

Depth map estimation with consistent normals from stereo images^{*}

Alexander Malyshev

University of Bergen, Bergen, Norway

Abstract. The total variation regularization of non-convex data terms in continuous variational models can be convexified by the so called functional lifting, which may be considered as a continuous counterpart of Ishikawa's method for multi-label discrete variational problems. We solve the resulting convex continuous variational problem by the augmented Lagrangian method. Application of this method to the dense depth map estimation allows us to obtain a consistent normal field to the depth surface as a byproduct. We illustrate the method with numerical examples of the depth map estimation for rectified stereo image pairs.

Keywords: Consistent normal vector field · Point cloud · Augmented Lagrangian method.

1 Introduction

Estimation of the depth map for a three-dimensional (3D) scene is a major task of computer vision [18, 17]. For example, depth maps can be generated from several two-dimensional (2D) color or grayscale images of a 3D scene taken by one or more cameras positioned at different space locations. More specifically, given a set of 2D images, which consists of rectified stereo pairs of images taken by a stereo pair of cameras, we have to estimate depth maps for each stereo pair of images and merge them into one 3D point cloud. The point cloud is used afterwards to reconstruct a 2D surface of the scene.

Efficient surface reconstruction methods [4, 8, 10, 11], which produce 2D surfaces from 3D point clouds, most often use approximate normal vectors to the reconstructed surfaces. Good surface reconstructions are obtained by means of consistent normal maps [9, 2, 7], which respect edges and other feature points on the reconstructed surface, i.e. the normal vectors are smooth within the smooth parts of the surface but discontinuous across the edges and corners on the surface.

We estimate the depth map in the form of the global minimum of suitable functionals, which are non-convex in general. The global variational approaches to depth estimation can be more attractive than the faster local procedures because they are more robust to various image corruptions and provide wider

^{*} Supported by the Research Executive Agency of the European Commission, grant 778035 - PDE-GIR - H2020-MSCA-RISE-2017.

regularization conditions. However, finding global minima of non-convex functionals may be very difficult or even not feasible in practice.

Our estimation of the depth map uses the continuous variational model

$$\arg \min_u \int_{\Omega} \alpha |u_x(x)| dx + \int_{\Omega} \rho(u(x), x) dx, \quad (1)$$

where the first term is called the total variation (TV) regularization introduced in [16]. The value $|u_x(x)| = \sqrt{u_{x_1}^2 + \dots + u_{x_n}^2}$ denotes the Euclidean length of the gradient vector $u_x = [u_{x_1}, \dots, u_{x_n}]$. The function $\rho(t, x)$ is supposed to be not convex in variable t . Since the point-wise global minima of $\rho(u(x), x)$ are very noisy functions, the total variation times a well-chosen parameter $\alpha > 0$ enforces necessary smoothness in the solution $u(x)$ of (1).

The authors of [13] have proposed an elegant convexification procedure called the functional lifting, which reduces the non-convex model (1) to a convex variational model but with an extra dimension. Further developments and extensions are found in [14, 12]. A similar method for discrete variational models has been earlier proposed by H. Ishikawa [5, 6]. The paper [13] also contains a detailed comparison of the continuous functional lifting with Ishikawa's method.

The present paper contributes to the estimation of the depth maps by extending the results from [13]. The numerical method of [13] is replaced by a faster method, which is a variant of the augmented Lagrangian method (ALM); cf. [19]. We point out that the convexification from [13] is not directly suitable for ALM and must be refined as in Theorem 1. A consistent normal field to the depth map surface is obtained as a byproduct of ALM; see formula (20).

2 Convex relaxation for continuous variational models

Similar to [13], the theoretical arguments below are not fully rigorous but rather informal. For instance, the expression of the total variation, formally valid for functions from the Sobolev space H^1 , is applied to functions, which are not necessarily in H^1 , etc.

The variational model (1) defines an unknown real-valued function $u: \Omega \rightarrow [a, b]$ of bounded variation in the rectangular domain $\Omega = [0, L_1] \times [0, L_2] \times \dots \times [0, L_n] \subset \mathbb{R}^n$. Note that the regularization parameter $\alpha > 0$ is introduced for convenience only.

