Pattern Recognition 42 (2009) 2306 – 2316

Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Edge-preserving smoothing using a similarity measure in adaptive geodesic neighbourhoods

Jacopo Grazzini*, Pierre Soille

Spatial Data Infrastructures Unit, Institute for Environment and Sustainability, Joint Research Centre—European Commission, TP 262 - via E.Fermi, 2749-21027 Ispra (VA), Italy

ARTICLE INFO

Article history:
Received 26 June 2008
Received in revised form 28 October 2008
Accepted 3 November 2008

Keywords:
Edge-preserving smoothing mathematical morphology
Spatial–tonal filtering
Multichannel image
Adaptive neighbourhood
Local pairwise similarity
Geodesic time
Geodesic mask

ABSTRACT

This paper introduces a novel image-dependent filtering approach derived from concepts known in mathematical morphology and aiming at edge-preserving smoothing of natural images. Like other adaptive methods, it assumes that the neighbourhood of a pixel contains the essential information required for the estimation of local features in the image. The proposed strategy essentially consists in a weighted averaging combining both spatial and tonal information. For that purpose, a twofold similarity measure is calculated from local geodesic time functions. This way, no prior operator definition is required since a weighting neighbourhood and a weighting kernel are determined automatically from the unfiltered input data for each pixel location. By designing relevant geodesic masks, two adaptive filtering algorithms are derived that are particularly efficient at smoothing heterogeneous areas while preserving relevant structures in greyscale and multichannel images.

1. Introduction

As a special case of filtering, image smoothing is commonly applied in a processing chain to improve the visual appearance in an image and to simplify subsequent image processing stages such as feature extraction, image segmentation or motion estimation [1,2]. The problem of image smoothing is to reduce undesirable distortions, due to the presence of noise or the poor image acquisition process, and that negatively affects image analysis and interpretation processes, while preserving important features such as homogeneous regions, discontinuities, edges and textures [3–5]. In this context, edge-preserving smoothing (EPS) techniques have been extensively used in computer vision and image processing in order to improve the performance of higher level processing stages [3,4,6]. There have been in particular substantial efforts in developing adaptive operators where the filter parameters can vary over different regions of the processed image. Indeed, such operators are desirable when both the image and the inherent noise have non-stationary behaviour. They can adapt to local image variation, so that they intrinsically allow the processing of image pixels with different strategies depending on the region they are positioned.

Strong relations have been established between a number of widely used adaptive EPS filters for digital image processing [4,7,8]. A generic idea underlying most of them is to update a pixel's intensity through a local weighted averaging of its neighbour pixels' intensities within a reduced processing window. Two important issues when designing this kind of filters are the selection of the processing window and the determination of the proper weights for averaging in accordance with the local features found in the image such as discontinuities or noise. In this paper, we propose an EPS approach derived from concepts known in mathematical morphology (MM) that does not require the definition of any processing window or any weights as it determines them automatically and adaptively from the input image. For that purpose, a pairwise discrete *geodesic time function* computed over an appropriate *geodesic mask* provides an adaptive neighbourhood and a local measure of the twofold spatial and tonal similarity around every pixel, so that the smoothing operation at a given pixel depends not only on the spatial location of its neighbour pixels but also on their tonal distance to it. Based on this approach, two efficient image-dependent algorithms are derived that are able to exploit different radiometric, geometrical and morphological image characteristics for EPS filtering. These algorithms are particularly suited to enhance the visual information in discrete images (e.g. remote sensed and medical data) while avoiding the creation of spurious artifacts through diffusion-like processes.

The rest of the paper is organised in the following manner. Section 2 presents a review of related works on adaptive EPS filtering and establishes a link with the EPS filter based on similarity

* Corresponding author.
E-mail addresses: Jacopo.Grazzini@jrc.it (J. Grazzini), Pierre.Soille@jrc.it (P. Soille).

0031-3203/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.patcog.2008.11.004

measures in adaptive geodesic neighbourhoods. In Section 3, we recall the fundamental notions of MM geodesic path and geodesic time within the discrete framework. We present the design of appropriate discrete adaptive neighbourhoods accounting for the local image content. In Section 4, we introduce the new filtering methodology based on the estimation of local geodesic time and derive from it two EPS algorithms. An extension of the algorithms to handle multichannel images is presented in Section 5. In Section 6, experiments are led on natural images and results are discussed. The conclusion and a description of future foreseen developments are presented in Section 7.

2. Related literature

This section establishes a relationship between several relevant areas of EPS filtering and our approach. The idea of adaptive filtering itself is not new, and many different methods have been proposed over the years for EPS. Detailed overviews and evaluations can be found in [4,6,9,10]. The most common form of smoothing of an image $f : \mathbb{Z}^2 \to \mathbb{Z}$ is a low-pass filtering, which can be expressed as the discrete convolution operation in the spatial domain [1]:

$$[f \otimes K](\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{Z}^2} K(\mathbf{x},\mathbf{y}) f(\mathbf{y}), \tag{1}$$

where $\mathbf{x} = (x, y)$ is the spatial location, $f(\mathbf{x})$ is the local intensity in $\mathbf{x}$ and K is a kernel function (also called 'window') that is assumed to be normalised: $\sum_{\mathbf{y}\in\mathbb{Z}^2} K(\mathbf{x},\mathbf{y}) = 1$, and of finite size as it is often desirable to estimate the true $f(\mathbf{x})$ from a local neighbourhood [8]. Depending on the functional form of K, different weighted averaging can be performed. The most straightforward and fastest technique consists in performing linear smoothing [1] by applying a kernel K invariant (i.e. with fixed size and weights) over space. Such approach assumes stationarity: the image is regarded as consisting of many regions in which the signal is stationary and ergodic. It yields good results when all the pixels in the window come from the same 'population' as the central pixel, as it uniformly smooths out the image. Difficulties arise when the window overlaps a discontinuity: on the boundaries between two regions, it leads to significant blurring caused by averaging of distinct populations. The main drawback is that image stationarity is not guaranteed, even within small neighbourhoods. The use of an invariant kernel K is not suited for images featuring structural information on various scales and with different shapes. This undesirable effect can be reduced using adaptive smoothing techniques that tune themselves to the contextual details of the image [3,11,12].

Adaptive techniques can apply both linear and nonlinear algorithms but are mainly characterised by filters that adjust to the local features, e.g. noise or discontinuities, detected over every location on the image. Such filters are more robust than non-adaptive filters and, therefore, they are appropriate for EPS. They can be subdivided into two main operator classes [13]: the adaptive–weighted operators [14] and the spatially adaptive operators [15]. The operator of the first class typically involves a fixed-size sliding window with coefficients depending on the local image statistics. The adaptive neighbourhood of the second class surrounds the central pixel, but its shape and size depend on the local image content rather than being arbitrarily defined. Following Eq. (1) the adaptive strategy implies to locally varying the kernel over image regions: around a pixel $\mathbf{x}$ to be updated, one uses a kernel $K = K_{\mathbf{x}}$ with proper weights and size depending on the actual image variability in the neighbourhood of $\mathbf{x}$. But how should one find the proper domain and weights for the kernel?

The issue of finding the proper area for the kernel was addressed by the so-called sigma filters [16] and the structure-adaptive anisotropic filter [17]. Ideally, the kernels should be wide in directions of homogeneous intensity and narrow in directions with important structural edges. To this end, other smoothing algorithms adapt the local kernel using both the location of the nearby samples and their intensity values [3,6]. The *bilateral filter*, introduced in [5] as a generalisation of the Gaussian convolution, exploits the idea of averaging pixel greylevels with weights depending on distances in the range and space domains. Namely, the kernel $K_{\mathbf{x}}$ takes into account two factors: spatial distances $d_s(\mathbf{x},\mathbf{y})$ and tonal distances $d_t(f(\mathbf{x}), f(\mathbf{y}))$. Introducing a tonal weight, the mixing of different intensity 'populations' is prevented. Such approach achieves both a strong denoising effect and an efficient sharpening of image. Strong connections have been established in the literature [4] between the bilateral filter and other nonlinear filtering strategies based on adaptive smoothing [3], M-estimators [6], Beltrami flow [18], local-mode finding [19] or mean-shift analysis [20]. In particular, the weighted averaging performed through Eq. (1) turns out to be an implementation of PDE-based diffusion [21], where the diffusion function is guided by local gradient strengths [7]. As a defining characteristic, iterative operations are often involved in these latter approaches, which, coupled with the fact that adaptive smoothing algorithms generally converge to a piecewise constant image [21], causes what we refer to as the termination problem, e.g. the difficulty to determine a suitable stopping time, at least in a general and automatic manner.

