

# Image filtering using morphological amoebas <sup>☆</sup>

Romain Lerallut <sup>\*</sup>, Étienne Decencière, Fernand Meyer

*Centre de Morphologie Mathématique, École des Mines de Paris, 35 rue Saint-Honoré, 77305 Fontainebleau, France*

Received 22 September 2005; received in revised form 6 March 2006; accepted 11 April 2006

## Abstract

This paper presents morphological operators with non-fixed shape kernels, or amoebas, which take into account the image contour variations to adapt their shape. Experiments on grayscale and color images demonstrate that these novel filters outperform classical morphological operations with a fixed, space-invariant structuring element for noise reduction applications. Tests on synthetic 3D images are then performed to show the high noise-reduction capacity of amoeba-based filters.

© 2006 Elsevier B.V. All rights reserved.

*Keywords:* Anisotropic filters; Noise reduction; Morphological filters; Color filters; 3D image processing

## 1. Introduction

Noise is possibly the most annoying problem in the field of image processing. There are two ways to work around it: either design particularly robust algorithms that can work in noisy environments, or try to eliminate the noise in a first step while losing as little relevant information as possible and consequently use a normally robust algorithm [Fig. 1](#).

There are of course many algorithms that aim at reducing the amount of noise in images. Most are quite effective but also often remove thin elements such as canals or peninsulas. Even worse, they can displace the contours and thus create additional problems in a segmentation application.

In mathematical morphology, we often couple one of these noise-reduction filters to a reconstruction filter that attempts to reconstruct only relevant information, such as contours, and not noise. However, a faithful reconstruction can be problematic when the contour itself is corrupted by noise, as seen in [Fig. 1](#). This can cause great problems in some applications which rely heavily on clean contour surfaces, such as 3D visualization, so a novel approach was proposed: morphological amoebas.



An amoeba (here *Amoeba proteus*) is a genus of protozoa that moves by projecting pseudopods and is a well-known representative unicellular organism. They are found in sluggish waters all over the world, both fresh and salt, as well as in soils and as parasites. They now begin a new life in the field of image processing.

## 2. Amoebas: dynamic structuring elements

### 2.1. Principle

#### 2.1.1. Classic filter kernel

Formally at least, classic filters work on a fixed-size sliding window, be they morphological operators (erosion, dilation) or convolution filters, such as the diffusion by a

<sup>☆</sup> This paper is an extended version of the one presented at the ISMM'05 congress [\[10\]](#).

<sup>\*</sup> Corresponding author. Tel.: +33 1 64694767.

E-mail address: [lerallut@cmm.ensmp.fr](mailto:lerallut@cmm.ensmp.fr) (R. Lerallut).




Fig. 1. Classic noise filtering (b) removes much contour information. Reconstruction (c) finds not only the contours, but also all the noise connected to the object.

Gaussian. If the shape of that window does not adapt itself to the content of the image (see Fig. 2), the results deteriorate. For instance, an isotropic Gaussian diffusion smooths the contours when its kernel steps over a strong gradient area.

2.1.2. Amoeba filter kernel

Having made this observation, Perona and Malik [1] (and others after them) have developed anisotropic filters that inhibit diffusion through strong gradients. Most early work on non-fixed shape structuring elements, such as [11] and [7] were restricted either in the types of operations performed (openings and closings) or in the type of images upon which they were to be used.

We were inspired by these examples to define a more general framework to develop filters, either morphological or not, whose kernels adapt to the content of the image in order to keep a certain homogeneity inside each structuring element (see Fig. 3) while at the same time keeping their size in check. Tomasi and Manduchi have described in [5] the coupling of a geometric distance between pixels with a distance between their values, which offers remarkable properties for our intended use.

The interest of this approach, compared to the one based on partial differential equations, is that it does not depart greatly from what we use in mathematical morphology, and therefore most of our algorithms can be made to use amoebas with little additional work. Most of the underlying theoretical groundwork for the morphological




Fig. 3. Closing of an image by an amoeba. The amoeba does not cross the contour and as such preserves even the small canals.

approach has been described by Jean Serra in his study [2] of structuring functions, although until now it has seen little practical use.