The function $\rho(t, x)$ can be non-convex in variable t , which creates serious difficulties when developing reliable numerical methods for solving (1). Fortunately, the non-convex variational model (1) can be reformulated into a convex form by adding an extra dimension, say t , to the available dimensions x_1, x_2, \dots, x_n . Such convexification uses special binary functions called the indicator functions of superlevel sets. Namely, for a given function $u(x)$, the indicator function of superlevel sets $\phi: [a, b] \times \Omega \rightarrow \{0, 1\}$ is defined as

$$\phi(t, x) = \begin{cases} 1, & \text{if } u(x) \geq t, \\ 0, & \text{if } u(x) < t. \end{cases} \quad (2)$$

The function $\phi(x, y)$ is binary and monotonically non-increasing with respect to the variable t . Owing to monotonicity, the original function $u(x)$ is reconstructed from $\phi(t, x)$ via the formula

$$u(x) = a + \int_a^b \phi(t, x) dt. \quad (3)$$

Theorem 1 (Functional uplifting). *If $u(x)$ is a solution of the minimization problem (1), then the indicator function of superlevel sets $\phi(t, x)$, constructed in (2), is a solution to the following minimization problem:*

$$\arg \min_{\phi \in \Phi_{\{0,1\}}} \int_{\Omega} \int_a^b \alpha |\phi_x(t, x)| - \rho(t, x) \phi_t(t, x) dx dt, \quad (4)$$

$$\begin{aligned} \Phi_{\{0,1\}} = \{ \phi \mid \phi(t, x) : [a, b] \times \Omega \rightarrow \{0, 1\}; \phi(a, x) = 1 \text{ and } \phi(b, x) = 0 \forall x; \\ \phi(t_1, x) \geq \phi(t_2, x) \text{ whenever } t_1 < t_2 \}. \end{aligned} \quad (5)$$

The converse is also true: if ϕ is the indicator function of superlevel sets associated with u and solves the minimization problem (4)-(5), then u solves (1).

Proof. Following [13], the co-area formula from [3] gives

$$\int_{\Omega} |u_x(x)| dx = \int_{\Omega} \int_a^b |\phi_x(t, x)| dt dx,$$

where $|\phi_x(t, x)| = \sqrt{\phi_{x_1}^2 + \dots + \phi_{x_n}^2}$. The identity $\phi_t(t, x) = -\delta(u(x) - t)$, where $\delta(\cdot)$ is the Dirac delta function, implies the equalities

$$\int_{\Omega} \rho(u(x), x) dx = \int_{\Omega} \int_a^b \rho(t, x) \delta(u(x) - t) dt dx = \int_{\Omega} \int_a^b -\rho(t, x) \phi_t(t, x) dt dx.$$

Theorem 2 (Convex relaxation). *If $\phi(t, x)$ is a solution of the relaxed (from $\Phi_{\{0,1\}}$ to Φ) minimization problem*

$$\arg \min_{\phi \in \Phi} \int_a^b \int_{\Omega} \alpha |\phi_x(t, x)| - \rho(t, x) \phi_t(t, x) dt dx, \quad (6)$$

$$\begin{aligned} \Phi = \{ \phi \mid \phi(t, x) : [a, b] \times \Omega \rightarrow [0, 1]; \phi(a, x) = 1 \text{ and } \phi(b, x) = 0 \forall x; \\ \phi(t_1, x) \geq \phi(t_2, x) \text{ whenever } t_1 < t_2 \}, \end{aligned} \quad (7)$$

then the binary function

$$\phi^{\theta}(t, x) = \begin{cases} 1, & \text{if } \phi(t, x) \geq \theta, \\ 0, & \text{if } \phi(t, x) < \theta, \end{cases} \quad (8)$$

is a solution of (4)-(5) for all threshold values $\theta \in (0, 1)$.

