Image feature enhancement based on the time-controlled total variation flow formulation

Ovidiu Ghita, Dana E. Ilea, Paul F. Whelan

Vision Systems Group, School of Electronic Engineering

Dublin City University, Dublin 9, Ireland

E-mail: ghitao@eeng.dcu.ie

Abstract

Data smoothing and feature enhancement are two important precursors to many higher-level computer vision applications such as image segmentation and scene understanding. Total variation (TV) flow algorithms are a distinct subcategory of diffusion-based filtering techniques that have been widely applied to reduce the level of noise in the image but not at the expense of poor feature preservation. In this paper we address a number of numerical aspects associated with the TV flow and in particular we are interested to redefine the TV flow regularization in order to reduce the effect of oscillations and improve the convergence of the implementations in the discrete domain. TV flow algorithms are implemented using iterative schemes and one difficult problem is the selection of appropriate criteria to identify the optimal number of iterations. In this paper we show that the application of a time-ageing procedure leads to an elegant formulation were the TV flow algorithms converge naturally to the optimal solution. To evaluate the performance of the proposed algorithm (referred in this paper to as time-controlled (TC)-TV flow), a large number of experiments on various types of natural images were conducted.

Keywords: TV flow, anisotropic diffusion, feature enhancement, numerical stability.

1. Introduction

One important step in the process of image segmentation is to reduce the errors caused by the image noise and local in-homogeneities. Typically, the image noise is reduced by the application of local averaging operators such as the Gaussian, but this approach leads to the introduction of image blur and the attenuation of important contextual features in the image such as edges. To compensate for these undesirable effects associated with linear smoothing strategies, non-linear approaches have been developed to achieve better feature preservation (Grahs et al., 2002; Keeling and Stollberger, 2002; Sonka et al., 1998). Among non-linear techniques, the seminal work detailed in the paper by Perona and Malik (Perona and Malik, 1990) received a special attention from the vision community (Ghita et al., 2005; Ilea and Whelan, 2007; Smolka and Plataniotis, 2002; Weickert et al., 1998). Total variation (TV) flow algorithms have been viewed by many authors as a special case of the geometrical-driven anisotropic diffusion (Dibos and Koepfler, 1999; Gothandaraman et al., 2001; Petrovic et al., 2004; Rudin et al., 1992; Strong and Chan, 2003). The TV flow formulation has been evaluated from a numerical perspective by Breuβ et al., 2006 and they concluded that this data smoothing feature preserving technique is well-posed and it leads to constant signals in finite times. However, in their paper they stressed that the discrete representations of the standard TV flow show instability in areas defined by zero gradients and the implementations in the discrete (image) domain should be approached with care.

In this paper we follow the work detailed in (Breuβ et al., 2006) and we focus on a number of numerical aspects associated with the TV flow formulation with a view of improving the numerical stability and the convergence time. Also in line with other non-linear smoothing strategies such as anisotropic diffusion (Ilea and Whelan, 2007; Smolka and Plataniotis, 2002; Weickert, 1998), the TV flow formulation is implemented as an iterative scheme where the number of iterations is typically a user defined parameter (Andreu et al., 2001; Breuβ et al., 2006; Petrovic et al., 2004). In this paper we show that the application of a time-ageing procedure to the time step size parameter, implements a data smoothing framework where the evolution in time of the TV flow becomes predictable and the algorithm will converge naturally to the optimal result. This paper is organised as follows. In Section 2 the mathematical background behind the TV flow formulation is discussed. Section 3 details the numerical implementation of the TV flow and we approach the algorithmic solutions applied to improve the numerical stability. In Section 4 a number of experimental results are analysed while Section 5 concludes this paper.

2. TV Flow. Mathematical background

As it was mentioned in the previous section, the TV flow is a special case of the anisotropic diffusion function that has been first proposed by Perona and Malik, 1990 (PM).

$$\frac{\partial u}{\partial t} = \nabla\left(g\left(\frac{\nabla u}{\|\nabla u\|} \right) \right) \tag{1}$$

where $u(x,y,t)$ is the processed data at time t and ∇ is the gradient operator. In the PM formulation the diffusion function g is designed to control the smoothing based on the strength of the gradient. In the original paper (Perona and Malik, 1990), this function is implemented using exponential or reciprocal forms and the value of g is bounded in the interval (0, 1], $g \rightarrow 0$ when $\nabla u \rightarrow \infty$. In the TV flow formulation, the functional $g(\nabla u)$ is replaced with ∇u and this simple modification leads to a process where the data is diffused more in image areas with small discontinuities in the intensity data and the diffusive process is stopped in image locations where the gradients have large values.

$$\frac{\partial u}{\partial t} = \nabla\left(\frac{\nabla u}{\|\nabla u\|} \right) \tag{2}$$

In the one-dimensional (1D) case, this process can be formulated as the minimisation of the following functional called TV regularization (Breuβ et al., 2006; Grahs et al., 2002; Strong and Chan, 2003),

$$TV(u) = \int_{\Omega} (u - f)^2 + 2t \left\| \frac{\partial u}{\partial x} \right\| dx \,, \; u(t = 0) = f \tag{3}$$

where f is the initial data, t is a time parameter and Ω is the image domain. The main property of the function illustrated in equation (3) is that it does not penalise the image discontinuities and the smoothing is modulated by the descent of the gradient magnitude. In equation (2) it is noticeable that the TV flow formulation becomes unstable when $\nabla u \to 0$ and to circumvent this problem Breuβ et al., 2006 proposed a numerical implementation using the following regularization,

$$\frac{\partial u}{\partial t} = \nabla \left(\frac{\nabla u}{\sqrt{\beta^2 + \nabla u^2}} \right) \tag{4}$$

where β is a small positive regularization term.