Following the line of thought of [5], we introduce a single-step adaptive approach that associates to every pixel of the input image a weighted convolution of sample points within an adaptive neighbourhood, where the weights depend not only on the spatial distances but also on the tonal distance to the considered pixel. The adaptive concept results from both the spatial adjustment of the operational kernel and the adjustment of the weights upon the kernel, as it estimates adaptively the neighbour sample points and the associated weights from the input data. Unlike the bilateral filter, it enables to account for the correlation between the location of the pixels and their values during the averaging process. Compared to other nonlinear techniques involving iterative operations [21], the proposed approach presents the advantage of not depending upon any termination time.

3. Defining adaptive geodesic neighbourhoods

Geodesic transforms are classical operators in image analysis [22–24]. We propose to use the geodesic time-based approach known in MM to build adaptive neighbourhood in the image graph.

3.1. Geodesic time on discrete graph

The *geodesic distance* between two points of a 2D connected set is defined as the length of the shortest (so-called *geodesic*) path(s) linking these points and remaining in the set [22,24]. The concept of *geodesic time* [25] generalises this idea for two pixels $\mathbf{x}$ and $\mathbf{y}$ on a greylevel image g (the *geodesic mask*): it is estimated as the smallest amount of time necessary for travelling from $\mathbf{x}$ to $\mathbf{y}$ over any path $\mathscr{P}$ lying on the hyperplane defined by the values of g. In this context, the image g is seen as a 'height map', e.g. a surface embedded in a 3D space, with the third coordinate given by its greylevel values.

As digital images are defined on discrete grids, the intensity values along a continuous path may not be known. This is why the geodesic time is related to cost functions in digital graphs and is practically computed on discrete paths [25,26]. Let us define a path $\mathscr{P} = \mathscr{P}_{\mathbf{x}\to\mathbf{y}}$, with length $l-1$ going from $\mathbf{x}$ to $\mathbf{y}$: $\mathscr{P}$ is a l-tuple $(\mathbf{x}_1, \ldots, \mathbf{x}_l)$ of pixels such that $\mathbf{x}_1 = \mathbf{x}$, $\mathbf{x}_l = \mathbf{y}$, and $(\mathbf{x}_{i-1}, \mathbf{x}_i)$ defines adjacent pixels for all $i \in [2, l]$. The time corresponds then to a cost function, denoted by Φ, for moving on the discrete greylevel mask g along $\mathscr{P}$. Each pair

$[\mathbf{x}_i, \mathbf{x}_{i+1}]$ of $\mathscr{P}$ is assigned a cost value $\tau_i(g) = \Phi(g, \mathbf{x}_i, \mathbf{x}_{i+1})$ defined as a function of the length $|\mathbf{x}_i - \mathbf{x}_{i+1}|$ and of the intensities $g(\mathbf{x}_i)$ and $g(\mathbf{x}_{i+1})$. Here, the length $|\mathbf{x}_i - \mathbf{x}_{i+1}|$ refers to the elementary step in the image graph: it is typically the Euclidean distance or the optimal Chamfer distance [27]. The total time for travelling over the image g along $\mathscr{P}$ is then equal to the sum of all the local costs.

$$\tau_g(\mathscr{P}) = \sum_{i=1}^{l-1} \tau_i(g). \tag{2}$$

Following, the geodesic time separating $\mathbf{x}$ and $\mathbf{y}$ is expressed by: $\tau_g(\mathbf{x}, \mathbf{y}) = \min\{\tau_g(\mathscr{P}) | \mathscr{P}$ is a path linking $\mathbf{x}$ to $\mathbf{y}\}$. This definition is easily extended for computing the geodesic time between a point $\mathbf{x}$ and a reference set Y of g: $\tau_g(\mathbf{x}, Y) = \min_{\mathbf{y} \in Y} \tau_g(\mathbf{x}, \mathbf{y})$ (with the special case $\tau_g(\mathbf{x}, Y) = 0$ when $\mathbf{x} \in Y$). By associating each point $\mathbf{x}$ of the mask g with its geodesic time to the reference set Y, the *geodesic time function* can finally be defined as

$$[\mathscr{T}_g(Y)](\mathbf{x}) = \tau_g(\mathbf{x}, Y). \tag{3}$$

There is a formal resemblance between this function and the continuous framework of the eikonal equation [28] since both involve the sweeping of the image by a propagating front. In particular, the reciprocal of the local speed of propagation in [28] is analogous to the cost function Φ underlying the calculation of $\mathscr{T}_g(Y)$.

3.2. Locally adaptive geodesic neighbourhoods

Using the previous definitions, a new family of adaptive neighbourhoods accounting for the local variability in an image f can be built. The key issue is the selection of a mask g and a cost function Φ that appropriately combine both spatial and tonal information in f.

Following the adaptive neighbourhood paradigm of [12], we derive two distinct classes of neighbourhoods that account for the local image content (Fig. 1):

- *Neighbourhoods based on the Σ-time*—Aiming at generalising the concept of *grey-weighted distance transform* [23,29], the time $\tau_g^{\Sigma}(\mathscr{P})$ necessary for travelling on a path is defined in [25] as the surface area under this path. By applying this definition locally and setting the mask to the image gradient: $g = \nabla f$, pixels separated by high gradient magnitude values are considered to be 'further' away than those separated by low values. Indeed, this way, the incremental cost $\tau_i^{\Sigma}(\nabla f)$ is implicitly assumed to be[1]

$$\tau_i^{\Sigma}(\nabla f) = \tfrac{1}{2}(|\nabla f(\mathbf{x}_i)| + |\nabla f(\mathbf{x}_{i+1})|) \cdot |\mathbf{x}_i - \mathbf{x}_{i+1}|, \tag{4}$$

  so that the time propagates through the lowest values of $|\nabla f|$: the lower it is, the faster the propagation.
- *Neighbourhoods based on the Δ-time*—In order to account for both the distance between points and the roughness of the surface, the authors of [30] rather use the *weighted distance on curves space* so that $\tau_g^{\Delta}(\mathscr{P})$ minimises the changes in greylevel values. With this approach and a mask set to the image itself: $g = f$, two pixels are close if there exists a path linking them along which intensity variations are low. The cost $\tau_i^{\Delta}(f)$ is given by

$$\tau_i^{\Delta}(f) = \tfrac{1}{2}|f(\mathbf{x}_i) - f(\mathbf{x}_{i+1})| \cdot |\mathbf{x}_i - \mathbf{x}_{i+1}|. \tag{5}$$

  The intuitive interpretation of this definition is that it represents the minimal amount of ascents and descents to be travelled to reach a neighbour pixel [30]. It also represents a measure of the shortest path drawn on the projection of the 2D image onto the spatial–tonal domain.[2]

These definitions can be connected to various MM techniques [13,31]. In particular, the Σ-neighbourhoods can be seen as an extension of the geodesic dilation of [25], obtained by thresholding the geodesic time function computed over greylevel images. The Δ-derived windows coincide with the so-called morphological amoebas of [32] where the neighbourhoods are calculated by introducing a distance defined over greylevel values only. It is also closely related to the minimal path of [33]. More generally, this latter approach can be related to the notions of connected components, flat and quasi-flat zones [34], where spatially adaptive neighbourhoods are built through specific criterion mapping (e.g. on intensity or local contrast). The geodesic approach leads to efficient implementation because classical shortest path algorithms can be applied such as the Dijkstra's graph search algorithm [23,26]. The computational complexity of the Σ-time $\mathscr{T}_g^{\Sigma}(\mathbf{x})$ estimated from a single pixel $\mathbf{x}$ can be reduced to $O(n \log n)$, where n is the number of pixels reached from $\mathbf{x}$ in the spatial domain, through the use of priority queue data structures [24,26]. Such structures take advantage of the fact that images are finite and guarantee that pixels are processed only once. The Δ-time computation is also essentially Dijkstra's algorithm, slightly modified to allow for multiple passes over pixels. Indeed, the incremental cost $\tau_i^{\Delta}(g)$ depends on the cost of the previous $\mathscr{P}^{\Delta}_{\mathbf{x}_1 \to \mathbf{x}_i}$, thus one cannot in general redefine g as the sum of fixed weights, as it is the case when implementing τ_i^{Σ}. The geodesic mask is in fact constantly updated through the propagation of the time τ^{Δ} [26,30]. Therefore, the implementation of $\mathscr{T}_g^{\Delta}(\mathbf{x})$ has a complexity of $O(kn \log n)$, the graph connectivity index k being the maximal number of visits of a pixel.

We can observe the way neighbourhoods adapt to the image context on the geodesic level-sets of Fig. 2. In particular, in the first column, the closest neighbourhood of the marker pixel is essentially composed of pixels belonging to the road, e.g. to the same elongated structure. In the second column, the pixels across the road do not belong to the closest neighbourhood of the marker as the cost of crossing the road is high when considering either the gradient or the image variations. For noisy and homogeneous regions (third and fourth columns), the neighbourhoods are rather isotropic. The geodesic times can pull pixels belonging to the same class closer and propel those belonging to different classes further away. It has the ability to find the neighbourhood with an arbitrary shape and can preserve the intrinsic structure of the image. This is a particularly desirable property for local adaptive filtering.