Proof. The co-area formula

$$\int_{\Omega} |\phi_x(t, x)| dx = \int_{\Omega} \int_0^1 |\phi_x^{\theta}(t, x)| d\theta dx.$$

and differentiation of the identity $\phi(t, x) = \int_0^1 \phi^\theta(t, x) d\theta$ with respect to t as

$$\phi_t(t, x) = \int_0^1 \phi_t^\theta(t, x) d\theta$$

allow us to derive that the energy functional

$$E(\phi) = \int_{\Omega} \int_a^b \alpha |\phi_x(t, x)| - \rho(t, x) \phi_t(t, x) dt dx$$

satisfies the following identity for all $\phi \in \Phi$:

$$\begin{aligned} E(\phi) &= \int_0^1 \left\{ \int_{\Omega} \int_a^b \alpha |\phi_x^\theta(t, x)| - \rho(t, x) \phi_t^\theta(t, x) dt dx \right\} d\theta \\ &= \int_0^1 E(\phi^\theta) d\theta. \end{aligned}$$

When ϕ is a minimizer of the functional E in Φ , the inequality $E(\phi^\theta) \geq E(\phi)$ holds for all $\theta \in (0, 1)$ because $\phi^\theta \in \Phi$. If the measure of the set of all θ for which $E(\phi^\theta) > E(\phi)$ is larger than 0, then $E(\phi) = \int_0^1 E(\phi^\theta) d\theta > E(\phi)$. Hence the measure equals zero.

An argument similar to those in [1, 14] can be used to include θ belonging to the exceptional set of zero measure.

3 The augmented Lagrangian method

Let us introduce a dual function $p(t, x)$ such that $p = (p_0, p_1) = \nabla \phi$, where $\nabla \phi = (\phi_t, \phi_x)$ is the full gradient of $\phi(t, x)$ and $|p_1| = \|p_1\|_2$ is the Euclidean norm of p_1 . Then the problem (6)-(7) is equivalent to the following constrained convex minimization problem

$$\min_{\phi, p} \int_{[a, b] \times \Omega} \alpha |p_1(t, x)| - \rho(t, x) p_0 dt dx, \quad (9)$$

$$p_0 = \phi_t, \quad p_1 = \phi_x; \quad p_0 \leq 0, \quad \phi(a, x) = 1, \quad \phi(b, x) = 0. \quad (10)$$

The convex variational problem (9)-(10) can be solved by means of the augmented Lagrangian method (cf. [19]) with the augmented Lagrangian

$$\mathcal{L}(\phi, p, \lambda) = \int_{[a, b] \times \Omega} \left[\alpha |p_1| - \rho p_0 + \langle \lambda, p - \nabla \phi \rangle + \frac{c}{2} \|p - \nabla \phi\|_2^2 \right] dt dx, \quad (11)$$

where $p = (p_0, p_1)$ and $\lambda = (\lambda_0, \lambda_1)$. The inner product $\langle \lambda, p - \nabla \phi \rangle$ is Euclidean. The constant $c > 0$ is sufficiently large but is not required to tend to infinity.

Each iteration of the ALM method consists of alternative minimizations with respect to ϕ and p and special updates of λ .

Algorithm ALM

1. Set $k = 0$ and initialize $p^0 = 0$ and $\lambda^0 = 0$.
2. Find the solution ϕ^{k+1} of the minimization problem

$$\phi^{k+1} = \arg \min_{\substack{\phi \\ \phi(a,x)=1 \\ \phi(b,x)=0}} \mathcal{L}(\phi, p^k, \lambda^k)$$

3. Find the dual variable p^{k+1} by solving the minimization problem

$$p^{k+1} = \arg \min_{\substack{p=(p_0,p_1) \\ p_0 \leq 0}} \mathcal{L}(\phi^{k+1}, p, \lambda^k)$$

4. Update multiplier λ in accordance with the augmented Lagrangian method

$$\lambda^{k+1} = \lambda^k + c(p^{k+1} - \nabla \phi^{k+1})$$

5. Set $k = k + 1$ and go to step 2.

Let us consider Steps 2 and 3 in detail.