3. Numerical implementation of the TV flow

The implementation of the TV flow in the discrete domain would require a simple approximation of the partial derivatives with the central differences in the (x,y,t) space. In this study, we have extended the approach detailed in Breuβ et al. 2006 to the two-dimensional (2D) grid G as follows,

$$G = \{(x, y, t) \mid x = i\Delta x, y = j\Delta y, t = n\Delta t, i, j, n \in \mathbf{N}, \Delta \mathrm{x} > 0, \Delta \mathrm{y} > 0, \Delta \mathrm{t} > 0\} \tag{5}$$

$$u_{i,j}^{n+1/2} = u_{i,j}^{n} + \frac{\Delta t}{\Delta x^2}\left(\frac{u_{i+1,j}^{n} - u_{i,j}^{n}}{\sqrt{\beta^2 + \left(\frac{u_{i+1,j}^{n} - u_{i,j}^{n}}{\Delta x}\right)^2}} - \frac{u_{i,j}^{n} - u_{i-1,j}^{n}}{\sqrt{\beta^2 + \left(\frac{u_{i,j}^{n} - u_{i-1,j}^{n}}{\Delta x}\right)^2}} \right) \tag{6}$$

$$u_{i,j}^{n+1} = u_{i,j}^{n+1/2} + \frac{\Delta t}{\Delta y^2}\left(\frac{u_{i,j+1}^{n+1/2} - u_{i,j}^{n+1/2}}{\sqrt{\beta^2 + \left(\frac{u_{i,j+1}^{n+1/2} - u_{i,j}^{n+1/2}}{\Delta y}\right)^2}} - \frac{u_{i,j}^{n+1/2} - u_{i,j-1}^{n+1/2}}{\sqrt{\beta^2 + \left(\frac{u_{i,j}^{n+1/2} - u_{i,j-1}^{n+1/2}}{\Delta y}\right)^2}} \right) \tag{7}$$

where n is the iteration index, Δx and Δy are the discrete spatial distances of the image grid, Δt is the time size step parameter and β is the regularization term. If we consider that the datapoint $u_{i,j}$ is a local maxima or minima, we can demonstrate that the discrete implementation of the TV flow is stable (for more details refer to Breuβ et al., 2006) if equation (8) is upheld. This equation indicates that the parameters β and Δt are not mutually independent and they need to be carefully selected.

$$\max\left(\frac{\Delta t}{\Delta x^2}, \frac{\Delta t}{\Delta y^2}\right) \le \frac{\beta}{2} \tag{8}$$

3.1. Time-controlled TV flow formulation

From the two parameters that control the smoothing process (β and Δt), the time size Δt has a larger influence on the intensity of the smoothing process and this observation is motivated by the fact that this parameter has a multiplicative effect. Surprisingly, in the original implementation (Breuβ et al., 2006) this parameter is kept constant during the iterative process and as a result with the increase in the number of iterations the amplitude of oscillations does not decrease monotonically. This can be demonstrated if we consider again that a datapoint $u_{i,j}$ is a local maxima or minima and we set the regularization parameter β to a value closer to zero. Thus, the maximal value that is applied to update the value of the datapoint $u_{i,j}^{n+1}$ in equations (6) and (7) is,

$$\max\left\{\left|u_{i+1,j}^{n}-u_{i,j}^{n}\right|,\left|u_{i,j}^{n}-u_{i-1,j}^{n}\right|,\left|u_{i,j+1}^{n}-u_{i,j}^{n}\right|,\left|u_{i,j}^{n}-u_{i,j-1}^{n}\right|\right\}\le\frac{2\Delta t}{\Delta x} \tag{9}$$

Equation (9) indicates that the oscillations are bounded to a finite value (for the sake of simplicity we assume that $\Delta x=\Delta y$), but there is no guarantee that they will follow a monotonic decrease. In fact, in our experiments we found that the amplitude of oscillations starts to decrease only after the TV flow process is applied for a large number of iterations. Nonetheless, this approach is not appropriate for two reasons. Firstly, with the increase in the number of iterations the computational load required to smooth the image data is high. Secondly, this precludes the application of numerical solutions to identify the optimal number of iterations.

To address these problems we propose an elegant solution where a time controlled (TC) ageing procedure is applied to decrease the value of the parameter Δt after each iteration (this approach is referred to as TC-TV flow). The time controlled ageing procedure has the role of cooling the smoothing process with the increase in the number of iterations and this can be implemented using either the linear (equation (10)) or exponential (equation (11)) function,

$$\begin{cases}\Delta t^{n+1}=\Delta t^{0},\, n=0\\ \Delta t^{n+1}=\Delta t^{n}*\gamma,\ n>0\ \text{ and }\ \gamma\in(0,1]\end{cases} \tag{10}$$

$$\Delta t^{n+1}=\Delta t^{0}*e^{-\frac{n}{\gamma}}\ ,\qquad \gamma>0 \tag{11}$$

where Δt^0 is the initial value of the time size parameter (at iteration zero). In our experiments the implementation detailed in equation (10) proved to be more stable with respect to the selection of the γ parameter and in this paper the reported results are obtained when the ageing procedure is implemented using the linear form. Figure 1 illustrates the results when the standard TV (S-TV) flow and our TC-TV flow algorithms were applied to a 1D noisy test signal. For clarity purposes, the results obtained after the application of the TV flow algorithms were superimposed on the original noiseless signal.