The shape of the amoeba must be computed for each pixel around which it is centered. Fig. 4 shows the shape of an amoeba depending on the position of its center. Note that in flat areas such as the center of the disc, or the background, the amoeba is maximally stretched, while it is reluctant to cross contour lines.

When an amoeba has been defined, most morphological operators and many other types of filters can be used on it: median, mean, rank filters, erosion, dilation, opening, closing, even more complex algorithms such as reconstruction filters, levelings, floodings, etc.




Fig. 2. Closing of an image by a large structuring element. The structuring element does not adapt its shape and merges two distinct objects.




Fig. 4. Shape of an amoeba at various positions on an image.

## 2.2. Construction

### 2.2.1. Amoeba distance

In general, a filtering kernel of radius  $r$  is formally defined on a square (or a hexagon) of that radius, that is to say on the ball of radius  $r$  relative to the norm associated to the chosen connectivity. We will keep this definition changing only the norm, using one that takes into account the gradient of the image, so that we get the behavior described in Fig. 5.

**Definition 1.** Let  $d_{\text{pixel}}$  be a distance defined between the values of the image, for example a difference of gray-value, or a color distance.

Let  $\sigma = (x = x_0, x_1, \dots, x_n = y)$  be a path between points  $x$  and  $y$ . Let  $\lambda$  be a real positive number. The value of the image at position  $x$  will be noted as  $\text{Image}(x)$ . The length of the path  $\sigma$  is defined as

$$L(\sigma) = \sum_{i=0}^{n-1} [1 + \lambda \cdot d_{\text{pixel}}(\text{Image}(x_i), \text{Image}(x_{i+1}))]$$

The “amoeba distance” with parameter  $\lambda$  is thus defined as:

$$\begin{cases} d_\lambda(x, x) = 0 \\ d_\lambda(x, y) = \min_{\sigma} L(\sigma), \text{ minimum taken on all paths between } x \text{ and } y \end{cases}$$

It is important to realize that  $d_{\text{pixel}}$  has no geometrical aspect, it is a distance computed only on the values of the pixels of the image. Furthermore, if  $n$  is the number of pixels of a path  $\sigma$ , then  $L(\sigma) \geq n$  (since  $\lambda \geq 0$ ), which bounds the maximal extension of the amoeba.

This distance also offers an interesting inclusion property:

**Property 1.** At a given radius  $r$ , the family of the balls  $\mathcal{B}_{\lambda,r}$  relative to the distance  $d_\lambda$  is decreasing (for the inclusion),

$$\begin{aligned} 0 \leq \lambda_1 \leq \lambda_2 &\Rightarrow \forall (x, y), d_{\lambda_1}(x, y) \leq d_{\lambda_2}(x, y) \\ &\Rightarrow \forall r \in \mathbf{R}^+, \mathcal{B}_{\lambda_1,r} \supset \mathcal{B}_{\lambda_2,r} \end{aligned}$$

which may be useful when building hierarchies of filters, such as a family of alternate sequential filters with strong gradient-preserving properties.

### 2.2.2. The pilot image

We have found that the noise in the image can often distort the shape of the amoeba. For this reason, we compute the *shape* of the amoeba on another image. Once the shape is computed, the values are sampled on the *original* image and processed by the filter (mean, median, max, min, ...).

Usually, the other image is the result of a strong noise removal filtering of the original image that dampens the noise while preserving as much as possible the larger contours. A large Gaussian works fairly well, and can be applied very quickly with advanced algorithms, however we will see below that iterating amoeba filters yields even better results.

Furthermore, using the same pilot image means that the amoebas in two successive runs will have the same shape, which result in the idempotence property in openings and closings.

