

# Non-linear Diffusion for Interactive Multi-scale Watershed Segmentation

Erik Dam and Mads Nielsen

IT-C

The IT University of Copenhagen

Glentevej 67

DK – 2400 Copenhagen NV

Denmark

{erikdam,malte}@it-c.dk

<http://www.it-c.dk>

**Abstract.** Multi-scale watersheds have proven useful for interactive segmentation of 3D medical images. For simpler segmentation tasks, the speed up compared to manual segmentation is more than one order of magnitude. Even where the image evidence does not very strongly support the task, the interactive watershed segmentation provides a speed up of a factor two. In this paper we evaluate a broad class of non-linear diffusion schemes for the purpose of interactively segmenting gray and white matter of the brain in T1-weighted MR images. Through a new scheme GAN, we show that diffusion similar to the nonlinear Perona-Malik scheme is superior to the other evaluated diffusion schemes. This provides a speed up factor of two compared to the linear diffusion scheme.

## 1 Introduction

In many medical segmentation tasks, the images do not sufficiently clearly outline the relevant anatomical structures for making simple automated segmentations. The counter-example is simple thresholding of bone structures in CT images [21]. However, such techniques do not work for most anatomical structures such as soft tissue in CT (due to varying and indistinguishable attenuation [10]), most structures in MR (due to the image inhomogeneities [10]), and most structures in PET/SPECT (due to noise [10]). Here, the alternatives are either performing a tedious manual outline slice per slice, or creating specialised algorithms heavily supported by prior information [6].

The interactive 3D multi-scale watershed segmentation tool,  $\nabla$  *Vision*, may successfully be applied in these situations [8]. The image scale-space is created by Gaussian convolution [35,17]. The watersheds of the gradient magnitude are computed independently at all scales. A linking procedure gives the simpler large scale watersheds the small scale localisation [26,20,15]. The linked watershed regions constitute a multi-scale partitioning of the images.

Ideally, a given anatomical structure may be outlined by a single region. However in most situations, the linked watersheds do not directly compare to the

anatomical structures.  $\nabla Vision$  lets the user arbitrarily change scale and select and deselect regions, and thereby sculpt the anatomical structure. All interaction is geometrical and thereby very intuitive for the clinician. The speed up compared to manual segmentation depends on the interactions required to outline the anatomical structure. Outlining soft tissue in CT may be orders of magnitude faster. Compared to computerised manual segmentation, a complicated task, the masseter (chewing muscle) in MR, shows a speed up factor of two [8,11].

The Gaussian scale-space is the least committed scale-space. A non-linear scale-space commits itself to certain intensity variations through the non-linear function and to certain local edge shapes through the diffusion structure. This is formalised through the connection between energy minimization methods [23] and non-linear diffusion in the biased non-linear diffusion [25]. In this light, one may argue that the use of the non-linear diffusion schemes is the first step in commitment towards using prior shape and intensity knowledge like in the active contour and core-based segmentation methods [6,30].

In this paper, we evaluate a large number of non-linear diffusion schemes for the multi-scale watersheds in 2D: non-linear isotropic Perona-Malik [29], non-linear anisotropic schemes [34], and a purely geometrical non-linear scheme (Mean Curvature Motion) [12]. In section 3.4, we argue that these schemes in a natural way span the space of diffusion schemes supporting segmentation.

The flavor of our paper is close to the comparison of diffusion schemes for segmentation performed on the *hyper-stack* [33,18,24]. The major differences are that the hyper-stack is based on isophote linking and that it constitutes an automated segmentation algorithm. In this paper we specifically evaluate how the *deep structure* of the various scale-spaces support the segmentation.

The multi-scale watershed segmentation method is introduced in section 2. In section 3 the evaluated diffusion schemes are presented. The evaluation is outlined in section 4 with results in section 5.

## 2 Multi-scale Watershed Segmentation


During rain the drops gather in pools. The topology of a landscape defines the regions of support for each pool — the *catchment basins*. The boundaries between the catchment basins are termed *watersheds*. On large scale, the watersheds of a landscape are the ridges and the catchment basins are the dales. The geographical concept *watershed* was introduced to mathematicians in [22].