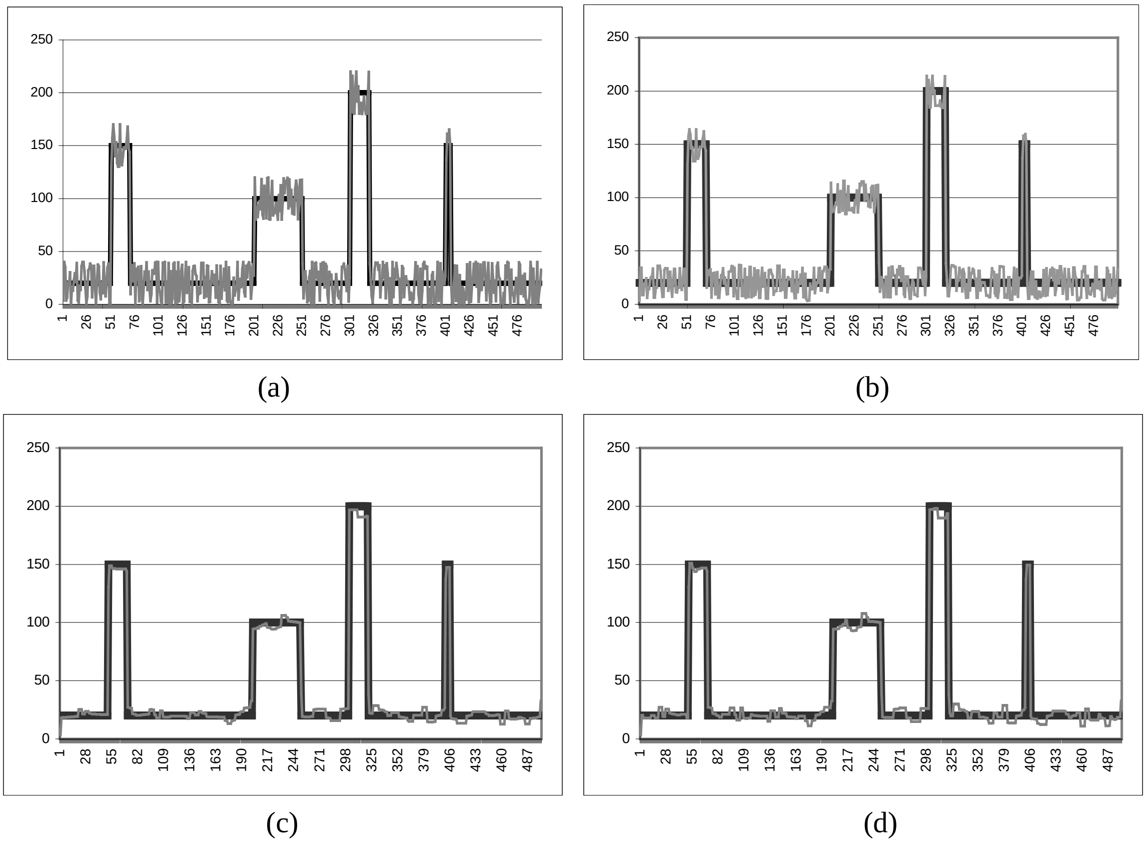

Figure 1. Results of the TV flow. (a) Original signal and the noisy signal. (b,c) The results from the standard S-TV flow when applied to 100 and 500 iterations respectively. (d) The result from our TC-TV algorithm (number of iterations detected automatically using equation 12).

In equations (10) and (11) the parameter γ controls to what extent the smoothing procedure is cooled and it has the effect of lowering the amplitude of the oscillations at each iteration (see equation (9)). This process is illustrated in Figure 2.

