3) is bounded from below, coercive, continuously differentiable and strictly convex.*

Proof. Due to

$$\begin{aligned} & J(\lambda u + (1 - \lambda)v) \\ &= \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sqrt{(D_{ij}^x(\lambda u + (1 - \lambda)v))^2 + (D_{ij}^y(\lambda u + (1 - \lambda)v))^2 + (\lambda\sqrt{\beta} + (1 - \lambda)\sqrt{\beta})^2} \\ &= \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sqrt{(\lambda D_{ij}^x u + (1 - \lambda)D_{ij}^x v)^2 + (\lambda D_{ij}^y u + (1 - \lambda)D_{ij}^y v)^2 + (\lambda\sqrt{\beta} + (1 - \lambda)\sqrt{\beta})^2} \\ &\leq \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sqrt{(\lambda D_{ij}^x u)^2 + (\lambda D_{ij}^y u)^2 + (\lambda\sqrt{\beta})^2} \\ &\quad + \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sqrt{((1 - \lambda)D_{ij}^x v)^2 + ((1 - \lambda)D_{ij}^y v)^2 + ((1 - \lambda)\sqrt{\beta})^2} \\ &= \lambda J(u) + (1 - \lambda)J(v) \end{aligned}$$

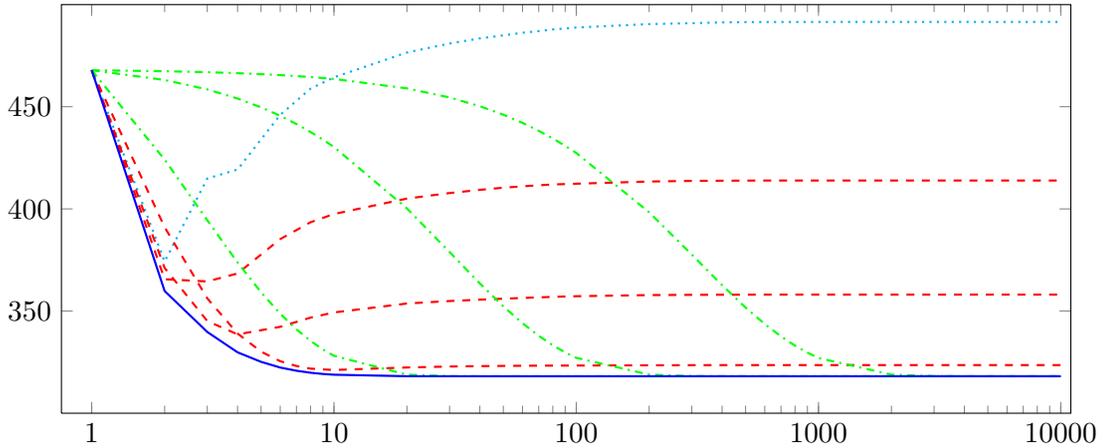


Figure 4.1. TV functional Baboon, explicit Euler method with step size $\tau = 0.5$ (cyan dotted line), explicit Euler method with $\tau = 0.4, 0.3, 0.2$ (red dashed lines, top to bottom), explicit Euler with $\tau = 0.1, 0.01, 0.001$ (green dash-dotted lines, left to right) discrete gradient method (blue) with step size $\tau = 2.5$ (in all experiments: $\alpha = 0.05$, $\beta = 0.001$)

for $\lambda \in [0, 1]$ and u, v two pictures, we find, that J is convex. With two different constant pictures, it is easy to see that J is not strictly convex. By checking that for $u \neq v$ and $t \in [0, 1]$, we have

$$F(t) = \frac{1}{2} \Delta x \Delta y \|tu + (1-t)v - u_0\|^2, \quad F''(t) = \Delta x \Delta y \|u - v\|^2 > 0$$

and hence the functional $\frac{1}{2} \Delta x \Delta y \|u - u_0\|^2$ is strictly convex. Therefore, V_α as a whole is strictly convex. From this first functional, coercivity is obvious. ■

Form Lemma 4.1, we immediately conclude the following corollary of our theorems.

Corollary 4.2. *The function V_α has a unique minimizer and the sequence generated by any discrete gradient method with step sizes $0 < c \leq \tau_n \leq M < \infty$ converges to the minimizer.*

In our examples below, we adopted another simplification which is common in image processing. The domain of the picture is scaled into a rectangular region such that $\Delta x = 1$ and $\Delta y = 1$. Then image data is scaled to $[0, 1]$. In all experiments we chose $\alpha = 0.05$ and $\beta = 0.001$. A first experiment with the Gonzalez discrete gradient shows that the discrete gradient method converges reliably to the minimizer of the total variation functional. In Figure 4.1 the value of the functional is plotted against the number of steps. For the step size $\tau = 2.5$, the discrete gradient method converges in about 10 steps to the same minimum value of the functional as the explicit Euler method in 40 steps with step size $\tau = 0.1$ or in 200 steps with step size $\tau = 0.01$ or in 2000 steps with step size $\tau = 0.001$. The equilibrium computed by the Euler method with step size $\tau = 0.001$ after 10000 steps is used as reference equilibrium picture. The explicit Euler method with step sizes $\tau = 0.2$ and larger stagnates at a larger value of the total variation functional and never converges to the correct value (cf. Figure 4.1). For these step sizes, the value of the functional oscillates, which can not be seen in the figure due to aliasing. The discrete gradient method cannot oscillate, since it is strictly decreasing.