4. Designing geodesic filtering kernels

According to Section 2, the most common strategy encountered in EPS consists in building local adaptive kernel functions [1,4]. We propose to exploit the information carried out by the geodesic time to estimate a local similarity measure accounting for the image variability and use it define the kernel.

[1] Note that the Σ-time is not a distance: according to Eq. (4), it is possible for the cost τ_i^{Σ} to be null between two adjacent pixels. However, it can be remedied by using an addition instead of a multiplication in the estimation of the elementary costs. The same applies to τ_i^{Δ}, see Eq. (5).

[2] The geodesic time introduced herein in the discrete case can also easily be extended to the continuous framework. Namely, when considering a continuous path $\mathscr{P}$ and a mask $g : \mathbb{R}^2 \to \mathbb{R}$, the geodesic Σ- and Δ-functions are, *resp.* expressed as: $\tau_g^{\Sigma}(\mathscr{P}) = \int_{\mathscr{P}} |g(\mathbf{t})| \, \mathrm{d}\mathbf{t}$ and $\tau_g^{\Delta}(\mathscr{P}) = \int_{\mathscr{P}} |(\mathrm{d}g(\mathbf{t}))/\mathrm{d}\mathbf{t}| \, \mathrm{d}\mathbf{t}$, with $|\cdot|$ a norm on $\mathbb{R}$.

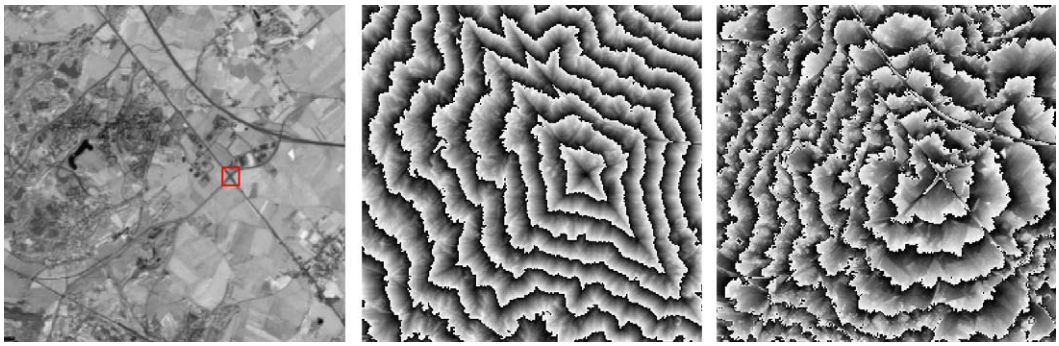

Fig. 1. Geodesic time propagation over a greylevel satellite image. Left: a marker pixel $\mathbf{x}$ is located at the intersection of two roads in the image (red mark); middle: $\mathscr{T}_g^{\Sigma}(\mathbf{x})$ estimated using the spatial gradient $g=\nabla f$; right: $\mathscr{T}_g^{\Delta}(\mathbf{x})$ computed over the image variations $g=f$. A cyclic greylevel palette has been used for visualising the geodesic fronts.

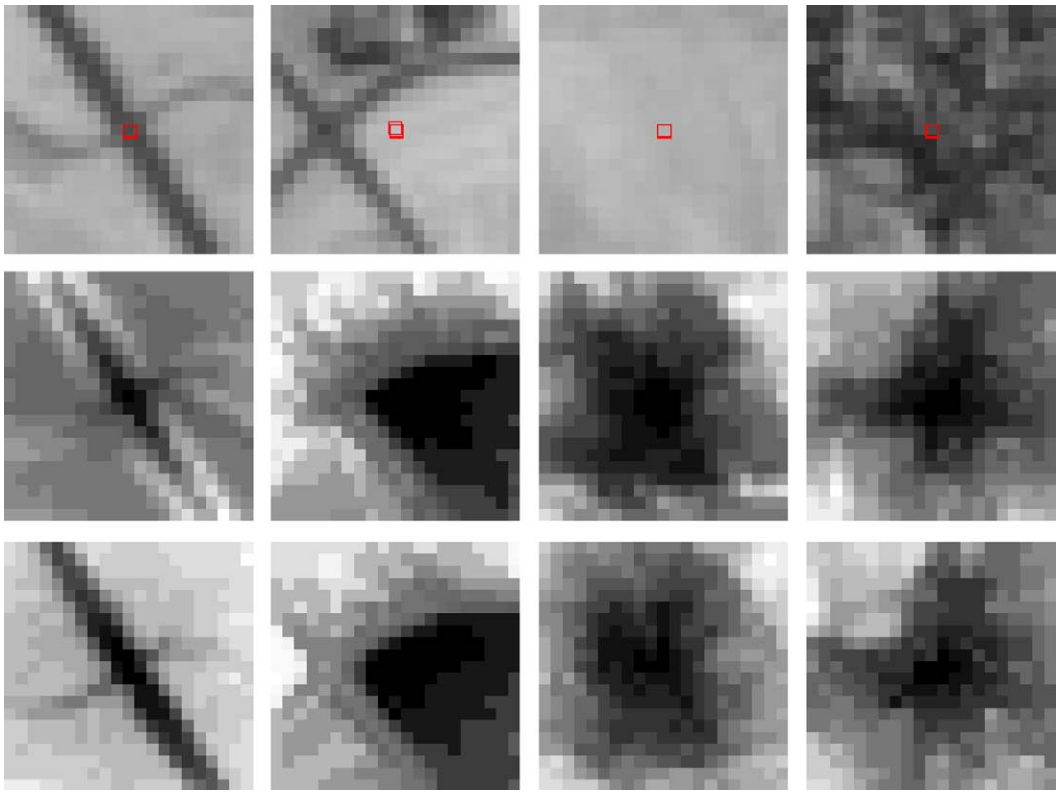

Fig. 2. Examples of local geodesic neighbourhoods. Top: excerpts (21×21 pixels) of the image of Fig. 1 with a marker pixel $\mathbf{x}$ located in the centre (red square); the level-sets of the time functions $\mathscr{T}_g^{\Sigma}(\mathbf{x})$ (middle) and $\mathscr{T}_g^{\Delta}(\mathbf{x})$ (right), *resp.* estimated using $g=\nabla f$ and f, are quantised in the range [0,15], from low (0) to high (15) time values. From left to right: $\mathbf{x}$ is, *resp.* located on a thin linear structure, near a strong discontinuity, inside a homogeneous textured area and in a noisy region.

4.1. Local pairwise similarity

In EPS, a similarity measure between the central pixel and its neighbourhood pixels is usually used to adjust the contribution of each pixel to the filtering kernel. Introducing such a measure is a good way to circumvent mixing different intensity populations, and, thus, critically determines the performance of the filter. A possible approach to measure similarity consists in using both the location of the nearby samples and their intensity values. We observe in particular that calculating the tonal weight d_t in bilateral filtering [5] implicitly introduces an estimate of the local gradient between neighbour pixels.

Similarly, we propose here a local pairwise similarity between a pixel $\mathbf{x}$ and any pixel $\mathbf{y}$ in its neighbourhood measured as a (positive monotonically) decreasing function of the geodesic time between them

$$K(\mathbf{x},\mathbf{y}) = \Psi([\mathscr{T}_g(\mathbf{x})](\mathbf{y})), \tag{6}$$

so that the shorter the time between $\mathbf{x}$ and $\mathbf{y}$, the stronger their similarity. The underlying idea is that the geodesic time function $\mathscr{T}_g(\mathbf{x})$ estimated at every pixel location $\mathbf{x}$ will define the intrinsic neighbourhood relationship(s) between $\mathbf{x}$ and its neighbours when the 2D image is projected onto the 3D (spatial–tonal) 'height map'. A large number of functions Ψ have been proposed in the literature for EPS techniques [6,8]. In bilateral filtering [5], the constituent weights are usually Gaussian functions, but others are not excluded [35].