## 3. Amoebas in practice

### 3.1. Adjunction

Erosions and dilations can easily be defined on amoebas. However, it is necessary to use *adjoint* erosions and dilations when using them to define openings and closings:

$$\begin{aligned} \delta(X) &= \cup_{x \in X} B_{\lambda,r}(x) \\ \epsilon(X) &= \{x / B_{\lambda,r}(x) \subset X\} \end{aligned}$$

These two operations are at the same time adjoint and relatively easy to compute, contrary to the symmetrical ones that use the transposition, which is not easy to compute for amoebas. See [2] for a discussion of the various forms of adjunction and transposition of structuring functions.

**Algorithms.** The algorithms used for the erosion and dilation are quite similar to those used with regular structuring elements, with the exception of the step of computing the shape of the amoeba.

Erosion (gray-level):

**for each pixel  $x$  do**

**compute** the shape of the amoeba centered on  $x$   
**compute** the minimum  $M$  of the pixels in the amoeba  
**set** the pixel of the output image at position  $x$  to value  $M$

Dilation (gray-level):

**for each pixel  $x$  do:**

**compute** the shape of the amoeba centered on  $x$   
**for each pixel  $y$  of the amoeba do:**  
**set** value( $y$ ) = max(value( $y$ ), value( $x$ ))

The opening using these algorithms can be seen as the gray-level extension of the classic binary algorithm of first




Fig. 5. Behavior of an amoeba on various relief types. Strong gradients should slow or even hamper the growth of the amoeba.

taking the centers of the circles that fit inside the shape (erosion), and then returning the union of all those circles (dilation).

### 3.2. Complexity

The theoretical complexity of a simple amoeba-based filter (erosion, dilation, mean, median) can be asymptotically approximated by:

$$T(n, k, \text{op}) = O[n * (\text{op}(k^d) + \text{amoeba}(k, d))]$$

where  $n$  is the number of pixels in the image,  $d$  is the dimensionality of the image (usually 2 or 3),  $k$  is the maximum radius of the amoeba,  $\text{op}(k^d)$  is the cost of the operation and  $\text{amoeba}(k, d)$  is the cost of computing the shape of the amoeba for a given pixel.

The shape of the amoebas is computed by a common region-growing implementation using a priority queue. Depending on the priority queue used, the complexity of this operation is slightly more than  $O(k^d)$  (see [3] and [4] for advanced queueing data structures).

Therefore, for erosion, dilation or mean as operators, we have a complexity of a little more than  $O(n * k^d)$  which is the complexity of a filter on a fixed-shape kernel. It has indeed been verified in practice that, while being quite slower than with fixed-shape kernels (especially optimized ones), filters using amoebas tend to follow rather well the predicted complexity, and do not explode. Tests have been

performed on 3D images, size  $512 \times 512 \times 100$ , with amoebas with sizes up to  $21 \times 21 \times 21$ , with computation time ranging from a few minutes to a few hours, depending on the parameters.

## 4. Results

### 4.1. Alternate sequential filters

The images of Fig. 6 compare the differences between alternate sequential filters (ASF) built on classic fixed shape kernels and ASFs on amoebas in the filtering of the image of a retina. The filter should be able to reduce the amount of background noise while preserving the shape of the vessels.

### 4.2. Median and mean

In the context of image enhancement, we have found that a simple mean or median coupled with an amoeba forms a very powerful noise-reduction filter. The images in Fig. 7 show the median and the mean computed on amoebas compared to those built on regular square kernels. The pilot image that drives the shape of the amoeba is the result of a standard Gaussian filter of size 3 on the original image, and the distance  $d_{\text{pixel}}$  is the absolute difference of gray-levels.




Fig. 6. Alternate sequential filters on classic kernels and on amoebas. The amoeba preserves extremely well the blood vessels while strongly flattening the other areas.




Fig. 7. Results of a “classic” median filtering and two amoeba-based filterings: a median and a mean on Edouard Manet’s painting “Le fifre”.

For the filters using amoebas, the median filter preserves well the contour, but the mean filter gives a more “aesthetically pleasing” image. In either case, the results are clearly superior to filterings by fixed-shape kernels, as seen in the Fig. 7.

