P. Levi, R.-J. Ahlers, F. May, M. Schanz (Eds.),
Mustererkennung 1998, Springer, Berlin, pp. 93–100, 1998.

Fast segmentation methods based on partial differential equations and the watershed transformation

Joachim Weickert

Department of Computer Science,
University of Copenhagen,
Universitetsparken 1,
DK-2100 Copenhagen, Denmark
E-mail: joachim@diku.dk

Abstract. Segmentation algorithms are presented which combine regularization by nonlinear partial differential equations (PDEs) with a watershed transformation with region merging. We develop efficient algorithms for two well-founded PDE methods. They use an additive operator splitting (AOS) leading to recursive and separable filters. Further speed-up can be obtained by embedding AOS schemes into a pyramid framework. Examples demonstrate that the preprocessing by these PDE techniques eases and stabilizes the segmentation. The typical CPU time for segmenting a 256^2 image on a workstation is less than 2 seconds.

1 Introduction

Segmentation is one of the bottlenecks of many image analysis and computer vision tasks ranging from medical image processing to robot navigation. Ideally it should be efficient to compute and correspond well with the physical objects depicted in the image.

In the last decade much research on PDE-based regularization methods has been carried out. Although the impressive results suggest that they might be attractive as a preprocessing step for many subsequent image analysis methods, not much research has been done which analyses the combination of PDE-based preprocessing methods with other techniques.

In the present paper we address this topic by combining two efficient PDE-based regularization methods with a watershed algorithm, a simple morphological technique for image segmentation.

This leads to segmentation algorithms which can be generalized in a straightforward way to arbitrary dimensional data sets. Their complexity is linear in the pixel number, and they produce identical results if the image is rotated by 90 degrees. A total CPU time of less than 2 seconds for segmenting a 256^2 image on a workstation makes them attractive for many time-critical applications.

The paper is organized as follows. Section 2 sketches the basic ideas behind two well-posed PDE-based regularization methods: a contrast-enhancing nonlinear diffusion filter and a contrast-diminishing diffusion–reaction method with a similarity term. In Section 3 we develop efficient and reliable numerical techniques for these methods. They are based on an additive operator splitting (AOS). For approximating the diffusion-reaction method, these AOS techniques are embedded into a pyramid framework. In Section 4 we discuss the watershed algorithm with region merging, and in Section 5 we illustrate the usefulness of the combined segmentation process by applying it to several examples. The paper is concluded with a summary in Section 6.

2 PDE-Based Regularization

Below two prototypes for well-posed PDE-based regularization techniques are sketched. The first one allows contrast enhancement, while the second one is contrast-diminishing and can be related to energy minimization methods. A more detailed treatment of diffusion filters can be found in [10].

2.1 The Nonlinear Diffusion Filter of Catté *et al.*

In the m-dimensional case the filter of Catté, Lions, Morel and Coll [3] has the following structure:
Let $\Omega := (0,a_1) \times .... \times (0,a_m)$ be our image domain and consider a (scalar) image $f(x)$ as a bounded mapping from Ω into the real numbers $\mathbb{R}$. Then a filtered image $u(x,t)$ of $f(x)$ is calculated by solving a diffusion equation with the original image as initial state, and reflecting boundary conditions:

$$\partial_t u \;=\; \sum_{l=1}^{m} \partial_{x_l}\Big(g(|\nabla u_\sigma|^2)\,\partial_{x_l} u\Big) \tag{1}$$

$$u(x,0) \;=\; f(x), \tag{2}$$

$$\partial_n u|_{\partial\Omega} \;=\; 0, \tag{3}$$

where n denotes the normal to the image boundary $\partial\Omega$.
The "time" t is a scale parameter: larger values lead to simpler image representations. In order to reduce smoothing at edges, the diffusivity g is chosen as a decreasing function of the edge detector $|\nabla u_\sigma|$, where ∇u_σ is the gradient of a Gaussian-smoothed version of u. We use the diffusivity

$$g(s^2) \;:=\; \begin{cases} 1 & (s^2 = 0) \\ 1 - \exp\left(\frac{-3.315}{(s/\lambda)^8}\right) & (s^2 > 0). \end{cases} \tag{4}$$