Minimization with respect to ϕ . Step 2 of the ALM algorithm concerns the optimization problem

$$\min_{\phi} \int_{[a,b] \times \Omega} \left[-\langle \lambda, \nabla \phi \rangle + \frac{c}{2} \|p - \nabla \phi\|_2^2 \right] dt dx$$

where $\phi(t, x)$ satisfies the boundary conditions $\phi(a, x) = 1$, $\phi(b, x) = 0$. Standard arguments from the variational calculus yield the Poisson equation

$$\phi_{tt} + \sum_{i=1}^n \phi_{x_i x_i} = (p_0 + \lambda_0/c)_t + \sum_{i=1}^n ((p_1)_i + \lambda_i/c)_{x_i} \quad (12)$$

with the Dirichlet and Neumann boundary conditions

$$\phi(a, x) = 1, \quad \phi(b, x) = 0, \quad (13)$$

$$\phi_{x_i}(t, x) = ((p_1)_i + \lambda_i/c)(t, x) \quad \text{if } x_i = 0 \text{ or } x_i = L_i. \quad (14)$$

Minimization with respect to p . Step 3 of the ALM algorithm requires solution to the pointwise minimization problem

$$\begin{aligned} & \arg \min_{p_0, p_1: p_0 \leq 0} \alpha \|p_1\|_2 - \rho p_0 + \langle \lambda, p - \nabla \phi \rangle + \frac{c}{2} \|p - \nabla \phi\|_2^2 \\ & = \arg \min_{p_0, p_1: p_0 \leq 0} \frac{2\alpha}{c} \|p_1\|_2 - \frac{2\rho}{c} p_0 + \|p - (\nabla \phi - \lambda/c)\|_2^2. \end{aligned} \quad (15)$$

Let us denote

$$q_0 = \phi_t - \lambda_0/c \text{ and } q_1 = \phi_x - \lambda_1/c. \quad (16)$$

Solution to $\min_{p_0 \leq 0} -\frac{2\rho}{c}p_0 + |p_0 - q_0|^2$ is $p_0 = \min(q_0 + \rho/c, 0)$.

A simple geometric argument reveals that the minimum point of the function $\frac{2\alpha}{c}\|p_1\|_2 + \|p_1 - q_1\|_2^2$ must have the form $p_1 = rq_1$ with $0 \leq r \leq 1$. For $r \in (0, 1)$ the necessary condition for extremal points is satisfied when

$$\frac{d}{dr} \left[\frac{2\alpha}{c}r + (r-1)^2\|q_1\|_2 \right] = \frac{2\alpha}{c} + 2(r-1)\|q_1\|_2 = 0.$$

Hence $r = 1 - \alpha/(c\|q_1\|_2)$ if $\|q_1\|_2 - \alpha/c > 0$. Otherwise, $r = 0$. As a result, solution to (18) is given by the pointwise formulas

$$p_0 = \min(q_0 + \rho/c, 0), \quad (17)$$

$$(p_1)_i = \begin{cases} [1 - \alpha/(c\|q_1\|_2)](q_1)_i, & \|q_1\|_2 - \alpha/c > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (18)$$

By Theorem 2, an approximation to $u(x)$ can be obtained in the form

$$u(x) = a + \int_a^b \phi^{1/2}(t, x) dt. \quad (19)$$

The normal field to the surface $t - u(x) = 0$ in the space with coordinates (t, x) is the set of the gradient vectors $n(t, x) = [1, -u_{x_1}, \dots, -u_{x_n}]^T$. We use formula (19) to express the derivatives of $u(x)$ in terms of $\phi^{1/2}(t, x)$ as $u_{x_i} = \int_a^b \phi_{x_i}^{1/2}(t, x) dt$. Since the derivatives $\phi_{x_i}(t, x)$ in the AL method are approximated by $(p_1)_i$, the normal field can be approximated as

$$n(t, x) = \frac{v}{\|v\|_2}, \quad \text{where } v(t, x) = \left[1, - \int_a^b (p_1)_1(t, x) dt, \dots, - \int_a^b (p_1)_n(t, x) dt \right]^T. \quad (20)$$

Formula (20) is one of the main contributions of the paper.