#### 4.3. Mean and median for color images

In the case of color images, the mean is replaced by the mean on each color component of the RGB color space. For the “median”, the point closest to the barycenter is chosen. Other distances or colorspace can be used, depending on the application, the type of noise and the quality of the color information (Fig. 8).

#### 4.4. Iteration

The quality of the filtering strongly depends on the image that determines the shape of the amoeba. The previous examples have used the original image filtered by a Gaussian, but this does not always yield good results (also see [6]).

A small detail of the image may be excessively smoothed in the pilot image, and thus disappears completely in the result image. On the other hand, noisy pixels may be left untouched if the pilot image does not eliminate them. A possible solution is to improve the quality of the pilot image, so that it helps the amoeba in preserving these features. Such an image should be well-smoothed in flat regions, while preserving as well as possible the contour information. One good method

to compute such an image would be of course to use an amoeba-based filter!

We will proceed in two steps: the first one follows the scheme described earlier, using the Gaussian-filtered original image as a pilot, with large amoebas, and outputs a well-smoothed image in flat areas while preserving as much as possible the most important contours. The second step takes the original image as input and the filtered image as a pilot, with smaller amoebas. These amoebas do not need to be as large as the first ones, since their shapes will be computed on a very smooth image, and therefore they will preserve well the finer details.

Although this refinement of the pilot image could be iterated, we have found in practice that once is enough to reduce the noise dramatically (see Fig. 9).

Pilot images obtained by a first amoeba-based filtering are also very useful for color images, since the amoeba-based pilot image provides better color coupling through the use of an appropriate color distance than simply merging the results of a Gaussian filtering of each channel independently.

Indeed, “perceptual distances” (such as LAB) often merge information from the separate channels in order to return a value closer to what a human would perceive.

## 5. Application to 3D images

### 5.1. 3D images and contour noise

The visualization algorithms for 3D volumic data often use local gradient information for the computation




Fig. 8. Color images: results of a “classic” median filtering, and two amoeba-based filterings: a median and a mean. As a simple extension of the grayscale approach, each channel of the pilot image has been independently smoothed by a Gaussian of size 3.

of the shading of the voxels and many applications, especially in the medical field, use mostly the contour information to visualize the various objects and as such are very vulnerable to noise on the contours (see Fig. 10). While noise reduction is important when processing 2D images, Boehm showed in [9] that a very small amount of noise can almost entirely prevent the visualization of 3D medical images when it is located at the contours (see Fig. 10).

When displaying directly 3D data (as opposed to studying a stack of 2D images), it is essential that the user be able to see the objects they are interested in. This is why most modern renderers include a complex transparency and shading model that makes it possible to peek far inside the image to display the interesting objects. To make a quantitative analysis easier, a synthetic image was created that presents many similarities with 3D scanner images, especially images of the cardiac region: strong textures and thin vessels to preserve.

Fig. 10 illustrates the problem due to strong noise in a 3D image and Fig. 11 shows that levelings are ineffective as they suppress the noise inside the objects but not on the contours. Most basic morphological tools are defined using set theory, and thus are adimensional. Amoebas are no exception to this rule, being the balls defined by a specific distance. Therefore, they can be used without modification on 3D images.

### 5.2. Median filtering with amoebas

One important aspect of the filtering of many medical images (MRI, scanner) is that those images are monospectral: coloring and, ultimately, tissue identification is done using a look-up table. This means that a shifting of the values may have a dramatic effect on the visualization and consequently on the interpretation of the images. This is why we have chosen to first test the median: with its property of returning only values existing in the image, the median lessens the risk of misinterpretation. However, traditional median filtering does not preserve well the contours, and may remove small details which may be crucial to a physician’s analysis, hence the use of amoebas.

For each pixel, the processing is done in two steps: first compute the shape of the amoeba centered on the pixel and then sample the values of the pixel inside the amoeba, feed them to the median operator and write the result at the center of the amoeba in the output image.

### 5.3. Results for 3D images

Fig. 12 shows the result of a median filter computed on amoebas. It is plain to see that most of the transparency effect has been preserved, which indicates that most of the noise has been filtered on the contours as well as inside the objects.