4.2. Geodesic kernels

EPS is performed by applying the weighted average of Eq. (1) with the similarity measure of Eq. (6) as the filtering kernel. Moreover, a

parameter α that controls the relative influences of tone and space in the calculation of K is introduced: the cost of crossing pixels is refined to $\alpha \cdot \tau_i$. Finally, using a Gaussian function G_σ with standard deviation σ (considering the concurrence of the parameter α, we set $\sigma = 1$), we are able to define new adaptive kernels:

- *Σ-kernel over image gradient*—The values of $|\nabla f|$ are regarded as the cost of crossing a pixel in the so-called Σ-kernel

$$K^{\Sigma}(\mathbf{x},\mathbf{y}) = G_\sigma\left[\alpha \cdot \sum_{i=1}^{l-1} \tau_i^{\Sigma}(\nabla f)\right], \tag{7}$$

  where the sum is performed along the (shortest) geodesic path $\mathscr{P}^{\Sigma}_{\mathbf{x}\to\mathbf{y}}$ with length $l-1$ linking the central pixel $\mathbf{x}$ to a neighbour pixel $\mathbf{y}$. As a consequence, in the averaging, higher weights are assigned to the nearby sample pixels that involve low gradient values along $\mathscr{P}^{\Sigma}_{\mathbf{x}\to\mathbf{y}}$, as compared to samples that are either further away from $\mathbf{x}$ or separated by higher gradient values.
- *Δ-kernel based on image variations*—Higher weights are assigned to the nearby sample pixels which are linked to $\mathbf{x}$ and with similar greylevel values. Hence, along a geodesic path $\mathscr{P}^{\Delta}_{\mathbf{x}\to\mathbf{y}}$, the Δ-kernel is defined as

$$K^{\Delta}(\mathbf{x},\mathbf{y}) = G_\sigma\left[\alpha \cdot \sum_{i=1}^{l-1} \tau_i^{\Delta}(f)\right]. \tag{8}$$

Performing the geodesic filtering with either of these two kernels is in fact equivalent to operate a progressive one-dimensional filtering, as the averaging is performed along a geodesic path. The shortest paths defined with these algorithms are similarly constrained to the surface of the 'height map': the path between two close pixels can be long, if there is a high 'ridge' or deep 'valley' in the intensity or gradient map between them. Still, Eqs. (4) and (5) clearly define distinct geodesic paths, and thus distinct models (see Fig. 1).

Even if the spatial distance (d_s in the bilateral filter) does not appear explicitly in Eq. (6), it is implicitly taken into account by the geodesic time functions as the length of the geodesic paths. Its role is to limit the spatial extent of the filter. The respective tonal distance (equivalent to d_t) enables the suppression of the contributions of pixels belonging to different connected components (see Fig. 3). This way, pixels from across a sharp feature are given less weight because they are not connected to the central pixel through any geodesic path and, therefore, they are penalised by the geodesic functions. Typically, if a pixel is located near an edge, then the intensity values of pixels on the same side of the edge will have much stronger influence on the filtering.

5. Extension to multichannel images

The study so far has dealt with scalar (greylevel) images. We discuss here some special aspects of the algorithm when it comes to the processing of multichannel images. Of particular interest is the integration of the contrast information contained in the various channels into one meaningful result.

Existing multichannel image processing solutions can be discerned between marginal methods, which operate on each channel separately, and vector methods, which process the pixels as vectors [1,36]. In particular, in applying our algorithm to multichannel image, one could envision a marginal procedure where the kernels K^{Σ} and K^{Δ} are applied unaltered to the different channels separately. However, in order to preserve the inherent correlation that exists between the different channels, the vector approach is preferred in general. Its major advantage is that it takes into account the actual multispectral edge information, so that further processing will be more efficient to preserve edges.

Our approach, in both its theoretical foundation (comparison of distance) and its implementation (use of priority queues), is highly dependent on the definition of an order on the range space. Unfortunately, it is not possible to define uniquely the ordering of multivalued data. The previous operators need, therefore, to be adjusted when dealing with multichannel images. Let us consider such an image $f : \mathbb{Z}^2 \to \mathbb{Z}^M$, with channels $(f_1, \ldots, f_M), M > 1$, we then propose the following filtering strategy to take into account the actual multichannel information:

- *For the Σ-filter*—A multichannel gradient should be introduced for extending the spatial gradient of Eq. (7). A way to achieve this is given by the *first fundamental form* [37]: $\Gamma = (\sum_m(\partial f_m/\partial x)^2, \sum_m(\partial f_m/\partial x)(\partial f_m/\partial y); \sum_m(\partial f_m/\partial x)(\partial f_m/\partial y), \sum_m(\partial f_m/\partial y)^2)$,which is a local measure of directional contrast based upon the gradients of the M bands,[3] and thus reflects the multispectral edge information of f. Indeed, the direction of maximal and minimal changes are given by the eigenvectors of this quadratic form while the corresponding (positive) eigenvalues $\lambda^+ \geqslant \lambda^-$ denote the rate of contrast change. In particular, the largest eigenvalue λ^+ is known to be the derivative energy in the most prominent direction. For a greylevel image ($M = 1$), it is verified that $\lambda^+ = |\nabla f|^2$, while $\lambda^- = 0$ [37]. Taking into account these observations, λ^+ is a natural estimate for the norm of the image gradient in Eq. (4). This leads to a kernel filter identical to the one proposed by [39] in the continuous framework. Other measures can be considered, e.g. $\lambda^+ - \lambda^-$ which is similar to λ^+ corrected by the energy contributed by noise [40]. In practice, we need to compute the 2×2 matrix Γ for each image point $\mathbf{x}$ in order to apply Eq. (4).
- *For the Δ-filter*—A multivariate ordering criterion (e.g. distance measures or similarity measures) should be defined for considering the local variations of the multichannel image [36]. In particular, the norm in Eq. (5) must be understood as a multispectral norm. In practice, we use the L^∞ norm on the different channels; then we have, when comparing the geodesic time, $\tau_i^{\Delta}(f) \leqslant t$ if and only if $|f_m(\mathbf{x}_i) - f_m(\mathbf{x}_{i+1})| \leqslant t$ for all $m = 1, \ldots, M$. As a consequence, the Δ-filter algorithm depends on the dimension of the tonal space, which may increase its processing time.

6. Experiments and discussion

In this section, experiments are conducted on digital images with the proposed EPS filters. The results are discussed and compared with the outputs of some standard EPS filtering techniques encountered in the literature.

6.1. Limitations of the original approach and suggested improvements

A limitation of both geodesic filters regards the treatment of the (scalar or vector) intensity $f(\mathbf{x})$ of the input central pixel $\mathbf{x}$. As already underlined in [19], by computing a spatial average, we make the local smoothing very dependent on the correctness of this value. Using $f(\mathbf{x})$ as the 'reference' tonal data for the estimation of the local geodesic time assumes that this value is more or less noise free: this is naturally a questionable assumption when building a noise suppression filter and it may have effect on the result. Especially it is not applicable when impulse noise affects the image. Following the

[3] Note that this expression is equivalent to that of the unsmoothed structure tensor $\nabla f^{\mathrm{T}} \nabla f$ generalised from scalar data to vector data [38].

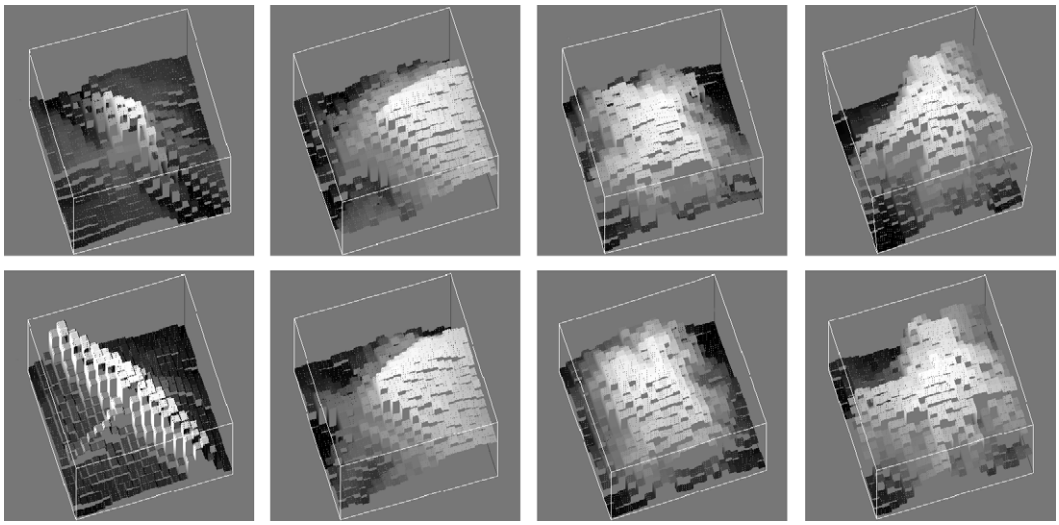

Fig. 3. Geodesic weighting kernels estimated on the excerpts of Fig. 2. The weights of the kernels K^{Σ} (top) and K^{Δ} (bottom) were quantised in [0,255], from (0) low to high (255) values, and visualised in 3D.

suggestion of [19], the filtered image can be improved through a prior processing step aiming at cancelling possible outliers present in the image, e.g. using a (possibly vector-valued) median filter [2,41] or a median filter followed by self-dual reconstruction [42]. It implicitly introduces an additional 'pilot image', similar to that of [32], which provides an initial estimate of the true value.