4 Numerical experiments

The meaning of the variable x is changed in Section 4. Instead of the points $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, we deal with the pixels $(x, y) \in \mathbb{R}^2$.

4.1 Depth map from stereo pairs of images

Our numerical experiments deal with estimation of the dense depth map; see, e.g., [17, 18] for more detail about depth estimation. Suppose that two functions $L(x, y)$ and $R(x, y)$ represent a pair of left and right 2-dimensional grayscale images with the horizontal coordinates $0 \leq x \leq L_x$ and vertical coordinates $0 \leq y \leq L_y$. We assume that the images L and R are rectified so that one can define the disparity map $t = t(x, y)$ as the parallel translation along the horizontal direction such that a point (x, y) in the left image L matches the

point $(x - t(x, y), y)$ in the right image R . The disparity map is converted into the depth map via the simple formula $depth = focus \cdot baseline / disparity$, where $baseline$ is the distance between the two camera positions, from which the images L and R have been taken. We assume that both cameras are identical, have focal length $focus$ and equally oriented in 3D space.

In order to estimate the disparity map for grayscale images, we use the simplest dissimilarity function

$$\rho(t, x, y) = |L(x, y) - R(x - t(x, y), y)|. \quad (21)$$

The range of the variable t in (21) for fixed variables x and y is the interval $\max(x - L_x, a) \leq t \leq \min(x, b)$ because the first argument $x - t$ of R must satisfy $0 \leq x - t \leq L_x$. The function $\rho(t, x, y)$ equals zero outside this range. More sophisticated dissimilarity functions can be found in [17, 18].

The dissimilarity function for RGB images can be constructed as the sum of absolute differences of intensities for all three color channels

$$\rho(t, x, y) = \sum_{C=\{R,G,B\}} |L_C(x, y) - R_C(x - t(x, y), y)|.$$

4.2 About solving the Poisson equation in rectangular domains

The Poisson equation (12) with the Dirichlet boundary conditions (13) and the Neumann boundary conditions (14) can be solved by several numerical methods.

When the size \mathcal{N} of the discretized function $\phi(t, x, y)$ is not very large, the Poisson equation can be solved by means of the fast discrete sine and cosine transforms. The arithmetic complexity of the fast Poisson solver is $O(\mathcal{N} \log \mathcal{N})$.

For very large \mathcal{N} , the Poisson equation can be efficiently solved by the multi-grid method, which has the linear arithmetic complexity $O(\mathcal{N})$.

4.3 Synthetic dataset

In order to demonstrate behaviour of the proposed method, we use a synthetic dataset, where all functions are constant along the axis y . This dataset is convenient for visualization because it is enough to plot only sections of such functions for a fixed y . These sections are the horizontal lines of the images displaying the functions. The depth function $d(x, y) = 10 + \sin(\text{mod}(x, \pi))$ of the synthetic dataset is defined for $|x| \leq 5\pi$ and defines the surface $z = d(x, y)$. Two identical cameras with focus equal to 1 are located at the points $(x, y) = (-2, 0)$ and $(x, y) = (2, 0)$. The coordinate axes x , y and z of both cameras are oriented in parallel to the axes x , y and z of the surface $z = d(x, y)$. The left and right digital images $L(x, y)$ and $R(x, y)$ have size 128×10 . The horizontal lines of L and R for a fixed y are shown in Fig. 1. Note that the most left 15 pixels of L and the most right 15 pixels of R on each horizontal line are outside the scene. The horizontal lines of the true disparity are shown on the right side of Fig. 2.

The augmented Lagrangian method has been implemented with the fast Poisson solver, which is based on the fast Fourier transform. The following parameters




Fig. 1. Horizontal lines of the rectified stereo images L and R




Fig. 2. The true disparity and its horizontal lines

have been used for computation by this method: $a = 13$, $b = 14.35$, $\alpha = 0.7$, $c = 0.1$. The number of nodes along the coordinate directions t , x and y are respectively 129, 128, and 16. The grid steps are $h_t = (b - a)/128$, $h_x = h_y = 1$. We have run 100 iterations of the augmented Lagrangian method. However, a sufficiently good convergence is achieved much earlier, say after 30 iterations.