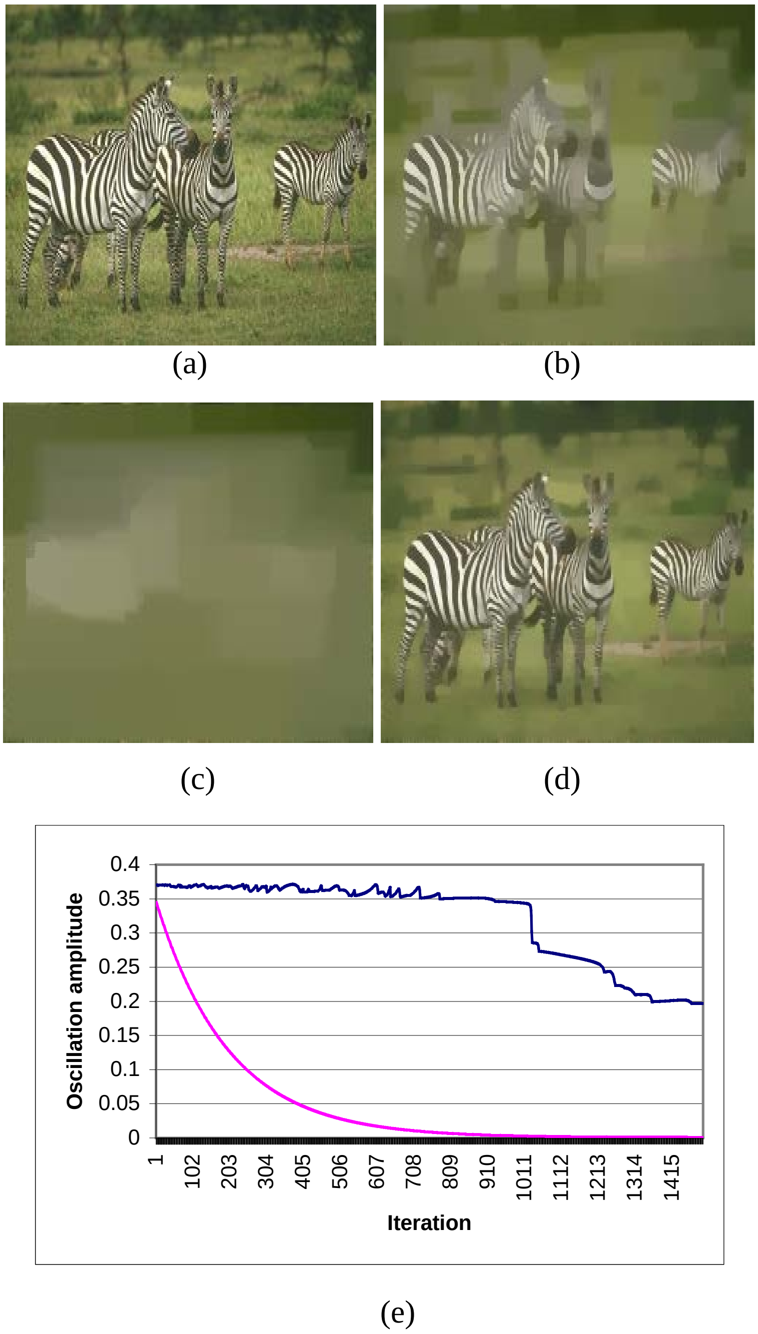

Figure 2. Convergence of the TV flow algorithm. (a) Original image (BSDB, 2001). (b) The result obtained after the application of the S-TV flow procedure (300 iterations). (c) The result obtained after the application of the S-TV flow procedure (1500 iterations). (d) The result obtained after the application of our TC-TV flow algorithm (1500 iterations). (e) The amplitude of oscillations with respect to the number of iterations (in blue/dark – S-TV flow implementation, in red/gray – TC-TV flow implementation). In this experiment the TV flow parameters are set to the following values: $\Delta t = 0.1$, $\gamma = 0.995$, $\beta = 0.1$. Note that all illustrations in this paper are best viewed in colour.

The application of the ageing procedure has another beneficial contribution, namely it allows us to derive a numerical solution for the identification of the optimal number of iterations. In equation (9) it can be observed that the absolute maximal updating value is bounded to $2\,\Delta t^{0}/\Delta x$, thus, taking in consideration that the oscillations can be either positive or negative we can set the convergence criteria in the following way,

$$\max_{i,j\in\Omega}\left(\left|u_{i,j}^{n+1}-u_{i,j}^{n}\right|\right)<\frac{\Delta t^{0}}{\Delta x} \tag{12}$$

where Ω is the image domain. The selection of the parameters Δt^{0}, β and γ is usually carried out based on experimentation.

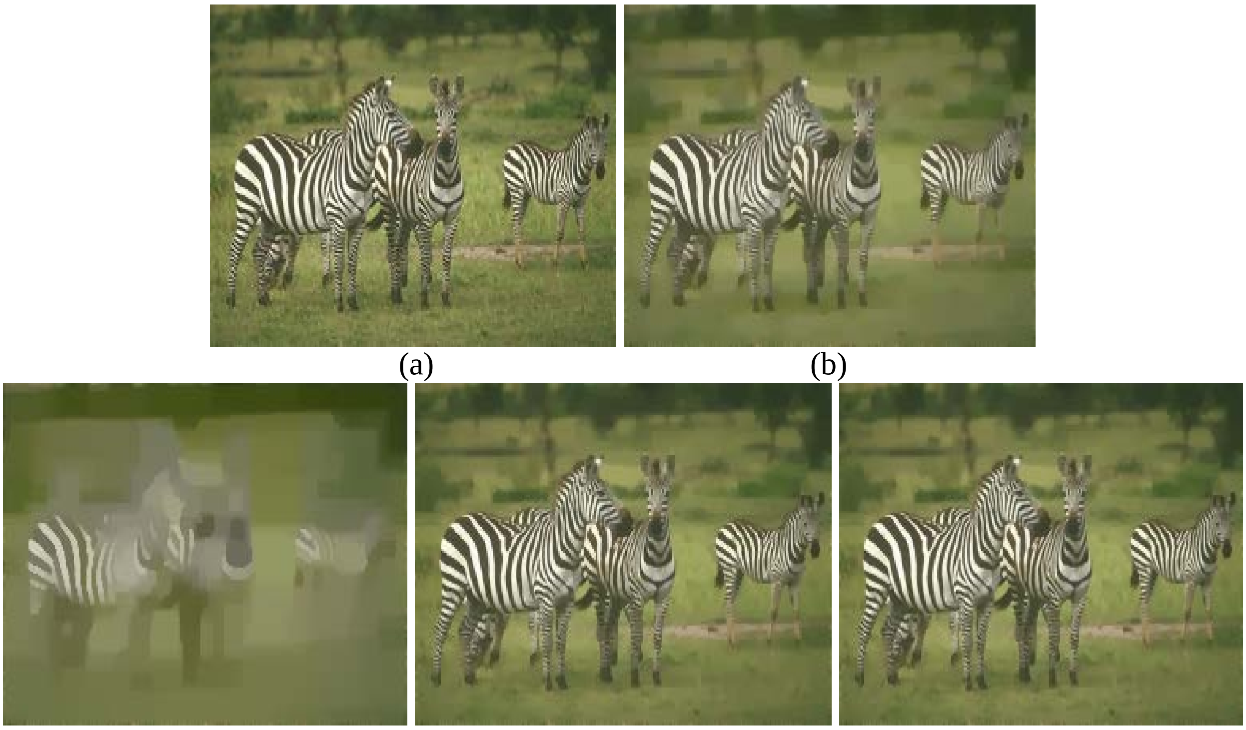

(a) (b)

(c) (d) (e)

Figure 3. Results of the TV flow when applied to the image depicted in Figure 2-a. (a-c) The results after the application of the S-TV flow procedure for 50, 300 and 1500 iterations respectively. (d,e) The results after the application of our TC-TV flow algorithm for 1500 iterations and for the number of iterations calculated using equation (12) respectively. In this experiment the TV flow parameters are set to the following values: Δt =0.05, γ = 0.995, β = 0.1.