Problems can also occur with the Σ-filter when the signal-to-noise ratio of the image is low (and, as a consequence, its derived local gradient vectors are noisy). To remedy this, one can perform a spatial regularisation using 2D Gaussian convolution for removing noise prior to filtering. However, it may cause important discontinuities to be blurred. It is in fact preferable to use the 2×2 matrix derived from the smoothed structure tensor [38] instead of the first fundamental form (see Section 5) for estimating the eigenvalues of the gradient tensor. The experiments in the following were led using this latter improvement. In addition, other discontinuity measures can be used to remedy overlocality, a weakness of spatial gradient.

Other problems concern the integration of the multichannel information. Our approach assumes that the information in the different channels is correlated to some degree. However, the spectral information might be highly correlated in some parts of the image and uncorrelated in other parts. Moreover, it makes the intrinsic assumption of comparable significance of each channel: it supposes that the level of noise is similar in all channels. It could be, therefore, useful to assign different weights to the components according to their significance and their level of noise, but this issue is not discussed in this paper.

6.2. Proposed implementation

In practice, due to memory and computational limitations, the support of the kernel K is limited in size, e.g. sample pixels that are further away than a given distance, say ω, to the central pixel $\mathbf{x}$ are not considered. As a consequence (see Section 3.2), calculating the Σ-time from every pixel in the image results in a complexity of $O(N \cdot \omega^2 \log \omega^2)$ where N is the total number of pixels, while the running time of the Δ-filter amounts to $O(N \cdot k\omega^2 \log \omega^2)$. An alternative approach consists in limiting the filtering to the sample pixels reached from $\mathbf{x}$ with a time inferior to a given threshold value: this is simply achieved by checking if the time is above this threshold when one extracts an element from the priority queue, and by stopping, in such case, the propagation process.

In terms of processing time, the implementation of the geodesic filters also depends on the selected distance ω. Running the Σ-filter on a 2.4 GHz-CPU to denoise a 3-bands image of size 400×400 pixels (see Fig. 5) took between 2.9 s when $\omega = 3$ and 33 s when $\omega = 11$. When applied on the same images and with identical parameters, the Δ-filter took between 5.8 and 51 s. These results however, do not incorporate any optimisation part. In particular, based on 1-byte encoded scalar images, a finite table of pre-computed G_{σ}-values can be used to further reduce the computational time. Interestingly, the best filters are usually obtained for small windows (typical values of $\omega = 5, 7$, see Fig. 6), which ensures a reasonable processing time.

6.3. Results and evaluation

The performance of the Σ- and Δ-filters in enhancement is first qualitatively evaluated through the subjective inspection of the visual appearance of filtered images (Figs. 4 and 5). Quantitative results based on the PSNR estimation [43] are also presented in Fig. 6, where the geodesic filters were applied on images with different noise distributions. Both filters result in rather satisfying filtered versions of the original images, smoothing homogeneous areas while preserving important structures such as edges (Figs. 4 and 5). Indeed, the generic filtering approach enhances features through the combined spatial and tonal actions represented in the geodesic similarity measure of Eq. (6). While the role of the (implicit) spatial weight in the kernels is to limit the spatial extent of the filter, the effect of the tonal weight is to suppress the contributions of pixels from different intensity populations (see Section 4.2). Therefore, since only pixels sharing similar intensity with the current pixel have significant weights in the averaging, edges are not diffused across and noise is effectively suppressed. Close inspection to the images after a processing step also shows that they are good at enhancing subtle texture regions and they suppress small elements corresponding to the main heterogeneities. Namely, when considering the edge map obtained after running our filters (Fig. 4(c)), it appears that the main edges are not altered, while the number of noisy edges are reduced. Indeed, the main features are usually still present in the filtered images, even in the case of strong smoothing (Fig. 4(c), right). This is

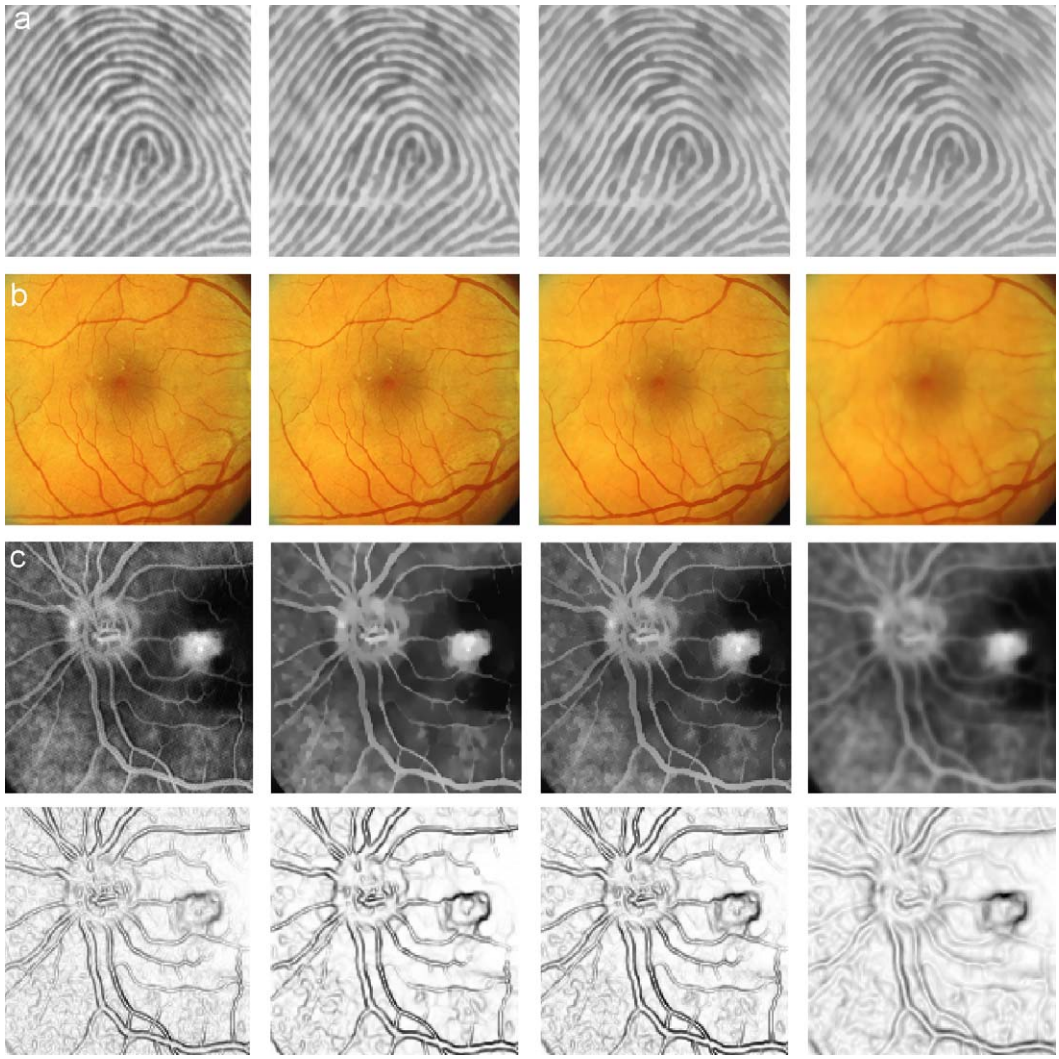

Fig. 4. Examples of filtering applied to greylevel and colour images by varying the parameters (α, ω) of the geodesic filters, as described in Sections 4 and 5 for colour images. In all experiments, the parameter $\sigma = 1$ of the Gaussian weighting function is constant. (a) Influence of the window parameter ω. From left to right: the Δ-filter is applied on a fingerprint image (left) with ω set to 3 (soft smoothing), 7 and 20 (strong); the parameter $\alpha = 2$ is constant. (b) Influence of the control parameter α. From left to right: the Σ-filter is applied on a colour retina image (left) by setting α to 7 (soft smoothing), 3 and 1 (strong); the parameter $\omega = 7$ is constant. (c) Highlighting the features in filtered images. Top, from left to right: smoothing of a greylevel retina image (left) using the Σ-filter with $(\alpha = 7, \omega = 11)$ and the Δ-filter with $(\alpha = 5, \omega = 7)$ and $(\alpha = 1, \omega = 3)$. Bottom: respective edge map computed as the magnitude of the Sobel gradient [1].

due to the fact that the filtering performed in this approach can be reduced to a one-dimensional process along geodesic paths, so that it is in particular able to preserve thin elongated structures. In general, the image structures are not geometrically damaged, which would be an issue for further processing like classification or segmentation. In this context, both the Σ-filter based on image gradient and the Δ-filter based on image variations show high capability at EPS when applied on noise-free images or on images perturbated with Gaussian distributed noise (Figs. 4–6). However, due to its intrinsic dependence on the intensity differences, the Δ-filter will rather enhance than smooth outlier pixels in the case of impulse noise.