Fig.3 displays the computed disparity versus the true disparity. The computed disparity is plotted using the blue solid line, and the true disparity using the red dash-dot line. We recall that the displayed disparity is defined in the coordinate system of the left image L . Note also that the most left 15 pixels lie outside the scene, i.e. the parts of both lines for $x = 1, 2, \dots, 15$ are fictitious.




Fig. 3. Disparity computed by the AL method vs. true disparity

This section of the paper aims to justify numerically that the augmented Lagrangian method produces consistent normals to the surfaces. The surface of the synthetic example is the disparity map $d(x, y)$, which is constant along y . The computed solution $u(x, y)$ turns out to be also constant along y . The computed function $(p_1)_2(x, y)$ is zero everywhere. We recall that the normal field $n(x, y)$ is computed by the formulas

$$n(x, y) = v(x, y) / \|v(x, y)\|_2, \text{ where } v(x, y) = [1, -h_t(p_1)_1, -h_t(p_1)_2].$$

Fig. 4 displays the normals along the horizontal line of the disparity map computed by the augmented Lagrangian method.




Fig. 4. Normals to the disparity surface computed by the AL method

4.4 Comparison with the numerical method of [13]

Apart from introducing the functional lifting for convexification, the authors of [13] propose an efficient numerical method for solving the resulting convex variational problem. Provided that $\rho(t, x, y) \geq 0$, they consider the convex model

$$\arg \min_{\phi \in D} \int_a^b \int_{\Omega} \left[\alpha \sqrt{\phi_x(t, x, y)^2 + \phi_y(t, x, y)^2} + \rho(t, x, y) |\phi_t(t, x, y)| \right] dt dx dy, \quad (22)$$

$$D = \{ \phi \mid \phi(t, x, y) : [a, b] \times \Omega \rightarrow [0, 1]; \phi(a, x, y) = 1 \text{ and } \phi(b, x, y) = 0 \}, \quad (23)$$

which is equivalently transformed into the primal-dual formulation

$$\min_{\phi \in D} \max_{p \in C} \int_a^b \int_{\Omega} \left(\phi_x p_1 + \phi_y p_2 + \frac{\rho}{\alpha} \phi_t p_0 \right) dx dy dt, \quad (24)$$

where

$$C = \{ p(t, x, y) \mid p(t, x, y) = (p_0, p_1, p_2)(t, x, y) : [a, b] \times \Omega \rightarrow [0, 1]; \quad (25)$$

$$|p_0(t, x, y)| \leq \frac{\rho(t, x, y)}{\alpha}; \sqrt{p_1^2(t, x, y) + p_2^2(t, x, y)} \leq 1 \}.$$

Such a transformation is obtained by using the vectors $[p_1(t, x, y), p_2(t, x, y)]$ coinciding with the normalized vectors $[\phi_x, \phi_y] / \sqrt{\phi_x^2 + \phi_y^2}$. Apparently, this representation does not lead to formula (20) for a consistent normal field.

The numerical algorithm for solving (24), proposed in [13], is called a primal-dual proximal point (PDPP) method and iteratively minimizes the functional in (24) with respect to the primal variable ϕ and then maximizes the same functional with respect to the dual variable p . Each iteration consists of the following two steps:

$$\begin{aligned} \text{primal step} \quad & \phi^{k+1} = \mathcal{P}_D(\phi^k + \tau_p \operatorname{div} p^k), \\ \text{dual step} \quad & p^{k+1} = \mathcal{P}_C(p^k + \tau_d \nabla \phi^{k+1}). \end{aligned}$$

The operator \mathcal{P}_D denotes the projection onto the set D , which can be computed by a simple truncation of ϕ^{k+1} to the interval $[0, 1]$ and setting $\phi(a, x, y) = 1$ and $\phi(b, x, y) = 0$. The operator \mathcal{P}_C denotes the projection onto the set C , which

can be computed via the formulas

$$p_1^{k+1} = p_1^{k+1} / \max \left(1, \sqrt{(p_1^{k+1})^2 + (p_2^{k+1})^2} \right)$$

$$p_2^{k+1} = p_2^{k+1} / \max \left(1, \sqrt{(p_1^{k+1})^2 + (p_2^{k+1})^2} \right)$$

$$p_0^{k+1} = p_0^{k+1} / \max \left(1, \frac{\alpha |p_0^{k+1}|}{\rho} \right).$$

The parameters τ_p and τ_d must guarantee stability of the method. The authors of [13] propose the choice $\tau_p = \tau_d = 1\sqrt{3}$. For the synthetic dataset, sufficient convergence is observed after 1000 iterations.