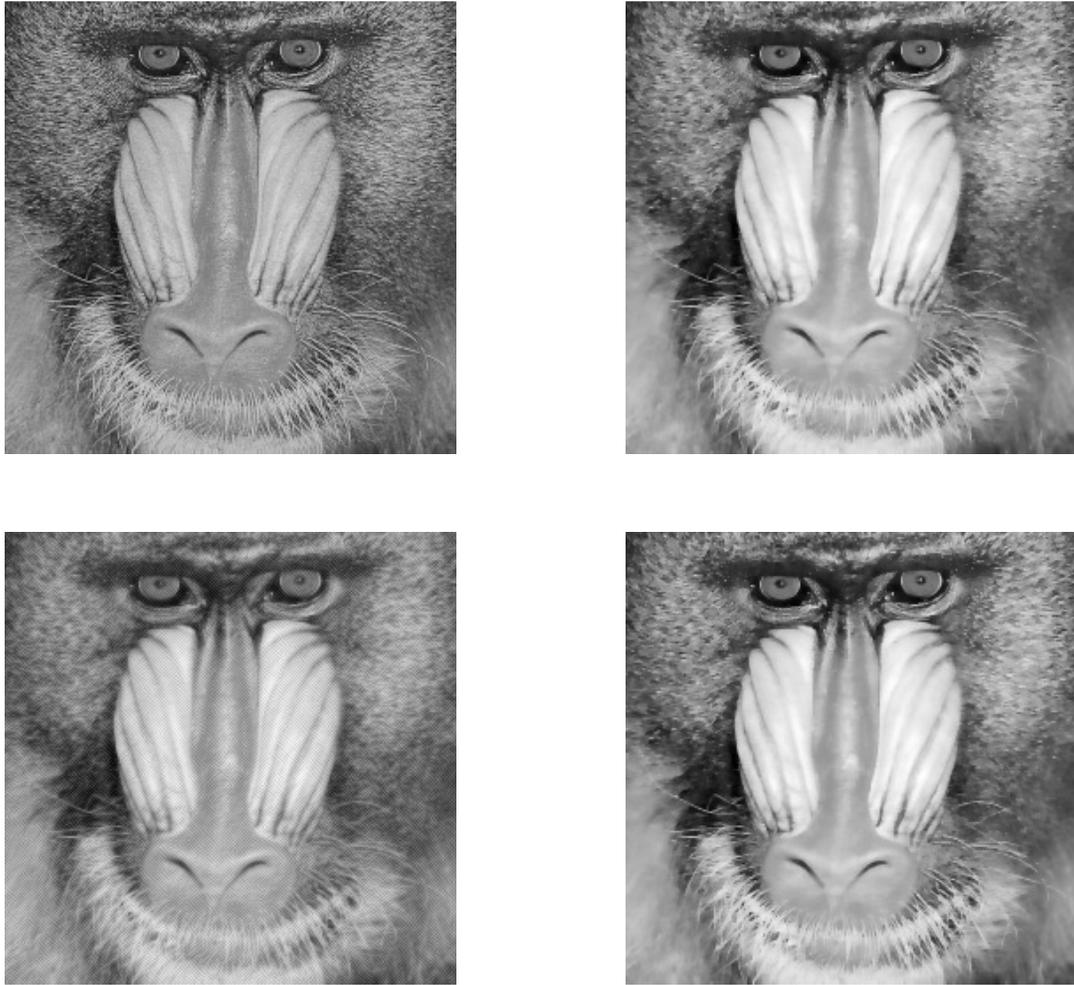


Figure 4.2. Original picture (top left) and TV-denoised reference picture with the explicit Euler method, step-size $\tau = 0.001$ after 10000 steps (top right), explicit Euler denoised image with step-size $\tau = 0.3$ after 10000 steps (bottom left) and discrete gradient denoised baboon with step-size $\tau = 2.5$ after 10 steps (bottom right).

In our experiment, we used the mandrill a. k. a. baboon picture of the USC-SIPI Image Database, [41], turned into a grayscale image. At a close inspection, the difference of the final denoised pictures shown in Figure 4.2 can be seen. The denoised image computed with the Euler method with step size $\tau = 0.3$ (Figure 4.2, bottom left-hand side) shows slightly more details than the reference picture (Figure 4.2, top right-hand side), while the denoised image computed with the discrete gradient method at step size $\tau = 2.5$ (Figure 4.2, bottom right-hand side) fits very well to the reference picture.

4.2. Grayscale image deblurring with TV regularization. Following [1, 9], we minimize the functional

$$T_{\alpha}(u) = \frac{1}{2} \int_{\Omega} ((Ku)(x) - u_0(x))^2 dx + \alpha TV(u),$$

where $TV(u)$ is defined as in (4.2). Under the smoothness assumption $u \in W^{1,1}$, this leads to the parabolic gradient system

$$u_t = \alpha \nabla \cdot \left[\frac{\nabla u}{|\nabla u|^\beta} \right] - K^*(Ku - u_0), \quad \text{with} \quad \frac{\partial u}{\partial \mathbf{n}} \Big|_{\partial \Omega} = 0.$$

The same procedure gives for the discretized functional

$$(4.5) \quad V_\alpha(u) = \frac{1}{2} \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} (Ku_{i,j} - (u_0)_{i,j})^2 + \alpha J(u),$$

where J is given as before in (4.4). We blur the image by convolution with the symmetric kernel K , given below as point spread function (PSF), which corresponds to the resulting image of a single bright pixel under the blurring transformation:

$$(4.6) \quad K = \frac{1}{49} \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}.$$

In order to ensure Neumann boundary conditions, the original image u is embedded in an image with four times the size of the original image, by reflecting the original image over the right-hand side and top boundaries. Then a two-dimensional convolution with the PSF in (4.6) is computed via the Fast Fourier Transform (FFT) and the resulting image of the correct size extracted. Details can be found in the nice introduction [22] to deblurring images. The original and the blurred image can be seen at the top of Figure 4.4. In all experiments we used $\alpha = 0.05$ and $\beta = 0.001$. In Figure 4.3, we first compare the value of the functional for subsequent steps of the Euler method with step size $\tau = 0.19$ (cyan dotted line) to the value of the functional for subsequent steps of the midpoint discrete gradient method (2.1) with step size $\tau = 2.5$ (blue solid line). While the explicit Euler method does not produce a useful result, the image that corresponds to the computed equilibrium by the discrete gradient method is shown in Figure 4.4, the right-hand side picture in the second row. For the solution of the implicit equation in the discrete gradient method, we have used the Newton method in the inner iteration with the exact Jacobian. The resulting linear system is solved by the Conjugate Gradient (CG) method (cf. [23]).