For such rapidly decreasing diffusivities, smoothing on both sides of an edge is much stronger than smoothing across it. This selective smoothing process prefers intraregional smoothing to interregional blurring. The factor 3.315 ensures that the flux $\Phi(s) := sg(s^2)$ is increasing for $|s| \leq \lambda$ and decreasing for $|s| > \lambda$. Thus, λ is a contrast parameter separating low-contrast regions with (smoothing) forward diffusion from high-contrast locations where backward diffusion may enhance edges [8]. After some time this filter creates segmentation-like results which are piecewise almost constant. For $t \to \infty$, however, the image reaches a flat steady-state [10]. Well-posedness results for this filter can be found in [3, 10] and a scale-space interpretation in terms of an extremum principle as well as decreasing variance, decreasing energy, and increasing entropy is given in [10].

2.2 A Diffusion-Reaction Filter

Sometimes nonlinear diffusion filtering is modified by supplement it with a reaction term which keeps the steady-state solution for $t \to \infty$ close to the original image f:

$$\partial_t u = \operatorname{div}\left(g(|\nabla u|^2)\,\nabla u\right) + \beta(f-u) \qquad (\beta > 0). \tag{5}$$

This equation can be regarded as the descent method of the energy functional

$$E_f(u) := \int_\Omega \left(\tfrac{\beta}{2}\cdot(u-f)^2 \,+\, \Psi(|\nabla u|)\right) dx, \tag{6}$$

where $\Psi(|\nabla u|)$ is a potential with $\nabla(\Psi(|\nabla u|)) = g(|\nabla u|^2)\nabla u$. The first term in the integrand of the energy functional is a similarity term which causes the regularized image to be close to the original one, while the second term tends to increase its smoothness.

From a theoretical viewpoint, it is advantageous to choose a convex potential function Ψ, since this guarantees well-posedness and stable algorithms [9]. For nonconvex potentials as in [6], several well-posedness questions are open. An example for a diffusivity leading to a convex potential is

$$g(|\nabla u|^2) := \frac{1}{\sqrt{1+|\nabla u|^2/\lambda^2}} \qquad (\lambda > 0). \tag{7}$$

This diffusivity is used in our experiments. It should be noted that diffusivities $g(s^2)$ with a corresponding convex potential $\Psi(s)$ yield a monotonously increasing flux function $\Phi(s)$. Thus, contrast enhancement is not possible.

3 Efficient Algorithms for PDE-Based Regularization

3.1 An AOS Scheme for the Pure Diffusion Filter

Let us first consider finite difference approximations to the m-dimensional diffusion filter of Catté *et al.*. A discrete m-dimensional image can be regarded as a vector $f \in \mathbb{R}^N$, whose components f_i, $i \in \{1, ..., N\}$ display the grey values at the pixels. Pixel i represents the location x_i. Let h_l denote the grid size in the l direction. We consider discrete times $t_k := k\tau$, where $k \in \mathbb{N}_0$ and τ is the time step size. By u_i^k and g_i^k we denote approximations to $u(x_i, t_k)$ and $g(|\nabla u_\sigma(x_i, t_k)|^2)$, respectively, where the gradient is replaced by central differences. The simplest discretization of the diffusion equation (1) with reflecting boundary conditions is given by

$$u_i^{k+1} \;=\; u_i^k + \tau \sum_{l=1}^{m} \sum_{j \in \mathcal{N}_l(i)} \frac{g_j^k + g_i^k}{2h_l^2}\,(u_j^k - u_i^k) \tag{8}$$

where $\mathcal{N}_l(i)$ consists of the two neighbours of pixel i along the l direction (boundary pixels may have only one neighbour). In vector–matrix notation (8) becomes

$$u^{k+1} \;=\; \Big(I + \tau \sum_{l=1}^{m} A_l(u^k)\Big)\, u^k, \tag{9}$$

where A_l describes the diffusive interaction in l direction. Unfortunately, the step size of this scheme has to be very small for stability reasons [11]: it must satisfy $\tau \le 1/\sum_{l=1}^{m} \frac{2}{h_l^2}$. For most practical applications, this restriction requires many iterations, such that the explicit scheme is rather slow.