Fig. 5. Disparity computed by the PDPP method vs. true disparity

To get the plot on the upper level of Fig.6, the PDPP method was iterated 2000 times, which is much longer than necessary for convergence. The normals have been computed by the formula $n(t, x, y) = [1, -u_x, -u_y]^T$, where $u(x, y)$ is the computed solution of (22).




Fig. 6. Normals computed by the PDPP method are shown on the upper level; normals computed by the AL method are shown on the lower level

The comparison of the normal fields in Fig. 6 demonstrates that both the augmented Lagrangian method and the PDPP method produce normal fields of similar visual quality.

4.5 Tsukuba dataset

Let us use the dataset Tsukuba from <http://vision.middlebury.edu/stereo> in order to demonstrate the working capacity of the AL method for calculation of a consistent normal field. The rectified stereo image pair Tsukuba of the smallest size 384×288 consists of two RGB images. The image from the left camera is shown as the left image in Fig. 7.




Fig. 7. The left image and the true disparity for the Tsukuba dataset

The corresponding image of the true disparity is also available at the Middlebury repository [15]; see the right image in Fig.7. The disparity map has 256 gray levels but it can be scaled to the interval $[0, 15]$ of integer numbers without any loss. The scaled image has only the following gray values (or labels): 0, 5, 6, 7, 8, 10, 11, 14. The interval (measured in pixels) containing all disparity values is $[a, b] = [0, 16]$.

The suitable parameters in the augmented Lagrangian (14) are $\alpha = 0.1$ and $c = 0.1$. The label space $[a, b]$ is divided into 32 subintervals of the length $h_t = (b - a)/32$. We have run only 100 iterations of the AL method. Fig.8 shows the computed disparity map and consistent normals for the Tsukuba rectified stereo pair.

5 Conclusion

We have developed a variant of the augmented Lagrangian method for a convex relaxation of the continuous variational problem with the total variation regularization. The most time consuming part of this method is solving a boundary value problem for the 3D Poisson equation. We solve it by the fast Poisson solver. The augmented Lagrangian method has a significantly faster convergence than the primal-dual proximity point method from [13].

An additional benefit of the augmented Lagrangian method consists in producing a consistent normal vector field to the solution surface as a byproduct. Numerical differentiation of the solution computed by the method from [13]

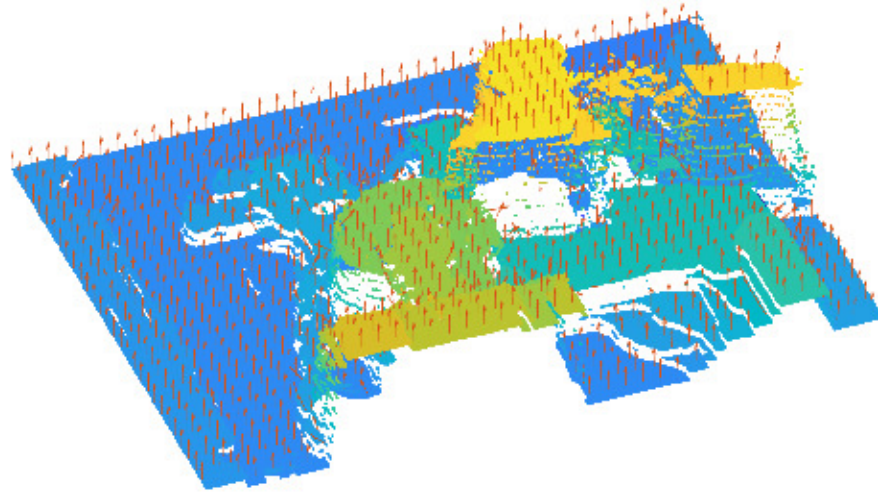


Fig. 8. Computed disparity and consistent normals for the Tsukuba dataset

seems to produce a normal vector field of similar quality. However, thorough verification of these properties requires further investigation.