The watersheds allow a simple partitioning of an image. However, for segmentation purposes the regions border should be defined as the watersheds of a dissimilarity measure instead of the original image. A simple, general, and non-committed dissimilarity measure is the gradient magnitude.

The structures that are outlined by this partitioning are defined with respect to the scale at which the gradient is calculated. Different scales are therefore needed to locate objects of different sizes. The theory of scale-space suggests looking at the *deep structure* [17,35,32] — how the catchment basins develop over scale.

Each catchment basin corresponds to a local minimum for the gradient magnitude. In [26,27], the generic events for the gradient magnitude minima are derived. The conclusion is that *fold annihilation*, *fold creation*, *cusp annihilation*, and *cusp creation* catastrophes are stable and therefore to be expected for typical images. For the catchment basins, this corresponds to the *annihilation*, *creation*, *merge*, and *split* events [26,7].

Linking of the catchment basins across scale combines the simplification at the detection scale with the fine scale precision at the localisation scale (see figure 1). The segmentation method presented in [26] uses these localised basins as building blocks for the segmentation. The user can shift the detection scale and thereby select building blocks appropriate for sculpting the desired objects.



**Fig. 1.** Linking of catchment basins across scale. The catchment basins at detection scale is linked down to the localisation scale and thereby get fine scale precision.

Similar approaches to the multi-scale watershed segmentation of [26] are presented in [20], [14], and [13]. However, an important contribution of [26] is that an intuitive interface is presented that allows the user to interact directly with the three-dimensional building blocks. This forms the basis for an implementation (the program  $\nabla$  *Vision* by *Generic Vision*) that has been tested by clinical researchers with promising results [16,8]. An important addition in this implementation is an extra “artificial” level of voxel-sized building blocks which allows arbitrary detail in the segmentation.

### 3 Diffusion Schemes for Multi-scale Watershed Segmentation

The original multi-scale watershed segmentation method relies on the linear Gaussian scale-space to simplify the image. This simplification determines how the catchment basins group into gradually larger building blocks corresponding to image structures at a given scale.

The linear scale-space for an image  $I(\mathbf{x})$  is described by the PDE  $\frac{\partial L(\mathbf{x};t)}{\partial t} = \Delta L(\mathbf{x};t) = L_{ii}(\mathbf{x};t)$  with the initial condition:  $L(\mathbf{x};0) = I(\mathbf{x})$ . The Gaussian convolution kernel with standard deviation  $\sigma = \sqrt{2t}$  is the Green’s function for the PDE.

The Linear Gaussian diffusion scheme (here denoted LG) has extremely nice theoretical properties. In particular, the causality property, the average gray level invariance property, and the fact that the image gets uniform intensity for scale tending to infinity ensures that the linear scale-space is applicable for the multi-scale watershed segmentation method. However, these properties do not ensure that it is an optimal diffusion scheme for the method.

In [26,27] the generic events for the gradient magnitude minima are derived. When the diffusion scheme is replaced by nonlinear schemes, the analysis of the generic events for the watershed regions is no longer applicable. Some of these nonlinear schemes have been analysed [9]. However, from a practical viewpoint, the linking of the discrete scale levels can handle nearly any diffusion scheme with suitable simplification properties due to robust matching of regions [7].

### 3.1 Regularised Perona-Malik

The classical Perona-Malik diffusion scheme is designed to preserve edges during the diffusion [29]. The regularisation due to [3] is denoted RPM:

$$\frac{\partial L(\mathbf{x}; t)}{\partial t} = \text{div}( p(|\nabla L_\sigma|^2) \nabla L ) \quad \text{where} \quad p(|\nabla L_\sigma|^2) = \frac{1}{1 + \frac{|\nabla L_\sigma|^2}{\lambda^2}} \quad (1)$$

The parameter  $\lambda$  is a threshold for the gradient magnitude required to make the scheme preserve an area (an edge). The  $\sigma$  determines the regularisation scale.

