

Mask Connectivity by Viscous Closings: Linking Merging Galaxies without Merging Double Stars

Ugo Moschini¹, Scott C. Trager² and Michael H. F. Wilkinson¹

¹ Johann Bernoulli Institute, University of Groningen, Groningen, The Netherlands

² Kapteyn Astronomical Institute, University of Groningen, Groningen, The Netherlands

Abstract. Second-generation connectivity opened the path to the use of mask images to freely define connectivity among the image components. In theory, any image could be treated as a mask image that defines a certain connectivity. This creates a new problem in terms of which image to use. In this paper, clustering masks suitable for the analysis of astronomical images are discussed. The connectivity defined by such masks must be capable of preserving faint structures like the filaments that link merging galaxies while separating neighboring stars. In this way, the actual morphology of the objects of interest is kept. This is useful for proper segmentation. We show that viscous mathematical morphology operators have a superior performance and create appropriate connectivity masks that can deal with the characteristic features of astronomical images.

1 Introduction

Second-generation connectivity [1–3] is an effective generalization of the notion of connectivity. It models directly the perceptual clustering or partitioning that is automatically made by a human observer while analysing an image or a scene. Clustering lets objects close to each other in morphological terms to be considered as a single entity, whereas partitioning can split objects that are barely connected. Algorithms that can compute attribute filters using second-generation connectivity based on Max-Trees have been already presented, providing both sequential and parallel implementations [3, 4] for shared-memory machines. Thus, second-generation connectivity is an enabling technology in real applications. In principle, any image can act as a mask that defines a connectivity, independently from how it is computed. Given this lack of restriction, the problem is to evaluate which masks work better according to the task at hand.

An important task in astronomy is to distinguish stars from galaxies and then to classify the way in which the galaxies interact. It would be desirable to achieve a clear separation of those objects that a human observer would immediately identify as separate entities. Fig. 1 shows two examples of merging galaxies connected by a filament. In such examples, nearby or double stars are visually interpreted as separate stars while merging galaxies connected by a filament

are identified as two linked objects. The state-of-the-art software used in the astronomical community used for the detection, segmentation, and classification of objects in astronomical images is called SExtractor [5]. SExtractor is based on an estimation of the background grey levels to establish a threshold above which an object is considered as such. Thresholding often removes not only noise in the background but also important, if faint, structures like the filaments between interacting galaxies. SExtractor would benefit from second-generation connectivity, exploiting the notion of connected components of an image rather than just cutting at a certain grey level threshold. The focus of this paper is on finding connectivity masks that can be used specifically in the analysis of astronomical images, moving from standard to viscous mathematical morphology operators to generate them.

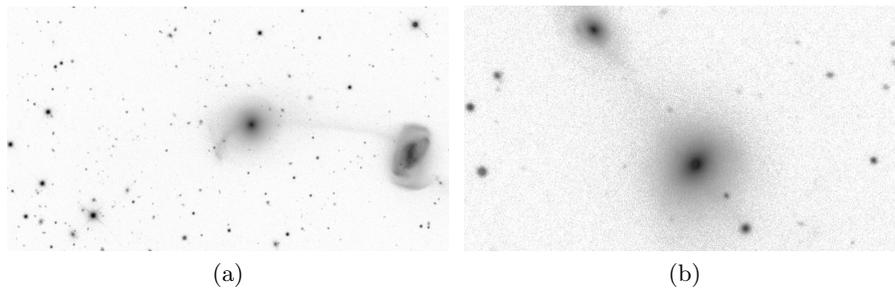


Fig. 1: Filaments in merging galaxies; images taken from the Sloan Digital Sky Survey.

Filaments connecting merging galaxies such as those in Figure 1 reveal information about the interaction between two galaxies. A filament is made of bright speckles and faint elongated structures. These components must be treated as a single connected entity and they should not be split up in smaller components, which individually cannot describe the structure of the filament or its presence. Clustering together the components of a filament, considering it a connected component, can be achieved with connectivity masks generated by extensive operators. The main problem is that, while we want to group together faint and diffuse structures broken up by noise, we want to keep closely-spaced bright structures, which have a far better signal-to-noise ratio, separated. We need operators which adapt the degree of clustering to the local image intensity. Viscous operators [6–8] seem ideally suited for the job.

The paper starts with the explanation of second-generation connectivity. Standard and viscous morphological operators are then presented. The next sections are dedicated to the comparison of suitable masks for astronomical purposes, showing how viscous morphological operators help in generating better masks. We finish with conclusions and a discussion of future work.

2 Second-generation connectivity

Let a binary image X be represented as a subset of some universal set V . Connected filtering operates on the partition of V generated by the *connected components* of X . In practice, V is often the discrete grid \mathbb{Z}^n , with $n = 2$ as the most common case. Given this grid-based representation of the image domain, it is natural to use adjacency relations to define connectivity. Therefore, in image processing the standard form of connectivity is graph-based path connectivity, in which a set is called connected if there exists a path made of adjacent (foreground) points within the set that connects every pair of its points. To facilitate this, we can extend our representation of a binary image as a subset of V to that of a subgraph of the undirected graph (V, E) , with $E \subseteq V \times V$ the set of edges.