In order to address the preceding problem we consider instead of (9) the *additive operator splitting (AOS) scheme*

$$u^{k+1} \;=\; \frac{1}{m} \sum_{l=1}^{m} \Big(I - m\tau A_l(u^k)\Big)^{-1} u^k. \tag{10}$$

The operators $I - m\tau A_l(u^k)$ describe one-dimensional diffusion processes along the x_l axes. Under a consecutive pixel numbering along the direction l they come down to strictly diagonally dominant tridiagonal matrices. They can be inverted in linear complexity by a Gaussian algorithm for tridiagonal systems. Moreover, (10) has the same first-order Taylor expansion in τ as (8): both methods are $O(\tau + h_1^2 + ... + h_m^2)$ approximations to the continuous equation. Since (10) is an *additive* splitting, all coordinate axes are treated in exactly the same manner. This is in contrast to conventional splitting techniques which are *multiplicative*. They may produce different results if the image is rotated by 90 degrees.

Recently a general framework for discrete nonlinear diffusion scale-spaces has been discovered, which guarantees that the discretization reveals the same scale-space properties as its continuous counterpart [10]. One can verify [11] that the AOS scheme creates such a discrete nonlinear diffusion scale-space for every (!) step size τ. As a consequence, it preserves the average grey level μ, satisfies a causality property in terms of a maximum–minimum principle, and converges to a constant steady state. Moreover, the process is a simplifying, information-reducing transform with respect to many aspects: The p-norms $\|u^k\|_p := (\sum_{i=1}^{N} |u_i^k|^p)^{1/p}$ and all even central moments $M_{2n}[u^k] := \frac{1}{N}\sum_{j=1}^{N}(u_j^k - \mu)^{2n}$ are decreasing in k, and the entropy $S[u^k] := -\sum_{j=1}^{N} u_j^k \ln u_j^k$, a measure of uncertainty and missing information, is increasing in k (if all f_j are positive).

Many nonlinear diffusion problems require only the elimination of noise and some small-scale details. Usually this can be accomplished with less than 10 iterations in sufficient precision. On an HP 9000/889 workstation, this does not take more than 1 second for a 256^2 image. For further algorithmic details on AOS schemes, the reader is referred to [11].

3.2 A Pyramid AOS Scheme for the Diffusion–Reaction Method

It is possible to extend the preceding AOS framework also to the diffusion–reaction filter (5). It is not hard to see that this can be achieved by using

$$u^{k+1} \;=\; \frac{1}{m}\sum_{l=1}^{m}\Big(I-\frac{m\tau}{1+\tau\beta}A_l(u^k)\Big)^{-1}\,\frac{u^k+\tau\beta f}{1+\tau\beta}. \tag{11}$$

In contrast to the pure diffusion filter, however, we are now interested in approximating the steady-state solution for $t \to \infty$. Even with large time step sizes, the diffusion process will mainly act within a fairly small vicinity around each pixel. Thus, many iterations are required if the image size is large. In order to speed up the process, we may embed the AOS schemes into a pyramid framework. The idea is as follows:

- downsample the image by creating a Gaussian pyramid [2] with the smoothing mask $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$
- adapt the filter parameters to the downsampled image; experiments have shown that λ should be divided by 1.5 and β by 4 if one reduces the image size by a factor 2.
- start with a coarse level (e.g. a 4×4 image), and apply AOS iterations until convergence is obtained.
- expand this solution to the next finer level by linear interpolation, and use it as initial value for AOS iterations at this level.
- proceed in the same way until convergence at the finest level is reached.

Typically, only a few iterations (around 10) are necessary in order to have convergence in terms of a sufficiently small residue at each level. The total complexity remains linear, since the Gaussian pyramid decomposition can be performed in linear complexity. For regularizing a 256^2 image on an HP 9000/889 workstation, a CPU time of 1–1.5 seconds is sufficient.