For instance, if the parameter Δt^0 is set to a large value, the algorithm will converge faster but the results returned by the TV flow will show a significant amount of blur. Conversely, if the parameter Δt^0 is set to a low value such as 0.05 or lower then the algorithm will preserve better the features in the image, but this advantage is obtained at the expense of a high computational cost since the algorithm will require a large number of iterations to converge. Figure 3 illustrates the results obtained when the S-TV and TC-TV flow algorithms were applied to the image depicted in Figure 2-a and the time size parameter (Δt^0) is set to 0.05. From the experimental results depicted in Figures 2 and 3 it can be observed that the selection of the time size parameter plays an important role for the S-TV algorithm while the results returned by the TC-TV flow are virtually unaffected by the selection of this parameter (see the results depicted in Figures 2-d and 3-d). Figure 3 also illustrates the results obtained after the TC-TV flow is applied for 1500 iterations and for the number of iterations calculated automatically using equation (12). It can be observed that the results depicted in Figures 3-d and 3-e show similar levels of smoothing and this is motivated by the fact that due to the monotonic descent of the ageing function illustrated in equation (10) the smoothing effect has a negligible effect after the algorithm reaches the number of iterations calculated using equation (12).

The regularization parameter β has a lower influence on the performance of the TV flow and is typically set to a low value in the interval [0.05, 0.2]. The parameter γ controls the descent of the time size parameter and to avoid a step decrease of Δt that will result in a premature termination of the TC-TV flow it has to be set to a positive value close to 1 ($\gamma \leq 1$). Figure 4 shows a number of results after the application of the TC-TV flow to the image depicted in Figure 2-a when the ageing parameter γ is varied. The results depicted in Figure 4 indicate that a selection of the ageing parameter γ close to 1 (see Figure 4-c) would generate a smoothing framework that is similar with the S-TV flow (the monotonic descent of the ageing function is very slow). On the other hand, a selection of the ageing parameter γ to a value lower than 0.99

would force the TC-TV flow algorithm to a premature termination as illustrated in Figure 4-a. In our experiments we have observed that the optimal values for the ageing parameter γ are in the interval [0.991, 0.999].

Figure 4. Performance of the TC-TV flow when the ageing parameter γ is varied. (a) γ = 0.99. (b) γ = 0.995. (c) γ = 0.9995. To isolate the effect of the ageing parameter γ on the performance of the TC-TV flow in this experiment the time size and regularization parameters were maintained constant and set to the following values: Δt^0 =0.1, β = 0.1.

4. Experiments and Results

In this study we focus our investigation on the adaptive smoothing of colour images. To achieve this we have extended the smoothing model depicted in equations (6) and (7) to cover data where each datapoint is formed by a three-dimensional (3D) vector whose elements are defined by the colour components. One of the main issues associated with the formulation detailed in Section 2 is the poor convergence of the standard TV flow algorithms and to compensate for this problem we have implemented a cooling procedure that acts upon the time size parameter to force the amplitude of oscillations to follow a monotonic descent with the increase in the number of iterations. As indicated in Section 3 the introduction of this ageing procedure allows us to identify numerically the number of iterations (see equation (12)) and in our experiments we evaluate the effectiveness of this procedure. Figures 5 to 7 depict a number of experimental results when the standard S-TV and our TC-TV flow implementations were applied both to noiseless data and to images that are artificially corrupted with noise (Gaussian noise with a standard deviation of 20 gray-levels on each colour channel). In the experimental results shown in Figures 5 to 8 the number of iteration for TC-TV flow is determined automatically using equation (12).

The results depicted in Figures 5 and 6 indicate that the standard S-TV flow introduces an undesired level of blur while our TC-TV flow algorithm is able to achieve superior results. This undesired blur is obvious in image regions defined by the cart's wheel (see Figures 5-f (S-TV flow) and 5-g (TC-TV flow)) and in the image areas representing the fur in Figure 6 (see the close up details shown in Figures 6-f (S-TV flow) and 6-g (TC-TV flow)). The improved performance of the TC-TV flow algorithm when compared with the standard algorithm was expected since the strength of the smoothing process is decreasing with the increase in the number of iterations. We have also found in our experiments that the TV flow algorithms produce robust results even in cases when they are applied to images corrupted with a high level of image noise. Figure 7 depicts the results returned by the standard and TC-TV flow implementations when applied to an image corrupted with Gaussian noise with standard deviation of 20 gray-levels on each colour channel. Another set of experiments was conducted to evaluate the computational overhead associated with the TV flow algorithms analysed in this study and the results are depicted in Table 1. While the computational time required to compute the standard S-TV flow for one iteration is constant, the computational load when the algorithm is executed for 250 and 1000 iterations can be obtained by multiplying the computational time required to execute the algorithm for 500 iteration with a factor of 0.5 and 2 respectively. The experiments have been conducted using a Pentium 1.7 GHz PC, 500 MB RAM memory and running Windows 2000.