We repeat the experiment with the step sizes $\tau = 0.185, 0.18, 0.175$ (red dashed lines, top to bottom). The explicit Euler method converges for these step sizes. The values of the functional of the explicit Euler method can also be seen in Figure 4.3. The explicit Euler method obviously converges to an incorrect equilibrium value of the function for these step sizes. The explicit Euler method with the step sizes $\tau = 0.1, 0.01, 0.001$ (green dash-dotted lines, left to right) converge to the same equilibrium value of the functional as the discrete gradient method does for the step size $\tau = 2.5$ (blue, solid line). But after 10000

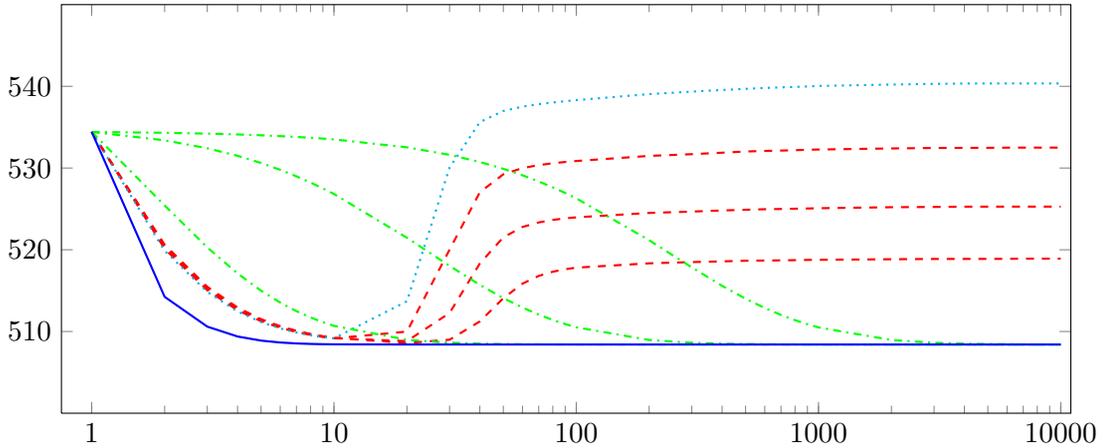


Figure 4.3. TV deblurring functional Baboon, explicit Euler method with step size $\tau = 0.19$ (cyan dotted line), explicit Euler with $\tau = 0.185, 0.18, 0.175$ (red dashed lines, top to bottom), explicit Euler with $\tau = 0.1, 0.01, 0.001$ (green dash-dotted lines, left to right), discrete gradient method (blue solid line) with step size $\tau = 2.5$ (in all experiments: $\alpha = 0.05$, $\beta = 0.001$).

Euler steps with $\tau = 0.001$, the reached value of the functional is 508.4131 and still larger than 508.4069, the value reached by the discrete gradient method with step size $\tau = 2.25$ in step 30 (blue, solid line). An experiment with the discrete gradient method shows that the equilibrium picture computed by the discrete gradient method is the same for the different step sizes $\tau = 2.5, 0.1, 0.185, 0.18, 0.175, 0.01, 0.001$. The equilibrium picture of the discrete gradient method after 10 steps can be seen in Figure 4.4 in the bottom row on the right-hand side. The picture with step-size $\tau = 0.01$ after 10000 steps has been used as a reference for the correct equilibrium and can be seen at the bottom of Figure 4.4 on the left-hand side. The general observation is that the minimum value of the functional (and therefore the equilibrium image) found by the explicit Euler method clearly depends on the step size. This is unpleasant with respect to a reliable computation of minimizers (corresponding to smoothed images) for given smoothing parameters α . On the contrary, the discrete gradient method seems to find the correct minimum value of the functional for a broad range of step sizes, due to the preservation of dissipativity.

4.3. Grayscale image inpainting with TV regularization and adaptive step size. We consider the functional

$$T_\alpha(u) = \frac{1}{2} \int_{\Omega \setminus D} (u(x) - u_0(x))^2 dx + \alpha TV(u),$$

where TV is defined as in (4.2) and D is a subset of Ω in which no information about u_0 is available. This can be the case because either the image is damaged in D or the original scene in the image is occluded by something else in D . The task is to recover the original image in D by minimizing the functional above. This is called image inpainting. The associated