### 3.2 Anisotropic Nonlinear Diffusion

Weickert [34] defines the *anisotropic nonlinear diffusion* equation for a two-dimensional image  $I$  by the PDE:

$$\frac{\partial L(\mathbf{x}; t)}{\partial t} = \text{div}( D(J_\rho(\nabla L_\sigma)) \nabla L ) \quad \text{where} \quad L(\mathbf{x}; 0) = I(\mathbf{x}) \quad (2)$$

The diffusion tensor  $D \in C^\infty(R^{2 \times 2}, R^{2 \times 2})$  is assumed to be symmetric and uniform positive definite. The *structure tensor*  $J_\rho$  is evaluated at *integration scale*  $\rho$  (set to zero in the following), and the gradient  $\nabla L_\sigma$  at *sampling scale*  $\sigma$ .

The diffusion equation possesses simplification properties [34] that ensures that the diffusion schemes are applicable for the segmentation method.

For the following diffusion schemes, the diffusion tensor is defined in terms of the eigenvectors  $\bar{v}_1 \parallel \nabla L_\sigma$ ,  $\bar{v}_2 \perp \nabla L_\sigma$  and the corresponding eigenvalues  $\lambda_1$  and  $\lambda_2$ . Furthermore, Weickert presents a diffusivity function  $w_m$  designed to preserve edges more aggressively than the Perona-Malik diffusivity function  $p$ .

$$w_m(|\nabla L_\sigma|^2) = \begin{cases} 1 & |\nabla L_\sigma| = 0 \\ 1 - \exp\left(\frac{-C_m}{\left(\frac{|\nabla L_\sigma|^2}{\lambda}\right)^m}\right) & |\nabla L_\sigma| > 0 \end{cases} \quad (3)$$

Here  $m$  determines the aggressiveness of the diffusivity function, and  $C_m$  is derived [34,7] from  $m$  such that the flux magnitude function  $|\nabla L| w_m(|\nabla L_\sigma|^2)$  is increasing for  $|\nabla L|^2 < \lambda$  and decreasing for  $|\nabla L|^2 > \lambda$ .

### Isotropic Nonlinear Diffusion

Weickert [34] designs an isotropic nonlinear diffusion scheme (here denoted IND) by the following eigenvalues  $\lambda_1 = \lambda_2 = w_m(|\nabla L_\sigma|^2)$ . Intuitively, this is an increasingly aggressive version of the Perona-Malik scheme for  $m > 0.75$ .

### Edge Enhancing Diffusion

The anisotropic version is termed *edge enhanced diffusion* and defined by the eigenvalues  $\lambda_1 = w_m(|\nabla L_\sigma|^2)$  and  $\lambda_2 = 1$ . The choice  $m = 4$  (which implies  $C_m = 3.31488$ ) is used in [34] with visually appealing results. Here, we also exploit  $m = 2$  ( $C_m = 2.33666$ ) and  $m = 3$  ( $C_m = 2.9183$ ) for the edge enhancing diffusion scheme.

The isotropic nonlinear scheme enhances edges so aggressively that noise is preserved around edges for a long scale interval. The anisotropic schemes (here denoted EE2, EE3, and EE4 depending on the choice of  $m$ ) remedy this by smoothing along the edges.

## 3.3 Mean Curvature Motion

A number of morphological processes possess properties similar to diffusion schemes [2]. *Mean curvature motion* (here denoted MCM) is a special case of the anisotropic nonlinear diffusion scheme [7,34] where  $\lambda_1 = 0$  and  $\lambda_2 = 1$ . MCM is defined by the following PDE (where  $\kappa(L)$  is the mean curvature of the isophote landscape):

$$\frac{\partial L}{\partial t} = -\kappa(L) |\nabla L| \quad (4)$$

## 3.4 Generalised Anisotropic Nonlinear Diffusion

The diffusion schemes previously presented are defined by the diffusivity functions in the gradient direction and the isophote direction. Figure 2 illustrates this “space of diffusion schemes”. This inspires the new *Generalised Anisotropic Nonlinear* diffusion scheme (denoted GAN) defined by the following diffusivity functions  $\lambda_1$  and  $\lambda_2$ :