Table 1. Computational overhead associated with the standard (S-TV) and proposed TC-TV flow algorithms.

Image	Size	S-TV flow (500 iteration) (sec)	TC-TV flow (sec)
Figure 2-a	256 ×171	51.58	39.53
Figure 6-a	256 ×256	80.30	62.68
Figure 7-a	256 ×171	51.80	38.84
Figure 8-a1	256 ×171	52.18	34.43
Figure 8-a2	256 ×171	51.74	39.14
Figure 8-a3	257 ×269	82.74	62.50
Figure 8-a4	256 ×171	51.80	40.45

Figure 8 presents some additional results that illustrate the performance of our TC-TV flow algorithm when compared with that offered by the mean shift filtering (Comaniciu and Meer, 2002). These experiments were conducted on a set of standard test images from Berkeley database (BSDB, 2001). In all experiments the TC-TV flow parameters Δt^0, β and γ were set to 0.1, 0.1 and 0.995 respectively. The parameters that select the space (σ_s) and the range (σ_r) resolutions of the mean shift filtering were set to 8 and 4 respectively as recommended in Comaniciu and Meer, 2002. The experimental results depicted in Figure 8 indicate that the TC-TV flow algorithm outperforms the mean shift filtering when the results are analysed with respect to noise reduction and feature preservation.

(a) (b) (c)

(d) (e) (f) (g)

Figure 5. TV flow results. (a) Original image (BSDB, 2001). (b-d) Results returned by the S-TV flow algorithm when executed for 250, 500 and 1000 iterations respectively. (e) Result returned by the TC-TV flow algorithm (the number of iterations selected automatically using equation 12). (f, g) Close-up details from (b) S-TV flow (250 iterations) and (e) TC-TV flow respectively.

Figure 6. TV flow results. (a) Original image. (b-d) Results returned by the S-TV flow algorithm when executed for 250, 500 and 1000 iterations respectively. (e) Result returned by the TC-TV flow algorithm. (f, g) Close-up details from (b) S-TV flow (250 iterations) and (e) TC-TV flow respectively.

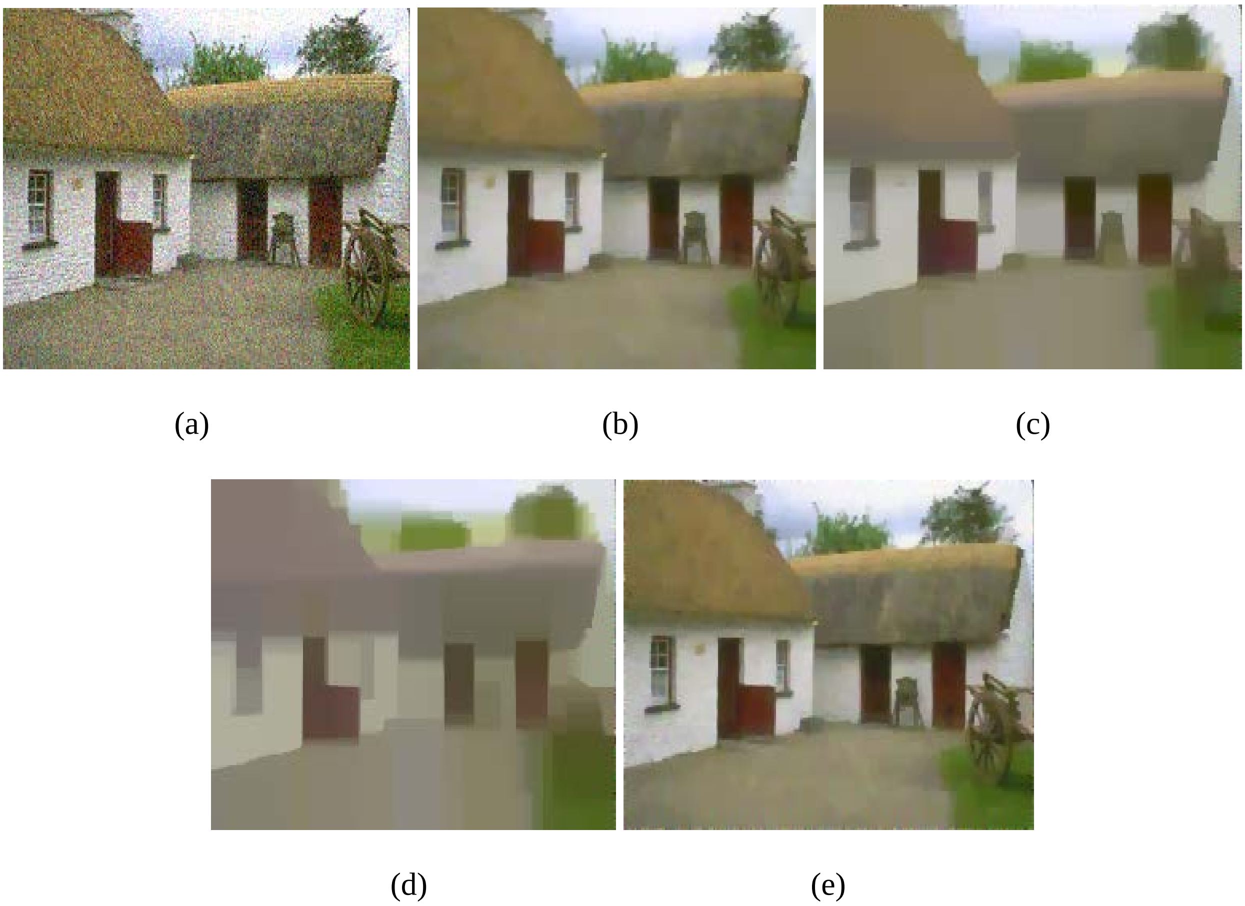

Figure 7. TV flow results. (a) Natural image corrupted with Gaussian noise (standard deviation 20 gray-levels on each colour channel). (b-d) Results returned by the S-TV flow algorithm when executed for 250, 500 and 1000 iterations respectively. (e) Result returned by the TC-TV flow algorithm.