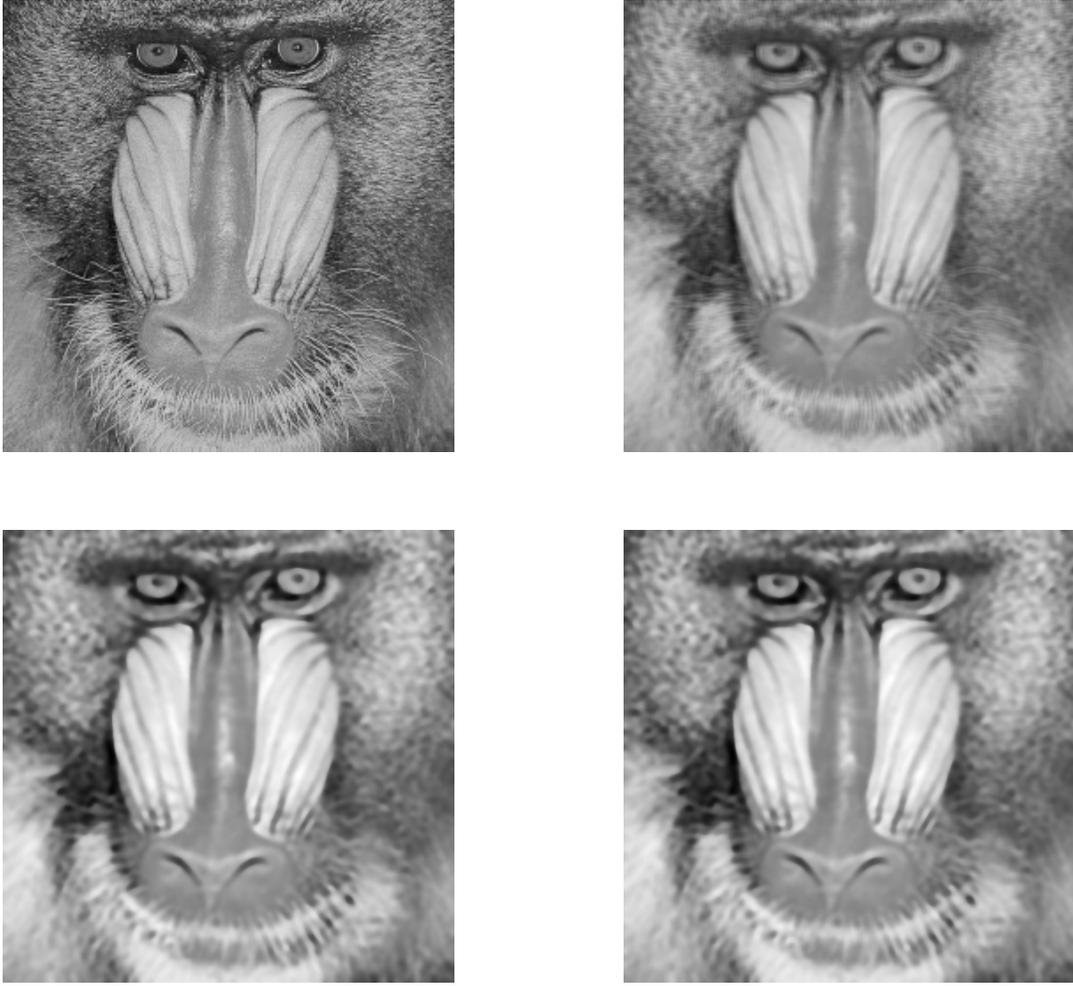


Figure 4.4. Original picture (top left), blurred picture (top right), TV-denoised reference picture with the explicit Euler method, step-size $\tau = 0.01$ after 10000 steps (bottom left), discrete gradient denoised baboon with step-size $\tau = 2.25$ after 10 steps (bottom right).

discretized functional is given by

$$(4.7) \quad V_\alpha(u) = \frac{1}{2} \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \mathbb{P}(u_{i,j} - (u_0)_{i,j})^2 + \alpha J(u),$$

where J is given as before in (4.4) and \mathbb{P} is a projection of the data fitting term onto the set of indices (i, j) that are outside the inpainting domain D .

In the following example, we compare the minimization of the TV-inpainting functional using the discrete gradient method with the Itoh & Abe gradient (2.2) and adaptive step size (as described below) with its minimization by the lagged-diffusivity method [1, 10]. In the latter, a minimizer of the TV-inpainting functional is characterized by a solution of the

corresponding Euler-Lagrange equation and the following fixed-point iteration is performed

$$(4.8) \quad 0 = \operatorname{div} \left(\frac{\nabla u_{n+1}}{\sqrt{|\nabla u_n|^2 + \epsilon}} \right) + \mathbb{1}_{\Omega \setminus D}(u_0 - u_{n+1}),$$

evaluating the nonlinearity in the previous time-step only, and where $\mathbb{1}_{\Omega \setminus D}$ is the characteristic function of the set $\Omega \setminus D$. Here ∇u is discretized as in (4.4) with backward finite differences and its negative adjoint the divergence div by forward finite differences.

For the Itoh & Abe discrete gradient approach we use a simple time step adaptation. In every iteration we compute two trial steps with time steps τ and 2τ and choose the one that decreases V_α most. If the chosen solution corresponds to the time step τ then we halve the time step for the next step, otherwise we double it.

The example in Figure 4.5 is a gray scale image of size 219×292 . The inpainting task is to remove the superimposed text from the image and replace it by the TV-interpolation of the surrounding gray values. Figure 4.6 reports the energy decrease for the Itoh-Abe discrete gradient method compared to the lagged-diffusivity iteration, and the evolution of the step sizes which were adaptively chosen throughout the discrete gradient iterations. In this experiment $\alpha = 0.0001$ and $\beta = 0.01$.

Note that the equations to be solved in an Itoh & Abe update for V_α under the Euclidean inner product as considered here uncouple to scalar equations. Yet, in Figure 4.6 it still appears to choose good descent directions even for large time steps. As one can also observe in Figure 4.6 the energy decrease with lagged-diffusivity is monotonic and faster than under the discrete gradient iteration. This qualitative behavior is representative for the application of lagged-diffusivity to convex functionals V_α . Monotonicity, however, breaks in the case of non-convex functionals for which we will see in Subsection 4.5 the discrete gradient method still preserves monotonic decrease.