4 Watershed Segmentation with Region Merging

The previously discussed PDE-based regularization techniques lead to piecewise almost constant image with less noise and less fine-scale details. In order to create a true segmentation, we have to process the regularized image by a technique which gives an edge map without dangling edges. This edge map should lead to a partitioning of the entire image into a finite number of regions, it should

handle the semantically important corners and junctions gracefully, and – last but not least – it should be fast.
We found a watershed technique [1] based on the gradient magnitude very useful for these purposes. Such a technique regards the gradient magnitude image as a landscape where the intensity values correspond to the elevation. Areas where a rain drop would drain to the same minimum are denoted as catchment basins, and the lines separating adjacent catchment basins are called watersheds.
We used an implementation as proposed by Fairfield [4]. The basic idea of this algorithm is sliding downhill on the gradient squared surface until one arrives at a local minimum. Then one replaces all pixels along this path by the image intensity at its corresponding extremum. This algorithm has linear complexity.
Watershed algorithms often create too many segments. Although this oversegmentation is less dominant in the PDE-regularized image than in the original one, it may still lead to problems. They can be circumvented by supplementing the watershed segmentation with a region merging step. In such a step, adjacent regions are merged if their contrast difference is below a specified threshold. This contrast parameter can be related to the contrast parameter λ of the PDE-based regularization, thus it does not constitute an additional parameter. First our algorithm identifies all candidates for merging, and performs the final region merging afterwards. Thus, the whole procedure is independent of the order in which the algorithm runs through the pixels, and it produces a segmentation which is invariant under image rotations by 90 degrees. Since the merging step is also of linear complexity, the total complexity remains linear.
A watershed segmentation of a 256^2 image with subsequent region merging takes about 0.5 CPU seconds on an HP 9000/889. Thus, the total segmentation time including the PDE-based regularization is less than 2 seconds.

5 Experiments

Figure 1 illustrates how preprocessing by nonlinear diffusion filtering greatly reduces the number of segments in a medical MR image. We also observe that under nonlinear diffusion the segment boundaries remain well located and need not be traced back in order to improve their localization as it is common with watershed segmentation based on linear diffusion [5, 7]. The segments correspond well with the real physical objects and segmentation of elongated objects does not create problems.
In Figure 2 it is shown that the merging step can be essential for avoiding the oversegmentation problems in the watershed algorithm. Nonlinear diffusion may create almost piecewise constant areas, but small fluctuations within such an area introduce many semantically irrelevant catchment basins. Such fluctuations can also be caused by quantization errors, e.g. by storing grey values in a bytewise manner. Merging adjacent regions with similar grey values constitutes a simple remedy for these problems.
Finally, Figure 3 gives a comparison between the two PDE-based regularization techniques. We observe a typical behaviour that can also be experienced with

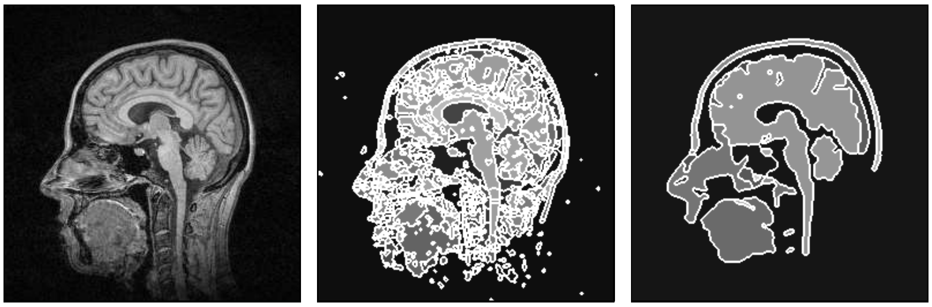

Fig. 1. LEFT: Test image. MIDDLE: Segmented. RIGHT: Diffused and segmented.