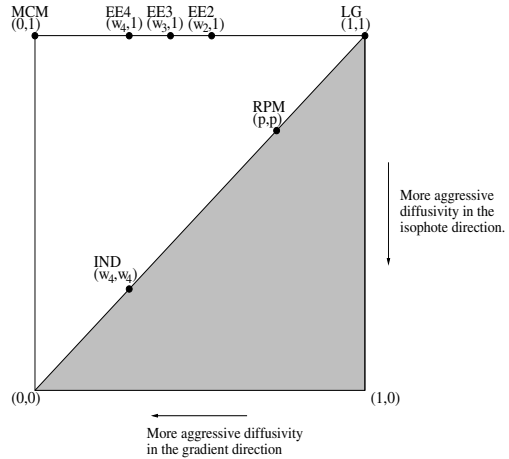
$$\begin{aligned} \lambda_1 &= w(m, \lambda, |\nabla L_\sigma|^2) \\ \lambda_2 &= \theta + (1 - \theta) \lambda_1 \end{aligned} \quad (5)$$

The Weickert diffusivity function (equation 3) is written  $w(m, \lambda, s^2)$  instead of  $w_m(\lambda, s^2)$  since  $m$  is to be perceived as a regular parameter of the diffusivity function. The parameter  $\theta$  determines the degree of anisotropy.

The GAN scheme is named *Generalised Anisotropic Nonlinear diffusion* since it offers a straightforward generalisation of the previously presented diffusion schemes. The choice of the parameters, in particular the aggressiveness parameter  $m$  and the anisotropy parameter  $\theta$ , allows the scheme to cover the white area in figure 2. The existing schemes are realised by the following:

**Fig. 2.** *Space of diffusion schemes.*

The gradient direction and isophote direction diffusivity functions determine the positions on the horizontal and vertical axes, respectively. The diffusion schemes with aggressive diffusivity functions are mapped closest to the lower, left corner. The gray area is populated by diffusion schemes that are not suited for segmentation-like purposes — they diffuse more across edges than along them. Thereby the catchment basins merge across possible object borders before merging inside the regions likely to correspond to the desired objects.



- LG      Linear Gaussian diffusion is defined by  $\lambda \rightarrow \infty$ .
- IND     Isotropic Nonlinear Diffusion is achieved for  $\theta = 0$ .
- RPM    Regularised Perona-Malik scheme is approximated by  $\theta = 0$  and  $m = 0.75$ .
- EEx    The Edge Enhancing schemes EE2, EE3, and EE4 are defined by  $\theta = 1$  and the corresponding  $m$ .
- MCM    Mean Curvature Motion is achieved for  $\theta = 1$  and  $\lambda \rightarrow 0$ .


## 4 Evaluation

An obvious evaluation is to let clinicians test the segmentation method on real segmentation tasks with the different diffusion schemes. However, this is not objective and requires extensive work by the clinicians. The alternative is to measure the quality of the segmentations with respect to a “correct segmentation” — the *ground truth*. The quality measure should be general, objective, and quantitative. For specific segmentation tasks, the quality measure could be defined in terms of specific features of the desired segmentations (shape, topology, etc.). However, for a general evaluation method, the measure must be simple and geometric.

In this paper, ground truth segmentations of white and gray matter for both real and simulated MRI brain scans are used (figures 3 and 4). The real data is from the *Internet Brain Segmentation Repository* [1]. The simulated data is from the *BrainWeb* [4,19,5]. The quality measure is simply defined by the relative error (misclassified pixels relative to number of pixels in ground truth object, where a pixel is uncorrectly segmented if it is included only in the segmentation or only in the ground truth).



**Fig. 3.** Simulated T1 MR brain scan with ground truth white and gray matter. The volume is  $181 \times 217 \times 181$  with coronal slice thickness 1 mm, intensity non-uniformity level 20%, noise level 9%. From <http://www.bic.mni.mcgill.ca/brainweb>



**Fig. 4.** Real T1 MR brain scan of a 55 year old male. The set contains 60  $256 \times 256$  slices with slice thickness 3.0mm. The slices are coronal with a flip angle of 40 degrees. From <http://neuro-www.mgh.harvard.edu/cma/ibsr>