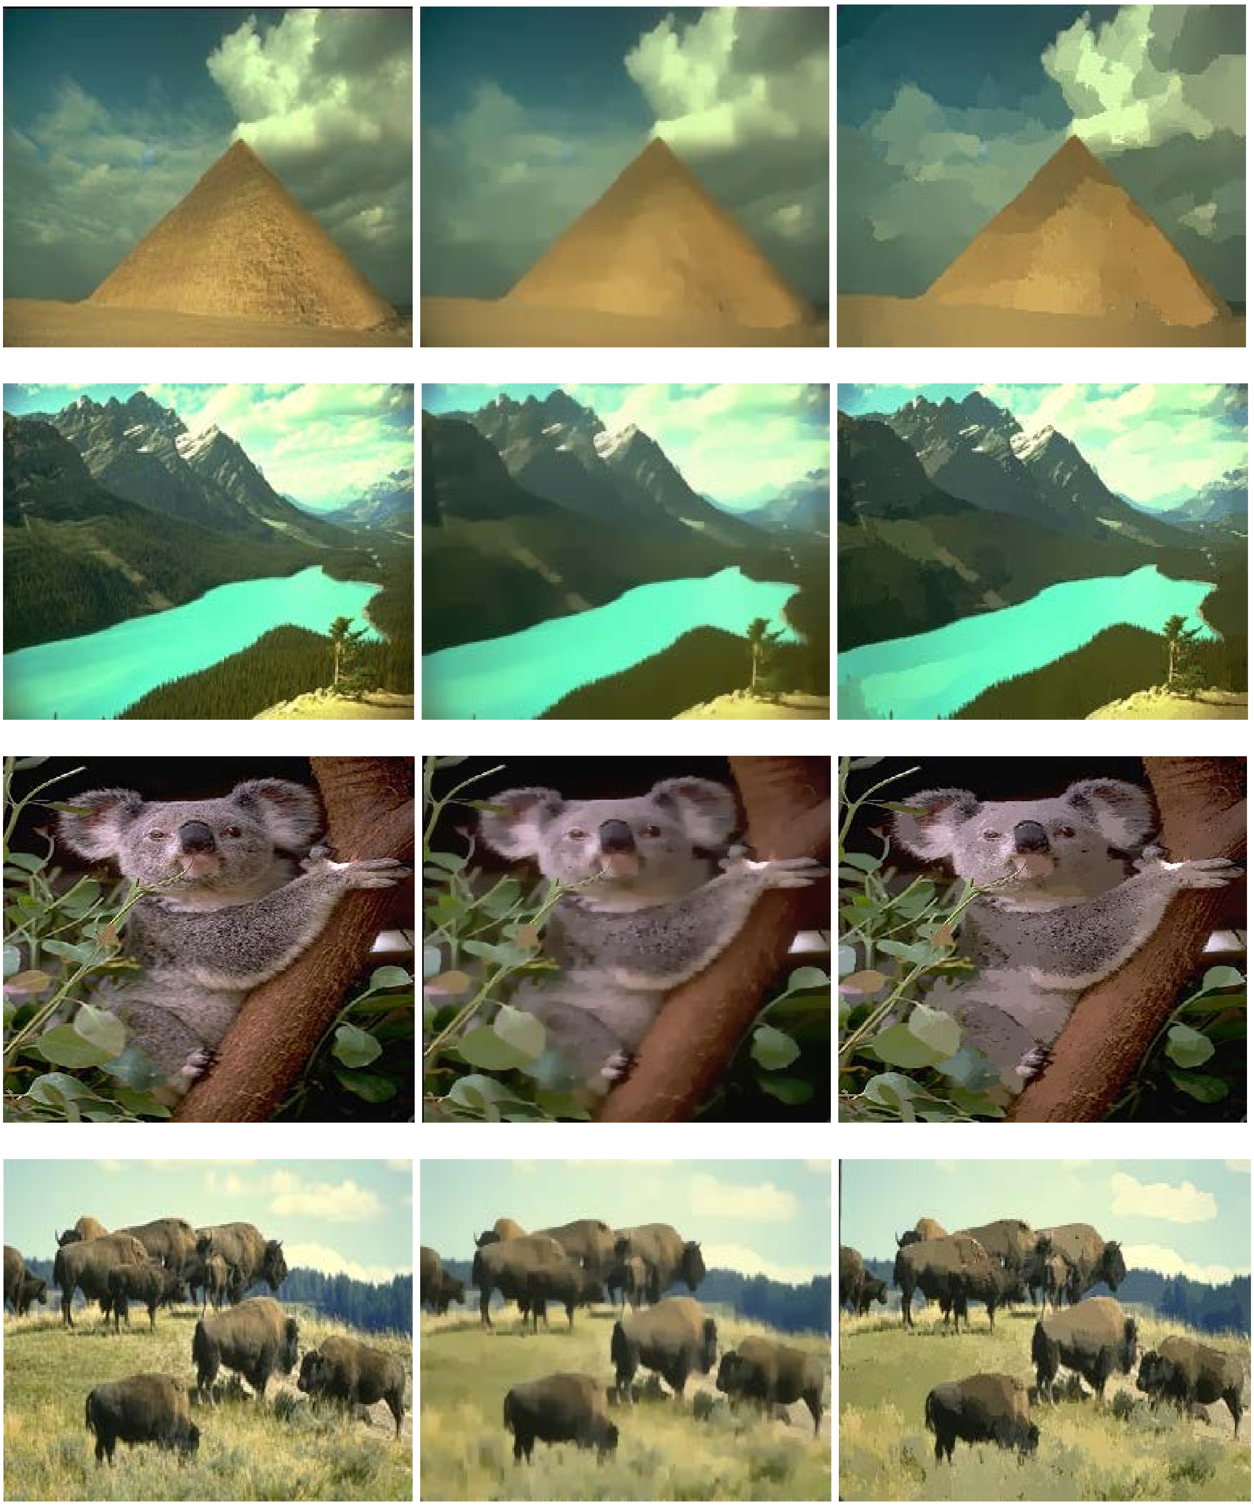

(a) (b) (c)

Figure 8. Additional results. Column (a) - Original images (BSDB, 2001). Column (b) - Results returned by the TC-TV flow algorithm. Column (c) Results returned by the mean shift filtering algorithm $(\sigma_s, \sigma_r) = (8,4)$.

5. Conclusions

The aim of this paper was to address some numerical aspects related to the stability of the discrete TV flow implementations. In this sense, we briefly detailed the methodology to implement the TV flow algorithms in the discrete domain and we have devised a number of improvements that increased the overall stability of the TV flow scheme. As indicated in other papers that have addressed the numerical stability of the TV flow algorithms (Andreu et al., 2001; Breuβ et al., 2006; Grahs et al., 2002), the most difficult problems associated with the implementation in the discrete domain is to control the level of oscillations and to select the optimal set of parameters. In this paper we demonstrated that the application of an ageing procedure that adjusts the value of the time size parameter with the increase in the number of iterations, leads to an efficient computational framework where the amplitude of oscillations decrease monotonically. We have also demonstrated that the application of the ageing procedure allows the development of a numerical solution to identify the optimal number of iterations. The experimental data indicated that the TV flow formulations are robust pre-processing schemes that can be successfully included in the development of a large number of algorithms ranging from adaptive data compression to image segmentation.

References

Andreu, F., Ballester, C., Caselles, V., Manzon, J.M., 2001. The Dirichlet problem for the total variation flow. Journal of Functional Analysis, 180, 347-403.

Breuß, M., Brox, T., Bürgel, A., Sonar, T., Weickert, J., 2006. Numerical aspects of TV Flow. Numerical Algorithms, 41(1), 79-101.

Comaniciu, D., Meer P., 2002. Mean shift: A robust approach toward feature space analysis, IEEE Trans. on Pattern Analysis and Machine Intelligence, 24(5), 603-619.

Dibos, F., Koepfler, G., 1999. Global total variation minimisation. SIAM Journal on Numerical Analysis, 37(2), 646-664.