A connected component is a connected set of maximal extent in the image. More explicitly, a connected component to which some point $x \in X$ belongs is the union of all the connected sets $C \subseteq X$ containing element x . Connected components are retrieved by a family of connected openings denoted as $\Gamma_x(X)$, $x \in V$.

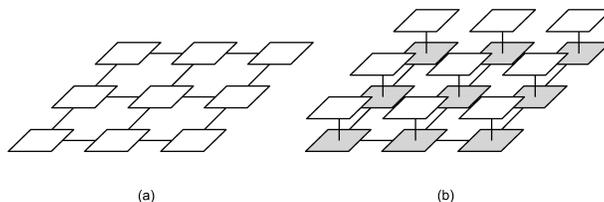


Fig. 2: Mask connectivity seen as modified path connection in a graph: (a) regular 4-connectivity; (b) mask connectivity. White nodes are original pixels, grey nodes mask pixels.

The concept of second-generation connectivity requires the creation of *connectivity masks* that define a connectivity over an image. The easiest way to visualize this is that we change the graph on which we work from (V, E) to (V', E') , as shown in Figure 2. We extend the set of vertices of the graph by adding a separate plane containing the mask image, and thus $V' = V \cup M$, where M is the set of mask-image pixels. We also remove the edges linking the original pixels together. The mask image pixels (shown in grey) are connected to each other in a conventional way. Furthermore, each original pixel is linked by an edge to the corresponding mask pixel. This means the set of edges $E' \subseteq M \times (V \cup M)$ and is disjoint from the original set of edges E . Thus, a pair of pixels x and y in the original image is connected at grey level h if there is a path $\{x_0, x_1, \dots, x_n\} \subset M$ with $(x_0, x), (x_n, y) \in E'$ and such that $f(x_i) \geq h$, for all x_i . Thus, the paths through the mask dictate the connectivity. When computing any properties of the connected components, the mask pixels are ignored, and only the original pixels contribute [9].

There is a simple relationship between the standard connectivity opening Γ_x and the mask-based one Γ_x^M ,

$$\Gamma_x^M(X) = \begin{cases} \Gamma_x(M) \cap X & \text{if } x \in X \cap M & (1a) \\ \{x\} & \text{if } x \in X \setminus M & (1b) \\ \emptyset & \text{otherwise} & (1c) \end{cases}$$

where M is the binary connectivity mask. A connectivity mask suitable for clustering can be created through any desired extensive operator. For a given mask M , the second-generation connected opening defined by $\Gamma_x^M(X)$, with X the image domain, returns as connected sets the intersection of the connected components in the mask image with the image itself. Therefore, the components in X are considered connected on the basis of the connectivity in the mask M rather than of the connectivity in X .

Using this approach, the topological aspects of the components of the mask image are exploited: the underlying connectivity is based on those and at the same time the shape of the original components is preserved. Components that are not connected in the original image can be connected through the underlying mask. For example, if the mask is computed from a dilation of the original, the concept of perceptual groups of image objects are perceived as clusters of connected components if the distance between them is lower than a given threshold, defined by the size of the structuring element. An example of the effect of a clustering mask is shown in Fig. 3 for the globular cluster Messier 13.

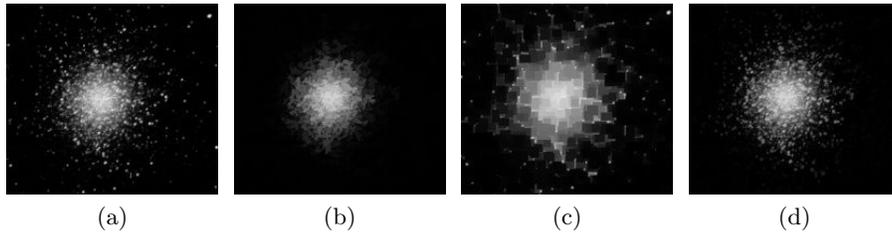


Fig. 3: Separating globular cluster Messier 13 from isolated stars by an area thinning preserving connected components with area between 50 pixels and a quarter of the image size: (a) the original image; (b) the filtered image using no clustering mask; (c) the mask generated through a standard closing; (d) the filtered image using (c) as a connectivity mask.