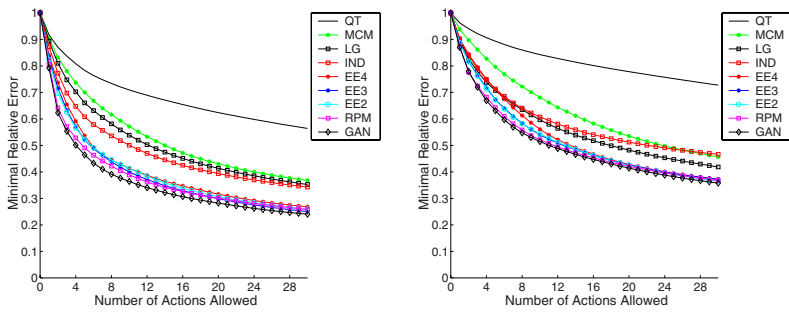
The pixel-sized building blocks allow the user to reach an arbitrary segmentation — therefore any diffusion scheme allow perfect segmentation. The evaluation of the semi-automatic method measures the user effort required to reach a specific *quality threshold*. In an evaluation with clinicians, the user effort could be measured as the time needed. For this evaluation, the effort is naturally measured as the minimal number of basic user actions required. The canonical actions are *selection* and *deselection* of the building blocks.

The quantitative evaluation allows automatic optimisation of the parameter sets for the diffusion schemes. For details on the evaluation method, the algorithm for establishing the optimal combinations of building blocks, the optimisation method, and the optimal parameter sets, see the technical report [7].

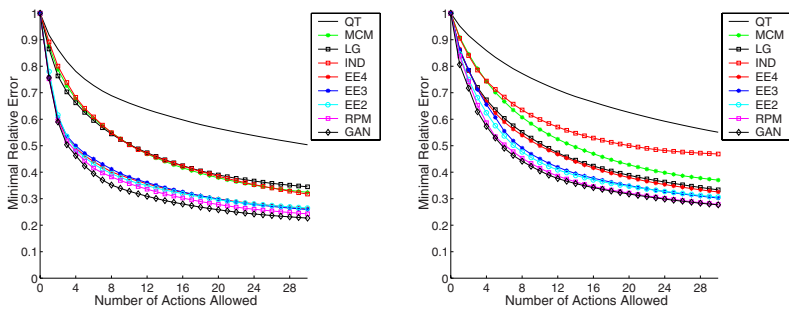
## 5 Results

The performance is illustrated by the error as a function of the number of actions for each scheme (with parameters optimised to the specific data set). This is put into perspective by the performance of a *Quad tree* linking scheme [31] (here denoted QT). The number of quad tree blocks required has a close relation to the *box-counting dimension* of the ground truth objects — also denoted the *Hausdorff dimension* [28]. Furthermore, the building blocks of the quad tree are not adapted to the geometry of the image. Thereby the performance of the quad tree linking gives a frame of reference for the performances of the diffusion schemes defined in terms of the complexity of the ground truth objects.

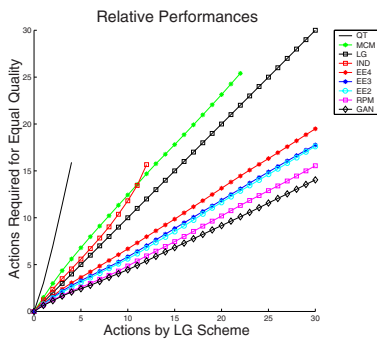
Figures 5 and 6 display the performance on the simulated and the real data, respectively. The best of the existing schemes is the regularised Perona-Malik scheme. The new GAN scheme is slightly better than this. However, the graphs do not give a clear notion of the quantitative differences in performance. Figure 7 delivers the desired relative performance indicator on average for all data sets. The relative performance is determined both for the training set used for



**Fig. 5.** Evaluation on simulated data from figure 3. White matter is left and gray matter is right. Slices 60, 80, 100, 120, and 140 from the data set are used.



**Fig. 6.** Evaluation on real data from figure 4. Left is white matter and right is gray matter. Slices 10, 20, 30, 40, and 50 from the data set are used.



Scheme	Performance	
	Training	Independent
QT	3.83	3.34
MCM	1.18	1.19
LG	1.00	1.00
IND	1.20	1.17
EE4	0.66	0.66
EE3	0.59	0.63
EE2	0.58	0.62
RPM	0.51	0.51
GAN	0.46	0.49

**Fig. 7.** Average performances with the Gaussian scheme as reference. For a given number of actions the performance of the LG scheme is noted — for each scheme is measured the actions required to get equal quality. The incline of a curve determines the number of actions required to obtain a given segmentation quality relative to the performance of the Gaussian scheme. This performance indicator is displayed in the table for each diffusion scheme. The new GAN scheme requires less than half as many actions compared to the Gaussian scheme for both the training set and an independent data set.



optimisation of the parameters for the diffusion schemes and for an independent data set.