Figure 4.5. Original image with superimposed writing and TV-inpainted image with DG method and adaptively chosen τ .

4.4. Multichannel image denoising with TV regularization. We check the fitness of the discrete gradient method for an application in the real world by an experiment with the

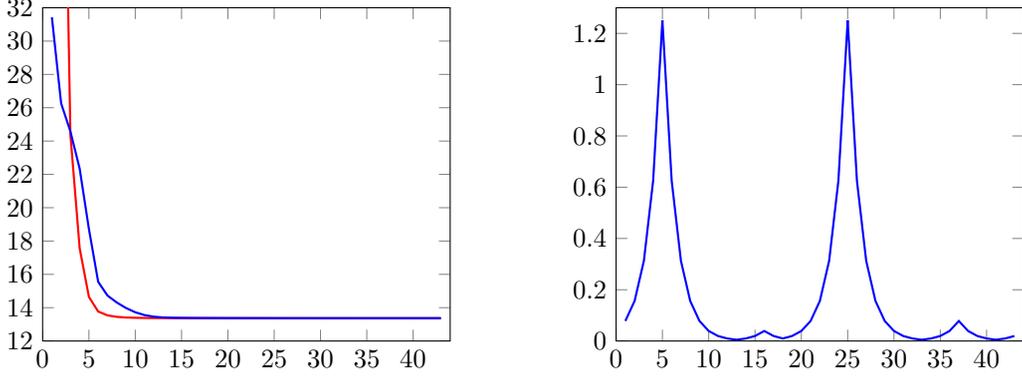


Figure 4.6. Left: Energy decrease for TV -inpainting result in Figure 4.5 with the Itoh-Abe DG method and adaptive step size (blue) and lagged diffusivity with $\tau = 0.1$ (red). Right: Adaptive step sizes for the Itoh-Abe discrete gradient.

discrete gradient method applied to an image processing task in macro photography. We use the multichannel model as described in [11] and first introduced in [5], which uses the TV -functional:

$$TV_2[u] = \left(\sum_{i=1}^p (TV[u_i])^2 \right)^{1/2} = \left(\sum_{i=1}^p \left(\int_{\Omega} |Du_i| dx \right)^2 \right)^{1/2}$$

where $TV(u_i)$ is defined as in (4.2) for p channels u_i , $i = 1, \dots, p$, in the denoising functional

$$T_{\alpha}(u) = \alpha TV_2[u] + \frac{1}{2} \int_{\Omega} \|u - u_0\|^2 dx.$$

With the global constants

$$c_i[u] = \frac{TV[u_i]}{TV_2[u]} \geq 0, \quad i = 1, \dots, p,$$

the Euler-Lagrange equilibrium system reads

$$-\alpha c_i[u] \nabla \cdot \left[\frac{\nabla u_i}{|\nabla u_i|_{\beta}} \right] + (u_i - u_{0,i}) = 0, \quad \frac{\partial u_i}{\partial \mathbf{n}} \Big|_{\partial \Omega} = 0, \quad i = 1, \dots, p.$$

Time-marching leads to the gradient system

$$\frac{d}{dt} u_i = \alpha \cdot c_i[u] \nabla \cdot \left[\frac{\nabla u_i}{|\nabla u_i|_{\beta}} \right] - (u_i - u_{0,i}) = 0, \quad \frac{\partial u_i}{\partial \mathbf{n}} \Big|_{\partial \Omega} = 0, \quad i = 1, \dots, p.$$

The discretized system is just given by using the discretized TV -functionals in TV_2 . The corresponding equations are then solved. The situation is analogous to the case of grayscale image denoising.



Figure 4.7. Original (left) and denoised (right) lice picture, © Volker Grimm

Proposition 4.3. *The functional V_α corresponding to the discretized multichannel TV denoising functional possesses a unique minimizer and the sequence generated by any discrete gradient method with step sizes $0 < c \leq \tau_n \leq M < \infty$ converges to the unique minimizer.*

Proof. With the help of Lemma 4.1, one can check that the discretized functional is bounded from below, coercive, continuously differentiable and strictly convex. The statement then follows from Theorem 3.2. ■

Encouraged by the results for the smaller test images before, we apply the discrete gradient method with the midpoint discrete gradient to a real world denoising problem. The picture at the top of Figure 4.7 is an original photography of some plant lice. The picture has been taken with a strong macro lens, the Canon MP-E 65mm macro lens, that exhibits an extremely low depth-of-field, ranging from 2.24mm at f/16 at 1x magnification, and a minimum of 0.048mm at f/2.8 at 5x magnification. As a camera, a Canon EOS 550D camera has been used, handheld in full sunlight, with an exposure time of 1/250 and f-stop number 14 at 3x magnification. The film speed has been set to ISO 6400, which was needed due to make an exposure time of 1/250 possible. The drawback of this approach to take macro photos without flash is that the high film speed produces a lot of noise due to the necessary amplification of the signal from the charge-coupled device (CCD) image sensor. This real-life noise can clearly be seen in the picture at the top of Figure 4.7 and in the picture detail on the left-hand side in Figure 4.8. The image size in width \times height is 5184×3456 . The overall denoising gradient system for an RGB picture therefore is of dimension $n = 3 \times 5184 \times 3456 = 53747712$.

Despite the size of the system, the discrete gradient method preserves the dissipativity and converges in 10 steps with step size $\tau = 0.2$ to the equilibrium picture. The image has been rescaled to pixel size $\Delta x = \Delta y = 1$, the image data has been in the interval $[0, 255]$, and the constants have been chosen as $\beta = 1$ and $\alpha = 100$. In Figure 4.7 and in the detail in Figure 4.8, one can see that the discrete gradient method successfully removes the noise from the original lice photography.