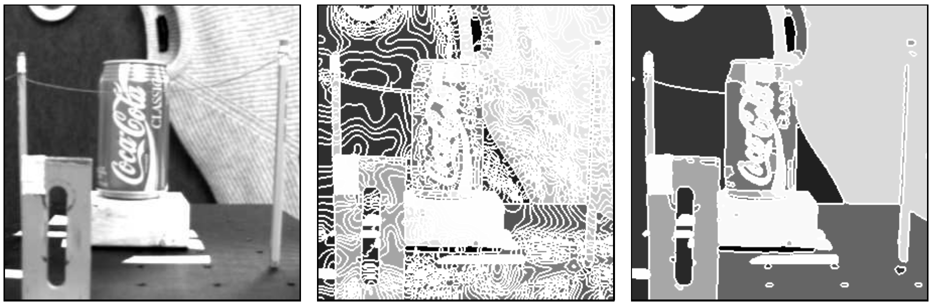

Fig. 2. LEFT: Test image. MIDDLE: Diffused, segmented without merging. RIGHT: Diffused, segmented with merging.

many other test images: the contrast-reducing diffusion-reaction method gives smoother results and is better capable of eliminating small-scale structures, while the edge-enhancing pure nonlinear diffusion process appears to be advantageous for recovering elongated structures. It leads to a very realistic segmentation of the man walking along the hallway.

6 Summary

We have presented efficient algorithms for two PDE-based regularizations. They greatly simplify subsequent segmentation tasks, such that already a simple watershed algorithm with region merging gives good results. Thanks to the AOS and pyramid AOS schemes these segmentation techniques are very fast. This makes them attractive for many time-critical applications. All axes are treated equally, since the result is independent of the pixel ordering, and the algorithm has linear complexity in any dimension.

Acknowledgements. This work was done within the EU-TMR network VIRGO. The hallway image was kindly provided by Ole Fogh Olsen and Mads Nielsen, and the watershed algorithm is a modification of an implementation by Hans Oltmans.

Fig. 3. Hallway scene. LEFT: Original image. MIDDLE: Diffusion–reaction, segmented. RIGHT: Diffusion, segmented.

References

1. S. Beucher, *Segmentation tools in mathematical morphology*, C.H. Chen, L.F. Pau, P.S.P. Wang (Eds.), Handbook of pattern recognition and computer vision, World Scientific, Singapore, 443–456, 1992.
2. P.J. Burt, E.H. Adelson, *The Laplacian pyramid as a compact image code*, IEEE Trans. Comm., Vol. 31, 532–540, 1983.
3. F. Catté, P.-L. Lions, J.-M. Morel, T. Coll, *Image selective smoothing and edge detection by nonlinear diffusion*, SIAM J. Numer. Anal., Vol. 29, 182–193, 1992.
4. J. Fairfield, *Toboggan contrast enhancement for contrast segmentation*, Proc. ICPR 10 (Atlantic City, June 16–21, 1990), Vol. 1, 712–716, 1990.
5. L.D. Griffin, A.C.F. Colchester, G.P. Robinson, *Scale and segmentation of grey-level images using maximum gradient paths*, Image Vision Comput., Vol. 10, 389–402, 1992.
6. N. Nordström, *Biased anisotropic diffusion - a unified regularization and diffusion approach to edge detection*, Image Vision Comput., Vol. 8, 318–327, 1990.
7. O.F. Olsen, *Multiscale watershed segmentation*, J. Sporring, M. Nielsen, L. Florack, P. Johansen (Eds.), Gaussian scale-space theory, Kluwer, Dordrecht, 191–200, 1997.
8. P. Perona, J. Malik, *Scale space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 12, 629–639, 1990.
9. C. Schnörr, *Unique reconstruction of piecewise smooth images by minimizing strictly convex non-quadratic functionals*, J. Math. Imag. Vision, Vol. 4, 189–198, 1994.
10. J. Weickert, *Anisotropic diffusion in image processing*, Teubner-Verlag, Stuttgart, 1998.
11. J. Weickert, B.M. ter Haar Romeny, M.A. Viergever, *Efficient and reliable schemes for nonlinear diffusion filtering*, IEEE Trans. Image Proc., Vol. 7, 398–410, 1998.