Fig. 9. Comparison between two pilot images: a Gaussian one, and one based on a strong amoeba-based filtering. With the amoeba pilot image the hand is better preserved, and the eyebrows do not begin to merge with the eyes, contrary to the Gaussian-based pilot image. Having both less noise and stronger contours in the pilot image also enables the use of smaller values on the lambda parameter so that the amoeba will stretch more in the flatter zones, and thus have a stronger smoothing effect in those zones, while preserving the position and dynamics of the contours.

Fig. 13 is a zoom of Fig. 12 centered on the thin structures. Once again we have good results in the preservation of those elongated structures. This is absolutely essential since this type of features include critical body parts such as blood vessels, nerves, etc.

It is extremely hard to quantify *in a meaningful way* the results of such a filtering, especially since the usual signal-to-noise ratio does not express well the fact that we may tolerate small variations in many places but not a few strong variations in critical areas. A more meaningful measure for this problem was developed by Boehm in [9]. This measure is tied to a 3D volume renderer (ray-caster) by comparing the visibility of each voxel in the filtered image with that of the corresponding voxel in the ideal image. When a voxel in the filtered image contributes in the same amount to the visualization as the corresponding voxel in the ideal image, then its associated quality measure is equal to one. The greater the difference between both contributions, the closer to zero the measure will be. Results are then averaged in the whole image as well as separately on each component.

Fig. 14 shows a comparison between an amoeba-based median filter and a classic alternate sequential filter (ASF). The results show clearly that while the ASF is as good as the amoeba on the larger structures (spheres B2, B3 and B5), it fails completely to preserve the thin structures such as the rings. Furthermore, there is a clear shift of the values, especially on the outermost rings, which may cause diagnostic errors.

#### 5.4. Simple optimizations

An important fact to take into account is that the radius parameter is like an amount of energy given to the amoeba. It can be used either to climb slopes (with a penalty given by the  $\lambda$  parameter) or it can be used to expand in flat areas. This amount of energy needs to be quite high so that the amoeba can jump over noisy pixels (though not too high so that it does not cross too much over strong gradient lines). However such a high energy means that in flat areas the amoeba will grow to a very large size, which means that not only will the shape be costly to compute but the resulting sample of pixel values will be quite large and so the filter operator will




Fig. 10. A synthetic 3D volume presenting many similarities with medical images of the cardiac regions. Left, original image. Right, image with addition of noise.




Fig. 11. Although the median filters much noise away, the reconstruction (needed to recover the position of the contours) reconstructs most of the noise on the borders of the object, removing most of the transparency.




Fig. 12. The median applied to the amoeba recovers most of the information very well, including thin details. The cleaner gradient on the contours results in a near-perfect transparency effect.






Fig. 13. The amoeba-based median filter recovers very well the contours, as well as the transparency.





Fig. 14. Comparison of the results of an amoeba-based median filter and an alternated sequential filter of size up to 3. The measure is performed on each component of the image (spheres B1–B5, parallelepiped P1 and P2, rings, and total average).

be accordingly long. A very simple yet dramatically effective optimization is thus to impose an upper bound on the size of the amoeba, the value of which depends on the type of noise and the characteristic size of the image elements. This can reduce the cost of computation by an order of magnitude without any detectable loss of effectiveness.

Another form of optimization is performed when applying a smoothing filter to compute the pilot image. This smoothing will reduce small noise without displacing too much the contours and enable the use of smaller amoebas for the same result.

## 6. Conclusion and future work

We have presented here a new type of structuring element that can be used in many morphological algorithms. By taking advantage of outside information, filters built upon those structuring elements can be made more robust on noisy images and in general behave in a “more sensible” way than those based on fixed-shape structuring elements. In addition, morphological amoebas are very adaptable and can be used on color images as well as monospectral ones and, like most morphological tools, they can be used on images of any dimension (2D, 3D, ...). Depending on the application, alternate sequential filters are very effective when looking for very flat zones, whereas median and mean filters output smoother images that may be more pleasing to the eye but could be harder to segment.