4.5. A non-convex example: grayscale image denoising with TV^p regularization, $0 < p < 1$. To motivate this regularization, we consider first ℓ_0 -minimization, where for $u \in \mathbb{R}^n$

$$\|u\|_0 = \text{the number of non-zero entries in } u,$$

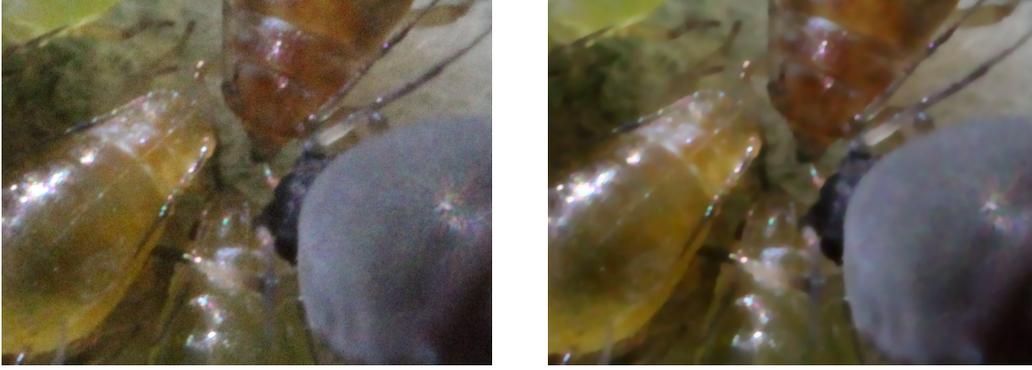


Figure 4.8. Zoom of original image (left) and denoised image (right), © Volker Grimm

which is designed to promote sparsity in u . Solving the ℓ_0 problem is in general NP-hard and therefore its convex relaxation, namely ℓ_1 -minimization, is considered in most sparse reconstruction approaches [8]. In this context, TV regularization can be seen as a convex approximation to a regularization which promotes sparsity of the gradient. Several papers indicate, however, that interesting regularization effects can be observed when studying regularizers that are ℓ_p norms in between ℓ_0 and ℓ_1 , namely penalties of the form

$$\|u\|_p^p = \sum |u_{i,j}|^p, \text{ with } 0 < p < 1,$$

compare [31] for instance. Recently this consideration has been extended to the case of the gradient in [24] where the authors study TV^p regularization, that is

$$TV^p(u) = \|\nabla u\|_p^p = \int_{\Omega} \left(\left(\frac{\partial u}{\partial x} \right)^2 + \left(\frac{\partial u}{\partial y} \right)^2 \right)^{p/2} d(x, y), \text{ with } 0 < p < 1,$$

Analogous to before, for discrete u we consider the discretized and smoothed TV^p functional

$$(4.9) \quad J(u) = \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \psi \left((D_{ij}^x u)^2 + (D_{ij}^y u)^2 \right),$$

with D^x , D^y given as in Section 4.1 and $\psi(t) = (t + \beta)^{p/2}$ with $0 < p < 1$, and its corresponding denoising functional

$$V_{\alpha}(u) = \frac{1}{2} \Delta x \Delta y \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} (u_{i,j} - (u_0)_{i,j})^2 + \alpha J(u).$$

As in Subsection 4.3 we employ the Itoh & Abe discrete gradient with adaptive step size selection. In Figure 4.9 we show a de-noising result with TV^p regularization (4.9) and $p = 0.8$ and $\beta = 0.2$. Since V_{α} is non-convex this time, we consider the behavior of the discrete gradient flow for two different initializations. We initialize the discrete gradient flow once with the noisy

image u_0 and once with a random initialization (randomly choosing the intensity in every pixel of the initial state). For both initializations the discrete gradient flow seems to converge to a decent critical point of V_α , where α was chosen 0.05 for $p = 0.8$ and $\alpha = 0.5$ for $p = 0.2$. In fact, in both cases both critical points seem to converge to a similar energy level, compare Figure 4.10. Note also, as p decreases, the gradient of the image at the computed minimum becomes sparser.