The results show that Perona-Malik is superior among the existing diffusion schemes — somewhat surprising even the anisotropic schemes show inferior performance. Furthermore, the optimal parameter sets for the GAN scheme reveal the optimal degrees of anisotropy and edge preservation aggressiveness. The anisotropy parameter  $\theta$  is close to zero for all data sets (0.06, 0.12, 0.0, and 0.0). The improved performance compared to RPM is due to slightly increased aggressiveness in the diffusivity function (the parameter  $m$  is 1.4, 1.0, 1.1, and 1.0 compared to the approximate value of 0.75 for RPM).

In the technical report [7], a number of other results are documented as well:

- With a tolerance area 1 pixel wide around the borders of the ground truth segments, 97% of the ground truth pixels can be segmented without the use of pixel-sized building blocks.
- A simple user heuristic requires around one third more user actions than the optimal actions used by the evaluation method for all diffusion schemes.
- The diffusion schemes require up to 30 scale levels for the discrete linking to be sufficiently closely discretised.
- Similar results are measured for higher number of actions, from simulated data with less noise and from real data of a schizophrenic brain.

## 6 Conclusion

We have constructed a generalised anisotropic diffusion scheme GAN capturing many known diffusion schemes. Tuning this for interactive multi-scale watershed segmentation of white/gray matter in T1-weighted MR slices of the brain shows that diffusion similar to regularised Perona-Malik is superior to the other diffusion schemes. Furthermore, the aggressiveness in the diffusion cut-off is more important than the degree of anisotropy. The best among the tested diffusion schemes yields a decrease in interaction time with more than a factor two compared to linear Gaussian scale-space. Our expectation is that the gain is even higher in a 3D implementation (if nothing else changes, the expected speed up factor in  $ND$  is approximately  $\sqrt{2}^N$ ). The conclusion is linked to the segmentation task of white/gray matter in the brain. For cases like vessels or abdominal organs, other diffusion schemes may be optimal.

## Acknowledgements

A large part of the work that this paper is based on was conducted at DIKU — the Department of Computer Science, University of Copenhagen, Denmark.

## References

1. The internet brain segmentation repository, 1999. MR brain data set 788\_6\_m and its manual segmentation was provided by the Center for Morphometric Analysis at Massachusetts General Hospital and is available at <http://neuro-www.mgh.harvard.edu/cma/ibsr>