Ghita, O., Robinson, K., Lynch, M., Whelan, P.F., 2005. MRI diffusion-based filtering: A note on performance characterisation. Computerized Medical Imaging and Graphics, 29(4), 267-277.

Gothandaraman, A., Whitaker, R.T., Gregor, J., 2001. Total variation for the removal of blocking effects in DCT based encodings. Proc. of the IEEE International Conference on Image Processing, Thessaloniki, Greece.

Grahs, T., Meister, A., Sonar, T., 2002. Image processing for numerical approximations of conservation laws: Nonlinear anisotropic artificial dissipation. SIAM Journal on Scientific Computing, 23(5), 1439-1455.

Ilea, D.E., Whelan, P.F., 2007. Adaptive pre-filtering techniques for colour image analysis. Proc. of the International Machine Vision & Image Processing Conference, IEEE Computer Society Press.

Keeling, S.L., Stollberger, R., 2002. Nonlinear anisotropic diffusion filtering for multiscale edge enhancement. Inverse Problems, 18, 175-190.

Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. IEEE Trans. on Pattern Analysis and Machine Intelligence, 12(7), 629-639.

Petrovic, A., Escoda, O.D., Vandergheynst, P., 2004. Multiresolution segmentation of natural images: From linear to non-linear scale-space representations. IEEE Trans. on Image Processing, 13(8), 1104-1114.

Rudin, L.I., Osher, S., Fatemi, E., 1992. Nonlinear total variation based noise removal algorithms. Physica D, 60, 259-268.

Smolka, B., Plataniotis, K.N., 2002. On the coupled forward and backward anisotropic diffusion scheme for colour image enhancement, Lecture Notes in Computer Science, Springer, 2383, 175-200.

Sonka, M., Hlavac, V., Boyle R., 1998. Image processing, analysis and machine vision. 2nd edition, PWS Boston.

Strong, D.M., Chan, T.F., 2003. Edge-preserving and scale-dependent properties of total variation regularization. Inverse Problems, 19, 165–187.

Weickert, J., ter Haar Romeny, B.M., Viergever, M.A., 1998. Efficient and reliable schemes for nonlinear diffusion filtering. IEEE Trans. on Image Processing, 7(3), 398-410.

Weickert, J., 1998. Anisotropic diffusion in image processing. Teubner Verlag, Stuttgart.

The Berkeley Segmentation Dataset and Benchmark (BSDB), 2001.

http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/