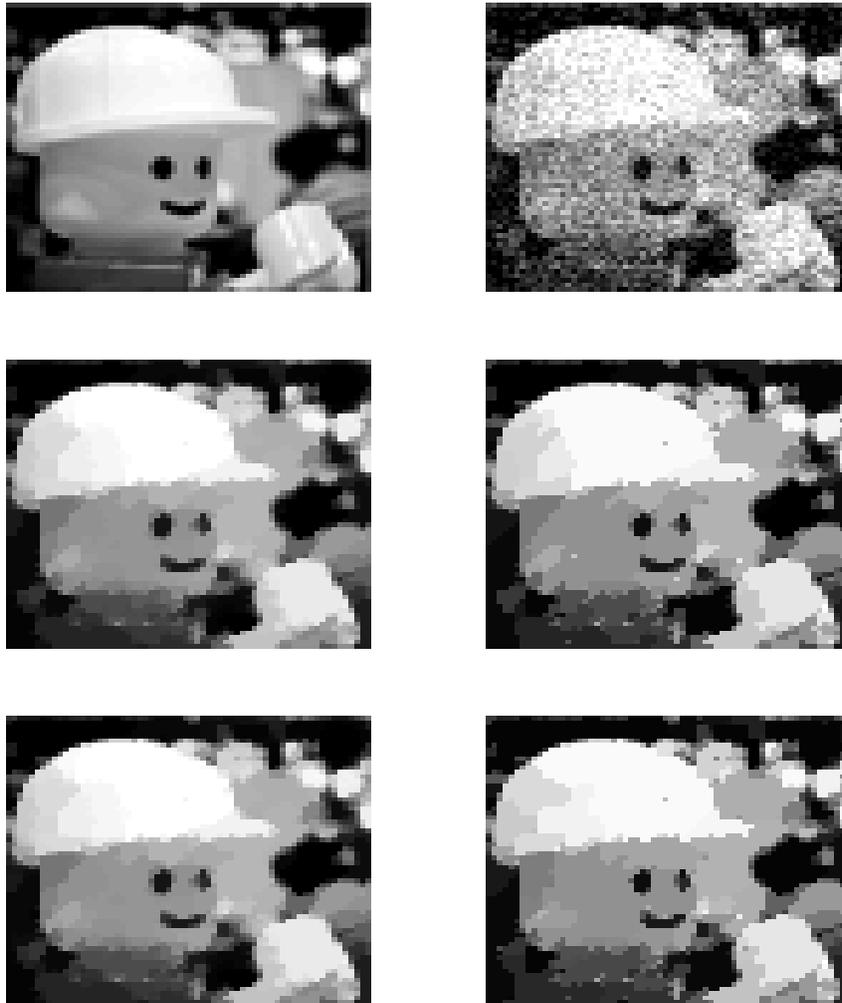


Figure 4.9. TV^p denoising with the Itoh & Abe discrete gradient for $p = 0.8$ (first column) and $p = 0.2$ (second column). First row: Original and noisy image; second row: TV^p denoising result with noisy image u_0 as initial condition; third row: TV^p denoising result with random initial condition.

5. Conclusion. We discussed discrete gradient methods, well-known in Geometric Numerical Integration for the preservation of dissipation in variational equations, with respect to their use in image processing. We assumed that V is smooth which is sufficient when it

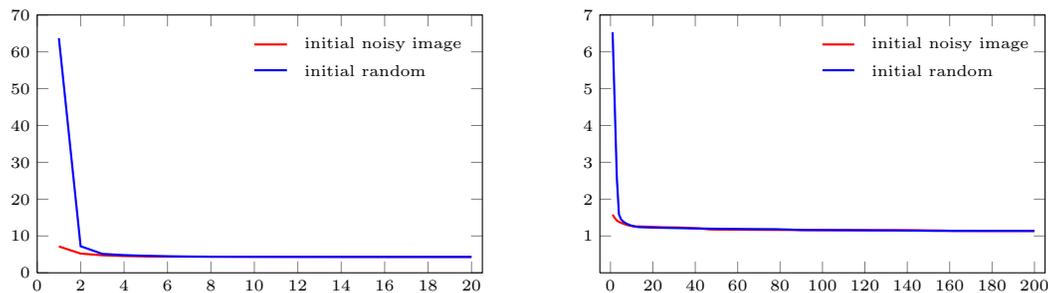


Figure 4.10. TV^p denoising with the Itoh & Abe discrete gradient: energy decrease for the result in Figure 4.9.

comes to actual computations. However, preliminary considerations by the standard techniques suggest this assumption may be weakened significantly. Note also, that in this paper we consider gradient flows of V with respect to the Euclidean inner product only. This can be however generalized, cf. [16], to gradient flows with respect to other inner products as they appear in image processing such as H^{-1} gradient flows [7, 32] or Wasserstein gradient flows [2, 6, 15, 17, 26, 34, 36, 38], just to name a few. We believe that the presented theory, that guarantees the convergence to the equilibrium of any discrete gradient method for a wide range of functionals and gradient flows used in image processing, as well as the conducted experiments indicate that discrete gradient methods could be very interesting for image processing tasks.

REFERENCES

- [1] R. Acar and C. R. Vogel. Analysis of bounded variation penalty methods for ill-posed problems. *Inverse Problems*, 10(6):1217–1229, 1994.
- [2] M. Benning, L. Calatroni, B. Düring, and C.-B. Schönlieb. A primal-dual approach for a total variation Wasserstein flow. In *Geometric Science of Information*, pages 413–421. Springer, 2013.
- [3] A. L. Bertozzi, S. Esedoglu, and A. Gillette. inpainting of binary images using the Cahn-Hilliard equation. *IEEE Transactions on image processing*, 16(1):285–291, 2007.
- [4] A. L. Bertozzi and A. Flenner. Diffuse interface models on graphs for classification of high dimensional data. *Multiscale Modeling & Simulation*, 10(3):1090–1118, 2012.
- [5] P. Blomgren and T. F. Chan. Total variation methods for restoration of vecor-valued images. *IEEE Trans. Image Process*, 7:304–309, 1998.
- [6] M. Burger, M. Franek, and C.-B. Schönlieb. Regularized regression and density estimation based on optimal transport. *Applied Mathematics Research eXpress*, 2012(2):209–253, 2012.
- [7] M. Burger, L. He, and C.-B. Schönlieb. Cahn-Hilliard inpainting and a generalization for grayvalue images. *SIAM Journal on Imaging Sciences*, 2(4):1129–1167, 2009.
- [8] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.
- [9] A. Chambolle and P. L. Lions. Image recovery via total variational minimization and related problems. *Numer. Math.*, 76:167–188, 1997.
- [10] T. F. Chan and P. Mulet. On the convergence of the lagged diffusivity fixed point method in total variation image restoration. *SIAM Journal on Numerical Analysis*, 36(2):354–367, 1999.
- [11] T. F. Chan and J. Shen. *Image processing and analysis: Variational, PDE, wavelet, and stochastic*