The global strength of smoothing can be controlled by the parameter α, which amplifies or attenuates the influence of the local contrast in parts of an image (Fig. 4). Small α values lead to an increase of the amount of blurring so that details have been sacrificed to the effect of denoising, producing a visual effect similar to the well-known cartoon-like effect (Fig. 4). When α increases, the geodesic time on the embedded surface becomes more sensitive to image deformations: the cost of crossing pixels increases. For intermediate α values, both filters result in a less diffusive effect. With higher values of α, almost all contrasts are preserved and filtering has very little effect on the image. Fig. 6 shows the respective influences of the parameters α and ω in the denoising process. A good compromise is generally found for kernels with size, ω between 5 and 7 (defining a set of 25–50 pixels for estimating the central filtered pixel) and smoothing factors, α between 5 and 7 (weighting the respective contribution of domain and range). In the particular case of Gaussian and uniform noise (Fig. 6, first two columns), these latter α values seem to be optimal for the considered Σ-filter, and do not depend much on the kernel size. On the contrary, filters using lower α values result in averaging the noise and, therefore, altering the structures when ω increases, whereas filters with high α values, and thus little smoothing effect, perform better denoising when ω increases.

6.4. Comparison with other methods

In Fig. 7, we present the filtered images obtained by applying different EPS techniques known in the literature on a single noisy

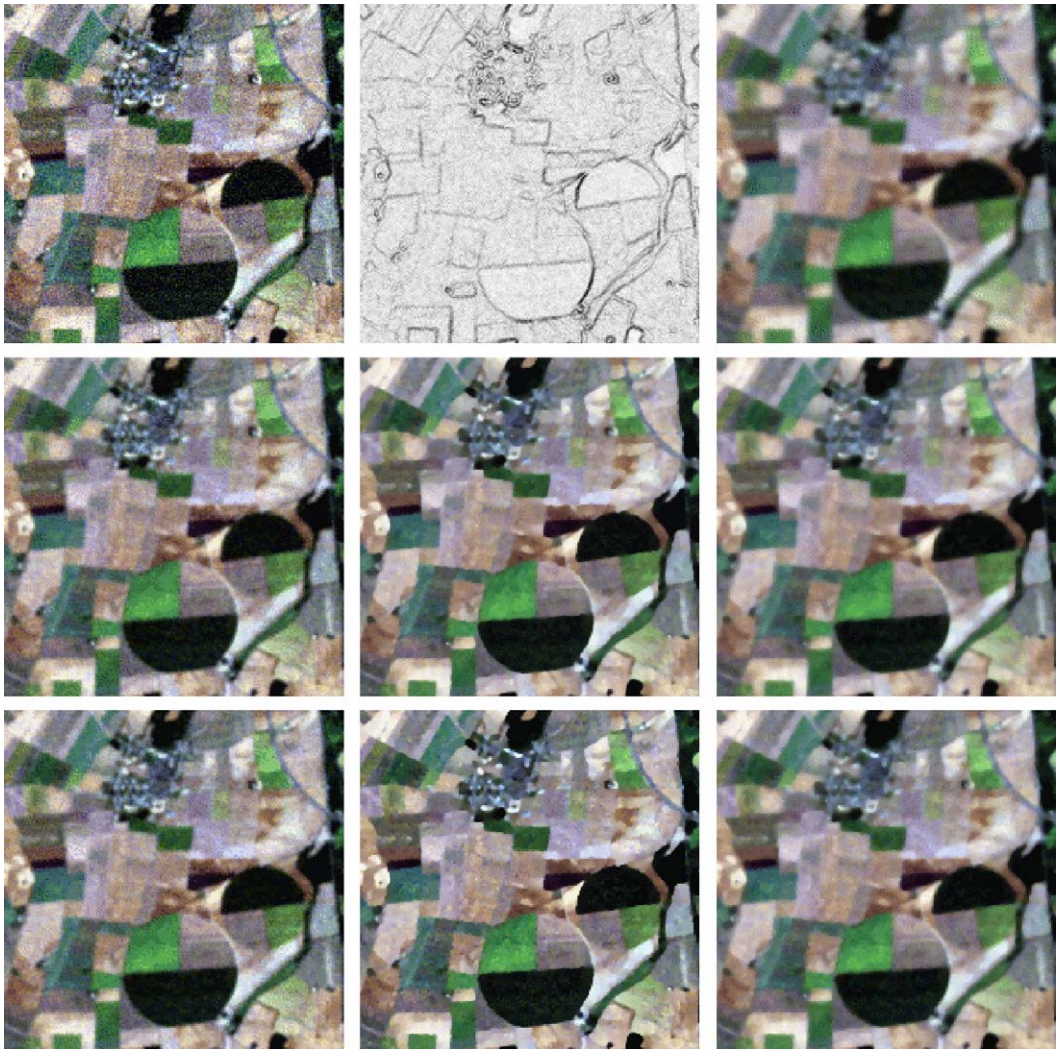

Fig. 5. Multispectral image denoising: geodesic filters compared to the amoeba filter of [32] when applied on a multichannel satellite image perturbated by a Gaussian distributed noise. First column: original noisy image (top) and images filtered using amoebas with radius 3 (middle) and 5 (bottom). Second column: λ^+ estimated from the structure tensor of the input image (top) and used by the Σ-filter with parameters $\alpha=3, \omega=3$ (middle) and $\alpha=5, \omega=5$ (bottom). Third column: Gaussian image (top) used as a 'pilot' for the Δ-filter with parameters $\alpha=3, \omega=3$ (middle) and $\alpha=5, \omega=5$ (bottom). The geodesic filters reach higher PSNR than the amoeba filter on this example.

image. For each method, we selected the outputs that were the most satisfying visually. The geodesic filters and the bilateral one [5] result in similar smoothed images. However, the latter uses, as a simple and intuitive choice for the adaptive kernel, separate terms for penalising the spatial and tonal distances. Breaking the filtering kernel into spatial and tonal terms (d_s and d_t) weakens the estimator performance since it limits the degree of freedom and ignores correlations between the location of the pixels and their values. In particular, pixels contributing to the filter are spatially unrelated to each other. By contrast, the proposed twofold similarity measure based on the geodesic time over an image-derived manifold enables to account for the correlations between the pixels' location and their values. Moreover, the (Euclidean) distance used in the bilateral filter [5], while being easier to calculate, does not take into account the image intensity values between two image pixels and thus ignores connectivity [18]: a pixel can have a relatively high weight although it belongs to a different object than that of the central pixel. The geodesic filters enable to penalise pixels that belong to a different connected component, as filtering is performed along geodesic paths. One advantage of the bilateral filter over our approach is that its computational complexity can be reduced to $O(K \cdot N \log N)$ with K independent of the window size ω [35].

The anisotropic diffusion [21] is an iterative technique that depends upon a termination time. Its major drawbacks include the fact that the solution has to be found using an often time-consuming iterative procedure and that it is very difficult to find a suitable stopping time. Our approach is a non-iterative estimation technique, which makes it more efficient and more stable. The adaptive Gaussian filtering of [44] adjusts locally the smoothing scale in a scale-space framework, through a minimal description length criterion, and it is not iterative. However, it also results in more blurred images (Fig. 7). The mean-shift algorithm [20] operates only on image intensities (be they scalar or vector-valued) and does not account for neighbourhood structure in image. Moreover, it is not fast because it requires many iterations to achieve the desired output. By adopting an adaptive strategy, our approach addresses the lack of flexibility of morphological operators like area open-closing, top-hat, based on the difficult choice of a structuring element [31]. Other morphological EPS filters based on self-dual reconstruction [42] are also able to smooth out texture and noise while preserving edge and corners (Fig. 7), but, while they consider information regarding the tonal distance and the connectivity, they do not integrate the spatial distance into the reconstruction process so that they usually flatten considerably the image. Note, finally, the difference between the Δ-filter and