2. L. Alvarez and J.-M. Morel. *Geometry-Driven Diffusion in Computer Vision*, chapter Morphological Approach to Multiscale Analysis: From Principles to Equations. Kluwer, 1994.
3. F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal of Numerical Analysis*, 29:182–193, 1992.
4. C.A. Cocosco, V. Kollokian, R.K.-S. Kwan, and A.C. Evans. Brainweb: Online interface to a 3D MRI simulated brain database. In *3rd Int. Con. on Functional Mapping of the Human Brain*, volume 5 of *NeuroImage*, page 425, 1997. <http://www.bic.mni.mcgill.ca/brainweb/>.
5. D.L. Collins, A.P. Zijdenbos, V. Kollokian, J.G. Sled, N.J. Kabani, C.J. Holmes, and A.C. Evans. Design and construction of a realistic digital brain phantom. *IEEE Transactions on Medical Imaging*, 17(3):463–468, June 1998.
6. T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Active shape models: Their training and application. *CVIU*, (1):38–59, 1995.
7. Erik Dam. Evaluation of diffusion schemes for watershed segmentation. Master's thesis, Uni. of Copenhagen, 2000. Report 2000/1 on <http://www.diku.dk/research/techreports/2000.htm>.
8. E. Dam, P. Johansen, O.F. Olsen, A. Thomsen, T. Darvann, A.B. Dobrzyniecki, N.V. Hermann, N. Kitai, S. Kreiborg, P. Larsen, M. Lillholm and M. Nielsen. Interactive multi-scale segmentation in clinical use. In *CompuRAD, European Congress of Radiology 2000*, March 2000.
9. James Damon. Generic properties of solutions to partial differential equations. *Arch. Rational Mech. Anal.*, (140):353–403, 1997.
10. Steve Webb (editor). *The Physics of Medical Imaging*. Inst. of Physics Publ., London, 1988.
11. S. Murakami et al. Three and four dimensional diagnosis for TMJ. In *Proc. of CAR 98*, 88–91, 1998.
12. M. Gage. An isoperimetric inequality with applications to curve shortening. *Invent. Math.*, (76):357–364, 1983.
13. John M. Gauch. Image segmentation and analysis via multiscale gradient watershed hierarchies. *IEEE Transactions on Image Processing*, 8(1):69 – 79, January 1999.
14. J.M. Gauch and S.M. Pizer. Multiresolution analysis of ridges and valleys in grey-scale images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(6):635–646, 1993.
15. L.D. Griffin and A.C.F. Colchester. Superficial and deep structure in linear diffusion scale space: isophotes, critical points and separatrices. *Image and Vision Computing*, 13(7):543–557, 1995.
16. Peter Johansen, Mads Nielsen, and Sven Kreiborg. The computation of natural shape, 1999. <http://www.diku.dk/research-groups/image/research/NaturalShape/>.
17. Jan J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.
18. André Koster. *Linking Models for Multi-scale Image Sequences*. PhD thesis, Utrecht, 1995.
19. R.K.-S. Kwan, A.C. Evans, and G.B. Pike. An extensible MRI simulator for post-processing evaluation. *Lecture Notes in Computer Science*, 1131:135–140, 1996.
20. Lawrence Lifshitz and Stephen Pizer. A multiresolution hierarchical approach to image segmentation based on intensity extrema. *IEEE PAMI*, 12(6):529 – 540, June 1990.
21. William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics*, 21(4):163–169, 1987.
22. J.C. Maxwell. *On Hills and Dales*. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 4th Series, 40(269):421–425, December 1870.
23. D. Mumford and J. Shah. Boundary detection by minimizing functionals. In *CVPR85*, 22–26, 1985.
24. W.J. Niessen, K.L. Vincken, J.A. Weickert, and M.A. Viergever. Nonlinear multiscale representations for image segmentation. *Comp. Vision & Image Understanding*, 66:233–245, 1997.
25. K.N. Nordstrom. Biased anisotropic diffusion: A unified regularization and diffusion approach to edge detection. *IVC*, 8(4):318–327, 1990.
26. Ole Fogh Olsen. Multi-scale segmentation of grey-scale images. Technical Report 96/30, Department of Computer Science, University of Copenhagen, 1996.
27. Ole Fogh Olsen and Mads Nielsen. Generic events for the gradient squared with application to multi-scale segmentation. In *Scale-Space Theory in Computer Vision*, volume 1252 of *Lecture Notes in Computer Science*, 101–112, Utrecht, The Netherlands, 1997.
28. Edward Ott. *Chaos in Dynamical Systems*. Cambridge University Press, 1993.
29. Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE PAMI*, 12(7):629 – 639, July 1990.
30. S.M. Pizer, C.A. Burbeck, J.M. Coggins, D.S. Fritsch, and B.S. Morse. Object shape before boundary shape: Scale-space medial axes. *JMIV*, 4:303–313, 1994.
31. H. Samet. The quadtree and related hierarchical data structures. *Surveys*, 16(2):187–260, 1984.
32. Jon Sporring, Mads Nielsen, Luc Florack, & Peter Johansen. *Gaussian Scale-Space Theory*. Kluwer Academic Publishers, Dordrecht, 1997.
33. Koen Vincken. *Probabilistic Multi-scale Image Segmentation by the Hyperstack*. PhD thesis, University of Utrecht, 1995.
34. Joachim Weickert. *Anisotropic Diffusion in Image Processing*. B. G. Teubner, Stuttgart, 1998.
35. Andrew P. Witkin. Scale-space filtering. In *Proceedings of International Joint Conference on Artificial Intelligence*, 1019–1022, Karlsruhe, Germany, 1983.