- methods*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2005.
- [12] T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. Image Process.*, 10(2):266–277, 2001.
 - [13] J. L. Cieřliński. Locally exact modifications of discrete gradient schemes. *Phys. Lett. A*, 377(8):592–597, 2013.
 - [14] D. Cohen and E. Hairer. Linear energy-preserving integrators for Poisson systems. *BIT*, 51(1):91–101, 2011.
 - [15] B. Düring and C.-B. Schönlieb. A high-contrast fourth-order pde from imaging: numerical solution by ADI splitting. *Multi-scale and High-Contrast Partial Differential Equations*, H. Ammari et al.(eds.), pages 93–103, 2012.
 - [16] E. Celledoni, V. Grimm, R. I. McLachlan, D. I. McLaren, D. O’Neale, B. Owren and G. R. W. Quispel. Preserving energy resp. dissipation in numerical PDEs using the “average vector field” method. *J. Comput. Phys.*, 231(20):6770–6789, 2012.
 - [17] S. Ferradans, N. Papadakis, G. Peyré, and J.-F. Aujol. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences*, 7(3):1853–1882, 2014.
 - [18] O. Gonzalez. Time integration and discrete Hamiltonian systems. *J. Nonlinear Sci.*, 6(5):449–467, 1996.
 - [19] V. Grimm, S. Henn, and K. Witsch. A higher-order PDE-based image registration approach. *Numer. Linear Algebra Appl.*, 13(5):399–417, 2006.
 - [20] E. Hairer and Ch. Lubich. Energy-diminishing integration of gradient systems. *IMA Journal of Numerical Analysis*, 34:452–461, 2014.
 - [21] E. Hairer, Ch. Lubich, and G. Wanner. *Geometric numerical integration*, volume 31 of *Springer Series in Computational Mathematics*. Springer, Heidelberg, 2010. Second Edition.
 - [22] P. C. Hansen, J. G. Nagy, and D. P. O’Leary. *Deblurring images; Matrices, spectra, and filtering*, volume 3 of *Fundamentals of Algorithms*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2006.
 - [23] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Research Nat. Bur. Standards*, 49:409–436 (1953), 1952.
 - [24] M. Hintermüller and T. Wu. Nonconvex TV^q -models in image restoration: Analysis and a trust-region regularization–based superlinearly convergent solver. *SIAM Journal on Imaging Sciences*, 6(3):1385–1415, 2013.
 - [25] T. Itoh and K. Abe. Hamiltonian-conserving discrete canonical equations based on variational difference quotients. *J. Comput. Phys.*, 76(1):85–102, 1988.
 - [26] J. Lellmann, D. A. Lorenz, C.-B. Schnlieb, and T. Valkonen. Imaging with Kantorovich–Rubinstein Discrepancy. *SIAM Journal on Imaging Sciences*, 7(4):2833–2859, 2014.
 - [27] T. Matsuo and D. Furihata. Dissipative or conservative finite-difference schemes for complex-valued nonlinear partial differential equations. *J. Comput. Phys.*, 171(2):425–447, 2001.
 - [28] R. McLachlan and R. Quispel. Six lectures on the geometric integration of ODEs. In *Foundations of computational mathematics (Oxford, 1999)*, volume 284 of *London Math. Soc. Lecture Note Ser.*, pages 155–210. Cambridge Univ. Press, Cambridge, 2001.
 - [29] R. I. McLachlan, G. R. W. Quispel, and N. Robidoux. Geometric integration using discrete gradients. *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.*, 357(1754):1021–1045, 1999.
 - [30] O. Michailovich, Y. Rathi, and A. Tannenbaum. Image segmentation using active contours driven by the Bhattacharyya gradient flow. *IEEE Trans. Image Process.*, 16(11):2787–2801, 2007.
 - [31] M. Nikolova, M. K. Ng, and C.-P. Tam. Fast nonconvex nonsmooth minimization methods for image restoration and reconstruction. *Image Processing, IEEE Transactions on*, 19(12):3073–3088, 2010.
 - [32] S. Osher, A. Solé, and L. Vese. Image decomposition and restoration using total variation minimization and the H^{-1} . *Multiscale Modeling & Simulation*, 1(3):349–370, 2003.
 - [33] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:629–639, 1990.
 - [34] G. Peyré, J. Fadili, and J. Rabin. Wasserstein active contours. In *Image Processing (ICIP), 2012 19th IEEE International Conference on*, pages 2541–2544. IEEE, 2012.
 - [35] G. R. W. Quispel and G. S. Turner. Discrete gradient methods for solving ODEs numerically while preserving a first integral. *J. Phys. A*, 29(13):L341–L349, 1996.
 - [36] J. Rabin, G. Peyré, J. Delon, and M. Berton. Wasserstein barycenter and its application to texture

- mixing. In *Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer, 2011.
- [37] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60:259–268, 1992.
- [38] B. Schmitzer and Ch. Schnörr. Object segmentation by shape matching with Wasserstein modes. In *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 123–136. Springer, 2013.
- [39] R. Strzodka, M. Droske, and M. Rumpf. Image registration by a regularized gradient flow. A streaming implementation in DX9 graphics hardware. *Computing*, 73(4):373–389, 2004.
- [40] A. M. Stuart and A. R. Humphries. *Dynamical systems and numerical analysis*, volume 2 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 1996.
- [41] The USC-SIPI Image Database, available at: <http://sipi.usc.edu/services/database/Database.html>.
- [42] J. Weickert. *Anisotropic diffusion in image processing*. European Consortium for Mathematics in Industry. B. G. Teubner, Stuttgart, 1998.
- [43] Ch. Xu and J. L. Prince. Snakes, shapes, and gradient vector flow. *IEEE Trans. Image Process.*, 7(3):359–369, 1998.