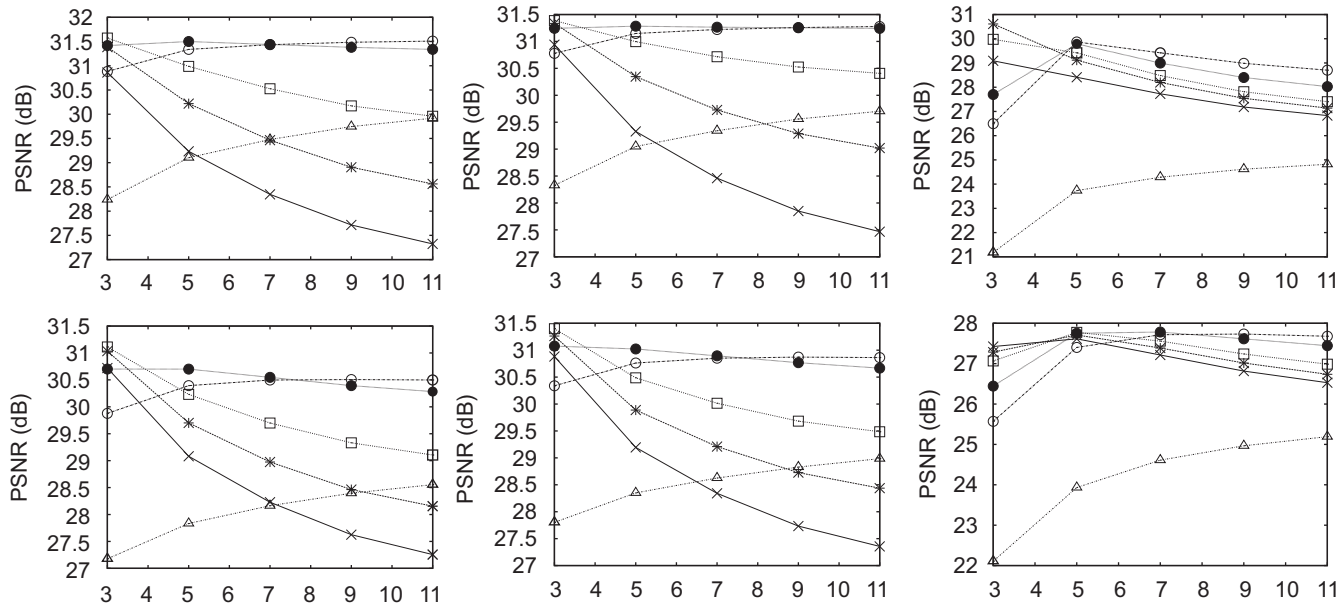

Fig. 6. The geodesic filters are applied on the greylevel image of Fig. 1, *resp.* perturbated with Gaussian noise (first column, with $PSNR = 24.14$ dB), uniform noise (second column, 24.93 dB) and impulse noise (third column, 15.82 dB). On each graph, the PSNR obtained with either the Σ-filter (top) or the Δ-filter (bottom) for $\alpha = 1$ (represented with symbol $\times$), 2 ($*$), 3 ($\square$), 5 ($\bullet$), 7 ($\circ$) and 15 ($\triangle$) are presented in function of the window size ω, varying from 3 to 11.

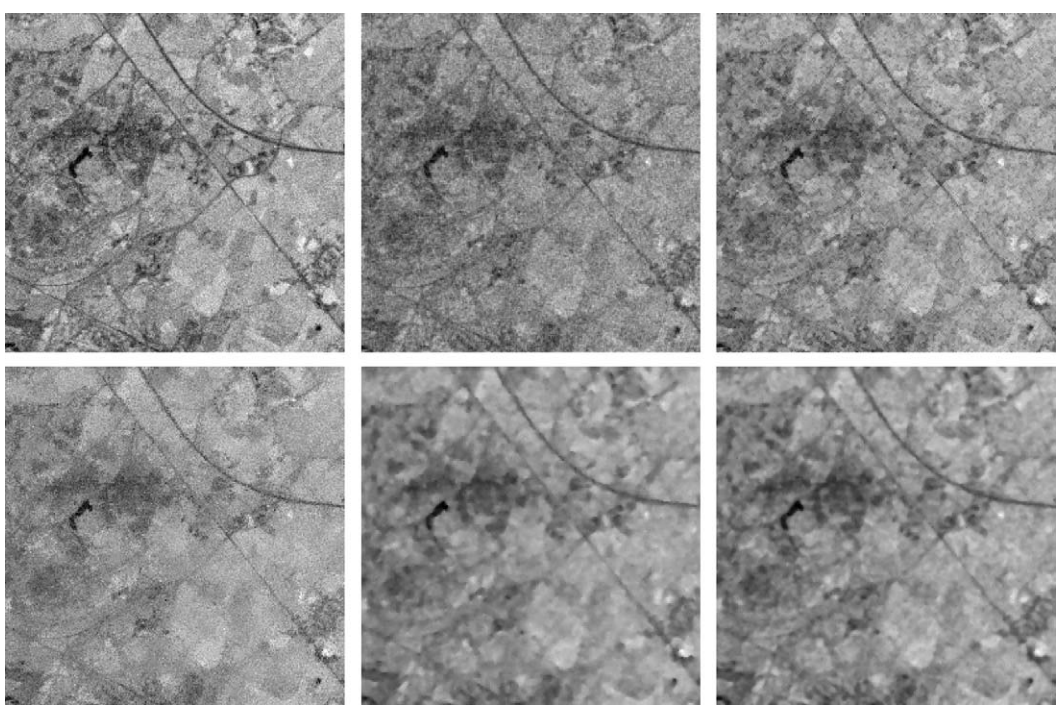

Fig. 7. Image denoising: the results of smoothing the satellite greylevel image of Fig. 1, corrupted with Gaussian distributed noise ($PSNR = 24.33$ dB, top left), are displayed (in this order) for: the Gaussian filter of [44] (27.89 dB), the morphological self-dual reconstruction from the median filter of [42] (29.74 dB), the bilateral filter of [5] (28.45 dB), and the geodesic Σ-filter (31.5 dB) and Δ-filter (31.11 dB).

the amoeba filters of [32]: not only the geodesic neighbourhoods are considered in the approach, but also the geodesic values themselves, as they are used to define the weights of the samples in the kernel, while the amoebas consider mean or rank statistics instead (Fig. 5).

7. Conclusion

In this paper, we introduced a new heuristic methodology for adaptive EPS from which we derive two efficient filtering algorithms. The basic idea is similar to that of spatial–tonal filtering approaches, which consist in employing both geometric and intensity closeness of neighbour pixels. The originality of our approach lies in the definition of a new similarity measure combining both spatial and tonal information and based on the local estimation of geodesic time functions. It performs around each pixel a weighted convolution, making use of: (i) a new set of weighting neighbourhoods through an adaptive selection of the pixels with non-negligible contribution to the filter, (ii) a new weighting function through an automatic estimation of the weights according to the local geodesic pattern. By designing relevant geodesic masks, we can define new smoothing filters that

enable the simplification and/or the denoising of images, depending on the input data and on the target purpose. The experiments performed herein show the effectiveness of the proposed approach in comparison to some standard EPS techniques. It can potentially preserve important structural elements, such as multichannel edges, and eliminate degradations. Like other spatial–tonal based techniques, the degree of smoothing in the image can be furthermore controlled to adjust the fidelity to the original image. By blurring small discontinuities and sharpening edges, the image structures are not geometrically damaged, which might be fatal for further processing. Therefore, this approach can be used as a preprocessing stage in feature extraction and/or image classification as it enables to create homogeneous regions, instead of pixels, as carriers of features. It is of particular interest for filtering data for which a discrete approach should be adopted, instead of a continuous one, in order to avoid creating spurious artifacts through diffusion-like processes. Potential applications are foreseen in the fields of medical imaging and remote sensing.

Possible improvements concern mainly the selection of the different parameters involved in the filtering strategy. One issue regards the spatial extent of the window used for estimating the local geodesic time functions: herein, we used a finite spatial window with size ω for limiting the calculations. Another issue concerns directly the estimation of the local geodesic time: the calculation could be refined by normalising locally the local cost τ_i of pixel crossing when propagating the geodesics. The selection of the smoothing control parameter α should also be further investigated. Further experiments should be led on strongly textured images and confronted to other recently developed techniques, in particular those that use a non-local approach—with neighbourhoods not necessarily spatially connected to the central pixel—for filtering [8]. Finally, it should also be compared to classical solutions to the differential wave-front propagation equations, especially to fast marching methods [28].

Acknowledgements

Both authors would like to thank the anonymous reviewers for their comments and suggestions that helped improving the quality of the present manuscript.