It is possible to use amoebas to create reconstruction filters and floodings that take advantage of the ability to parametrize the shape of the amoebas based on the image content. However, the behaviors of the amoebas are much more difficult to take into account when they are used in such complex algorithms. In particular, most of these algorithms are defined for structuring elements that have a radius of at most one pixel. In order to use them with amoebas, it is necessary to ensure that the amoebas have a maximum radius of one pixel. More complex implementations of both the amoebas and the algorithms can be used to work around this problem, but they are beyond the scope of this paper.

The results show that simple extensions of the scalar algorithms to the RGB space already yield excellent results, especially when iterating. The use of more “perceptual” distances (HLS or LAB) would probably prevent some unwanted blending of features, although this is as yet conjectural and will be the basis of further work.

The filtering of 3D images by morphological amoebas, though still in its infancy, seems very promising. Expressing the coupling between image data and geometry through a kernel makes it possible to implement a much larger range of filters to an image than was possible before. Future work will include quantitative as well as qualitative comparison on real medical 3D images. Another area where improvement is to be expected is the computation of the shape of the amoeba. Not only should it be possible

to use elaborate gradient estimation such as proposed in [6], but also providing more complex behaviors for the amoebas, such as an incompressible minimum element, to guarantee at least some diffusion, or on the contrary a minimum size requirement to prevent diffusion through small holes.

We hope to be able to run comparisons between amoebas and well-known anisotropic filtering methods, both in terms of quality and computation time. However such comparisons are not really meaningful without expert tuning of both algorithms, which was not available to us by press time although the subject will certainly be the matter of a subsequent paper.

Finally, it is important to notice that this amoeba framework is general enough to accommodate other types of distances. The distance presented here couples geometry and grey levels (or color distances), but other similar schemes can be expressed in terms of amoebas. For instance, the approaches presented in [7] and [8], which offer very interesting results, can be implemented by amoebas with the appropriate distance, which thus inherit all the possibilities available to the ones described in this paper. These various approaches show the viability and the vitality of the amoeba framework, as well as its applicability to many fields of research.

## References

- [1] 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.
- [2] J. Serra et al. *Mathematical Morphology for Boolean Lattices*, vol. 2, Academic Press, London, 1988, Chapter 2, pp. 37–46.
- [3] B.V. Cherkassky, A.V. Goldberg, Heap-on-Top Priority Queues, TR 96-042, NEC Research Institute, Princeton, NJ, 1996.
- [4] A. Brodnik et al., Worst case constant time priority queue, *Symposium on Discrete Algorithms* (2001) 523–528.
- [5] C. Tomasi, R. Manduchi, Bilateral filtering for gray and color images, in: *Proceedings of IEEE International Conference on Computer Vision*, Bombay, India, 1998.
- [6] F. Catté et al., Image selective smoothing and edge detection by nonlinear diffusion, *SIAM Journal of Numerical Analysis* 29 (1) (1992) 182–193.
- [7] U.M. Braga-Neto, Alternating sequential filters by adaptive-neighborhood structuring functions, in: *Proceedings of International Symposium on Mathematical Morphology ISMM'96*, 1996, pp. 139–146.
- [8] J. Debayle, J.-C. Pinoli, Multiscale image filtering and segmentation by means of adaptive neighborhood mathematical morphology, in: *Proceedings of IEEE International Conference on Image Processing*, Genoa, Italy, 2005.
- [9] M. Boehm, Contribution à l'amélioration du rendu volumique de données médicales 3D, PhD Thesis, Centre de Morphologie Mathématique, École des Mines de Paris, 2004.
- [10] R. Lerallut, E. Decencière, F. Meyer, Image processing using morphological amoebas, in: *Proceedings of the 5th International Symposium on Mathematical Morphology*, Kluwer Academic Publishers, Dordrecht, 2005.
- [11] J.G. Verly, R.L. Delanoy, Adaptive mathematical morphology for range imagery, *IEEE Transactions on Image Processing* 2 (2) (1993) 272–275.