References

1. Berkels, B.: An unconstrained multiphase thresholding approach for image segmentation. In: Tai, X.C., Moerken, K., Lysaker, M., Lie, K.-A. (eds.) *SSVM 2009 “Scale Space and Variational Methods in Computer Vision”*, LNCS 5567, pp. 26–37, Springer (2009)
2. Castillo, E., Liang, J., Zhao, H.-K.: Point cloud segmentation via constrained non-linear least squares surface normal estimates. Chapter 13. In: Breuss, M., Bruckstein, A., Maragos, P. (eds.) *“Innovations for Shape Analysis: Models and Algorithms”*, Springer, New York (2013)
3. Fleming, W., Rishel, R.: An integral formula for total gradient variation. *Arch. Math.*, **11**, 218–222 (1960)
4. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Surface reconstruction from unorganized points. In: *Proc. 19th Annual Conf. Computer Graphics and Interactive Techniques, ACM SIGGRAPH*. *Computer Graphics*, **26**, pp. 71–78 (1992)
5. Ishikawa, H.: Exact optimization for Markov random fields with convex priors. *IEEE Trans. Pattern Anal. Mach. Intell.* **25**(10), 1333–1336 (2003)
6. Ishikawa, H.: Graph Cuts—Combinatorial Optimization in Vision. In: Lézoray O., Grady, L. (eds.) *“Image Processing and Analysis with Graphs”*, pp. 25–63, CRC Press, (2012)
7. Jakob, J., Buchenau, C., Guthe M.: Parallel Globally Consistent Normal Orientation of Raw Unorganized Point Clouds. *ComputerGraphicsForum*, **38**(5), 163–173 (2019)

8. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proc. Fourth Eurographics Symposium on Geometry Processing, Eurographics Association, pp. 61–70, (2006)
9. König, S., Gumhold, S.: Consistent propagation of normal orientations in point clouds. In: Proc. Vision, Modeling, and Visualization Workshop, Braunschweig, Germany, 9 pp. (2009)
10. Lempitsky V., Boykov, Y.: Global optimization for shape fitting. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 1–8, IEEE Xplore (2007)
11. Liang, J., Park, F., Zhao, H.: Robust and efficient implicit surface reconstruction for point clouds based on convexified image segmentation. *J. Sci. Comput*, **54**(2–3) 577–602 (2013)
12. Möllenhoff, T., Laude, E., Möller, M., Lellmann, J., Cremers, D.: Sublabel-accurate relaxation of nonconvex energies. In: IEEE Conf. Computer Vision and Pattern Recognition (CVPR 2016), IEEE Xplore, pp. 3948–3956, (2016)
13. Pock, T., Schoenemann, T., Graber, G., Bischof, H., Cremers, D.: A convex formulation of continuous multi-label problems. In: Forsyth, D., Torr, P., Zisserman, A. (Eds.) ECCV 2008, Part III, LNCS 5304, pp. 792–805, Springer-Verlag (2008)
14. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: Global solutions of variational models with convex regularization. *SIAM J. Imaging Sci.*, **3**(4), 1122–1145 (2010)
15. <http://vision.middlebury.edu/stereo>
16. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *IEEE Trans. Pattern Anal. Mach. Intell.* **60**(1), 259–268 (1992)
17. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comp. Vis.* **47**(1-3), 7–42 (2002)
18. Szeliski, R.: *Computer Vision: Algorithms and Applications*. Springer (2010)
19. Wu C., Tai, X.-C.: Augmented Lagrangian Method, Dual Methods, and Split Bregman Iteration for ROF, Vectorial TV, and High Order Models. *SIAM J. Imaging Sciences* **3**(3), 300–339 (2010)