References

[1] B. Jähne, Digital Image Processing: Concepts, Algorithms and Scientific Applications, fourth ed., Springer, Berlin, 1997.
[2] I. Pitas, A. Venetsanopoulos, Nonlinear Digital Filters: Principles and Applications, Kluwer Academic Publishers, Norwell, USA, 1999.
[3] P. Saint-Marc, J.-S. Chen, G. Medioni, Adaptive smoothing: a general tool for early vision, IEEE Transactions on Pattern Analysis and Machine Intelligence 13 (6) (1991) 514–529.
[4] P. Mrázek, J.J. Weickert, A. Bruhn, On robust estimation and smoothing with spatial and tonal kernels, in: R. Klette, R. Kozera, L. Noakes, J. Weickert (Eds.), Geometric Properties from Incomplete Data, Springer, Dordrecht, The Netherlands, 2006, pp. 335–352.
[5] C. Tomasi, R. Manduchi, Bilateral filtering for gray and color images, in: Proceedings of ICCV, 1998, pp. 839–846.
[6] G. Winkler, V. Aurich, K. Hahn, A. Martin, Noise reduction in images: some recent edge-preserving methods, Pattern Recognition and Image Analysis 9 (4) (1999) 749–766.
[7] D. Barash, A fundamental relationship between bilateral filtering, adaptive smoothing and the non-linear diffusion equation, IEEE Transactions on Pattern Analysis and Machine Intelligence 24 (6) (2002) 844–847.
[8] A. Buades, B. Coll, J.-M. Morel, Neighborhood filters and PDE's, Numerische Mathematik 105 (1) (2006) 1–34.
[9] C. Chu, I. Glad, F. Godtliebsen, J. Marron, Edge-preserving smoothers for image processing, Journal of the American Statistical Association 93 (442) (1998) 526–556.
[10] A. Buades, B. Coll, J. Morel, A review of image denoising algorithms, with a new one, Multiscale Modeling & Simulation 4 (2) (2005) 490–530.
[11] M. Nitzberg, T. Shiota, Nonlinear image filtering with edge and corner enhancement, IEEE Transactions on Pattern Analysis and Machine Intelligence 14 (8) (1992) 826–833.
[12] R. Paranjape, R. Rangayyan, W. Morrow, Adaptive neighborhood mean and median image filtering, Journal of Electronic Imaging 3 (1994) 360–367.
[13] J. Debayle, J. Pinoli, General adaptive neighborhood image processing. Parts I and II, Journal of Mathematical Image Vision 25 (2) (2006) 245–284.
[14] D. Wang, A. Vagnucci, C. Li, Gradient inverse weighted smoothing scheme and the evaluation of its performance, Computer Vision, Graphics and Image Processing 15 (1981) 167–181.
[15] M. Nagao, T. Matsuyama, Edge preserving smoothing, Computer Graphics and Image Processing 9 (1979) 394–407.
[16] J. Lee, Digital image smoothing and the sigma filter, Computer Vision, Graphics and Image Processing 24 (1983) 255–269.
[17] G. Yang, P. Burger, D. Firmin, S. Underwood, Structure adaptive anisotropic image filtering, Image and Vision Computing 14 (1996) 135–145.
[18] A. Spira, R. Kimmel, N. Sochen, A short time Beltrami kernel for smoothing images and manifolds, IEEE Transactions on Image Processing 16 (6) (2007) 1628–1636.
[19] R. van den Boomgaard, J. van de Weijer, On the equivalence of local-mode finding robust estimation and mean-shift analysis as used in early vision tasks, in: Proceedings of ICPR, vol. 3, 2002, pp. 927–930.
[20] D. Barash, D. Comaniciu, A common framework for nonlinear diffusion, adaptive smoothing, bilateral filtering and mean shift, Image and Vision Computing 22 (1) (2004) 73–81.
[21] P. Perona, J. Malik, Scale space and edge detection using anisotropic diffusion, IEEE Transactions on Pattern Analysis and Machine Intelligence 12 (7) (1990) 629–639.
[22] C. Lantuéjoul, F. Maisonneuve, Geodesic methods in image analysis, Pattern Recognition 17 (1984) 177–187.
[23] B. Verwer, P. Verbeek, S. Dekker, An efficient uniform cost algorithm applied to distance transforms, IEEE Transactions on Pattern Analysis and Machine Intelligence 11 (4) (1989) 425–429.
[24] D. Coeurjolly, D. Miguet, L. Tougne, 2D and 3D visibility in discrete geometry: an application to discrete geodesic paths, Pattern Recognition Letters 25 (5) (2004) 561–570.
[25] P. Soille, Generalized geodesy via geodesic time, Pattern Recognition Letters 15 (12) (1994) 1235–1240.
[26] L. Ikonen, Priority pixel queue algorithm for geodesic distance transforms, Image and Vision Computing 25 (2007) 1520–1529.
[27] G. Borgefors, Distance transformations in digital images, Computer Vision, Graphics and Image Processing 34 (1986) 344–371.
[28] J. Sethian, A fast marching level set method for monotonically advancing fronts, Proceedings of National Academy of Science 93 (4) (1996) 1591–1595.
[29] G. Levi, U. Montanari, A grey-weighted skeleton, Information and Control 17 (1970) 62–91.
[30] L. Ikonen, P. Toivanen, Distance and nearest neighbor transforms on gray-level surfaces, Pattern Recognition Letters 28 (5) (2007) 604–612.
[31] F. Cheng, A. Venetsanopoulos, Adaptive morphological operators, fast algorithms and their applications, Pattern Recognition 33 (2000) 917–933.
[32] R. Lerallut, E. Decencière, F. Meyer, Image filtering using morphological amoebas, Image and Vision Computing 25 (4) (2007) 395–404.
[33] L. Cohen, R. Kimmel, Global minimum for active contour models: a minimal path approach, International Journal of Computer Vision 24 (1997) 57–78.
[34] P. Soille, J. Grazzini, Advances in constrained connectivity, in: Proceedings of DGCI, in: Lecture Notes in Computer Science, vol. 4992, Springer, Berlin, 2008, pp. 423–433.
[35] F. Durand, J. Dorsey, Fast bilateral filtering for the display of high-dynamic range images, ACM Transactions on Graphics 21 (3) (2002) 257–266.
[36] K. Plataniotis, A. Venetsanopoulos, Color Image Processing and Applications, Springer, Berlin, Germany, 2000.
[37] P. Scheunders, A multivalued image wavelet representation based on multiscale fundamental forms, IEEE Transactions on Image Processing 11 (5) (2002) 568–575.
[38] J. Weickert, Anisotropic Diffusion in Image Processing, Teubner, Stuttgart, Germany, 1998.
[39] L. Bertelli, B. Manjunath, Edge preserving filters using geodesic distances on weighted orthogonal domains, in: Proceedings of IEEE ICIP, 2007, pp. 321–324.
[40] G. Sapiro, D. Ringach, Anisotropic diffusion of multivalued images with applications to color filtering, IEEE Transactions on Image Processing 53 (11) (1996) 1582–1586.
[41] T. Sun, Y. Neuvo, Detail-preserving median based filters in image processing, Pattern Recognition Letters 15 (4) (1994) 341–347.
[42] P. Soille, Beyond self-duality in morphological image analysis, Image and Vision Computing 23 (2005) 249–257.
[43] I. Avcibas, B. Sankur, K. Sayood, Statistical evaluation of image quality measures, Journal of Electronic Imaging 11 (2) (2002) 206–223.
[44] G. Gomez, Local smoothness in terms of variance, in: Proceedings of BMVC, vol. 2, 2000, pp. 815–824.

About the Author—JACOPO GRAZZINI is currently a postdoctoral researcher at the Spatial Data Infrastructures of the Joint Research Center. He did a PhD in Computer Science at the French National Institute of Computer Science and Automatic (INRIA) and received his degree from University of Marne-la-Vallée in 2003. He holds a BS in Applied

Mathematics (1998), a MSc in Statistics and Stochastic Models (1999) and a MSc in Computer Science (2000) from the University Paris VII. His main research interests are signal and image processing for the analysis of remote sensed data. He is mainly dealing with non-linear methods applied to segmentation and feature extraction in satellite images and geospatial data.

About the Author—PIERRE SOILLE received the PhD degree from the Universite Catholique de Louvain, Belgium, in 1992 in collaboration with the Centre de Morphologie Mathematique of the Ecole des Mines de Paris and the 'Habilitation a Diriger des Recherches' degree from the Universite Montpellier II, France, in 1997. After his PhD, he pursued research on image analysis and mathematical morphology at the CSIRO Mathematical and Information Sciences, Sydney, the Centre de Morphologie Mathematique, Fontainebleau, France, and the Pattern Recognition Department of the Fraunhofer-Institut IPK, Berlin. During the period 1995–1998, he was an assistant professor with the Ecole des Mines d'Ales and EERIE, Nimes, France. In 1999, he was senior researcher with the Biotechnology and Biological Sciences Research Council, United Kingdom. Since December 1999, he has been with the Joint Research Centre of the European Commission, Ispra, Italy, where he is now working as principal researcher focusing on the development of advanced methodologies and algorithms for the generation of enhanced pan-European geospatial data sets.