3 Standard and Viscous operators

The goal of viscous mathematical morphology [6–8] is to adapt the activity of a structuring element to the local intensity information of an image. Morphological

filtering based on standard openings and closing applies the same transformation for all the image level sets. It is desirable to have different smoothing for different grey values, for example in situations in which it is known that components with high grey value do not require the same smoothing as components with lower grey value. A gray level image f can be thresholded into level sets $X_h = \{x \in E : f(x) \geq h\}$ where h represents the intensity value, ranging from the minimum to the maximum possible intensity value. The image can be reconstructed from all the threshold sets by taking their supremum $\bigvee_h h \cdot \chi_h(f)$. The *characteristic function* χ_h equals 1 if $x \in X_h$, else $\chi_h(x) = 0$. Standard dilation δ_λ and erosion ϵ_λ of an image f by a structuring element of size λ process all the threshold sets with the same structuring element and commute with grey scale thresholding:

$$\delta_\lambda(f) = \bigvee_h h \delta_\lambda[\chi_h(f)] \quad \epsilon_\lambda(f) = \bigvee_h h \epsilon_\lambda[\chi_h(f)] \quad (2)$$

In viscous mathematical morphology, the size of the structuring element is not a fixed value λ . It is a function $\lambda(h)$ defined for every intensity value h :

$$\delta_\lambda^{visc}(f) = \bigvee_h h \delta_{\lambda(h)}[\chi_h(f)] \quad \epsilon_\lambda^{visc}(f) = \bigvee_h h \epsilon_{\lambda(h)}[\chi_h(f)] \quad (3)$$

The function $\lambda(h)$ has been studied in two variants in previous work from Vachier and Meyer [7]. They define two functions to model two different situations. In the first case, regions of low intensity are dilated while points of higher intensity are left unchanged, and vice-versa in the second case. The function proposed for the former case was $\lambda(h) = \lambda_m(M - h)/M$ and $\lambda(h) = \lambda_m h/M$ in the latter, where λ_m is the maximum possible size of the structuring element and $[0, M]$ is the intensity range of the image. For the purposes of our work we are interested in the first case that let faint objects be dilated more, thus increasing the likelihood of identifying filaments. The equation that has been used in our work is $\Lambda(h)$, for $h \geq 1$:

$$\Lambda(h) = \lambda_m \frac{1}{h^p} \quad (4)$$

Figure 4 illustrates $\Lambda(h)$ for various p . Experiments were performed with different values for the parameter p , ranging in the interval from 0.2 to 0.5 in which good results were obtained. Since equation 4 must ideally create the biggest dilation for the lowest intensity value actually present in the image, the equation is rewritten as $\Lambda(h) = \lambda_m \frac{1}{(h - \min(f))^p}$, where $\min(f)$ is the lowest intensity value in the input image f , as a simple bias correction.

The dilation δ_Λ and erosion ϵ_Λ derived from equation 4 define an adjunction and produce openings γ^Λ and closings ϕ^Λ :

$$\delta_\Lambda(f) = \bigvee_h h \delta_\Lambda[\chi_h(f)] \quad \epsilon_\Lambda(f) = \bigvee_h h \epsilon_\Lambda[\chi_h(f)] \quad (5)$$

and

$$\gamma^\Lambda = \bigvee_h h \cdot \gamma_\Lambda \chi_h \quad \phi^\Lambda = \bigvee_h h \cdot \phi_\Lambda \chi_h \quad (6)$$

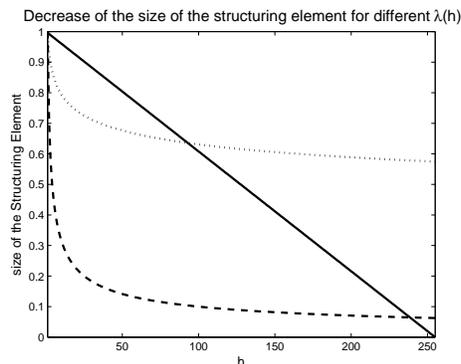


Fig. 4: Structuring element radius as a function of grey level h . Solid line shows the equation studied by Vachier and Meyer; dotted and dashed lines show equations $\Lambda(h)$, with $p = 0.1$ and $p = 0.5$ respectively.

where $\gamma_j = \delta_j \epsilon_j$ and $\phi_j = \epsilon_j \delta_j$ are the standard opening and closing by a structuring element of size j . The closing ϕ^A can be decomposed in the corresponding dilation and erosion:

$$\delta_A^{visc}(f) = \bigvee_h h \delta_{\Lambda(h)}[\chi_h(f)] \quad \epsilon_A^{visc}(f) = \bigvee_h h \epsilon_{\Lambda(h)}[\chi_h(f)]. \quad (7)$$

This means that a connectivity mask suitable for clustering can be obtained using a viscous closing implemented as $\phi^A = \epsilon_A^{visc} \delta_A^{visc}$, allowing for the preservation of high intensities while dilating lower intensities.

4 Experiments and Discussion

The experiments aim to evaluate viscous operators in the generation of connectivity masks suitable to be used for astronomical images. To identify a faint, granular, galactic filament, clustering with a fairly large structuring element is required at low grey levels. On the other hand, bright objects like stars do not need to be processed by any underlying connectivity mask since they should be identified as separate objects: little or no clustering is required in this case.

Reconstruction from markers [10] is used here to determine at which level two objects are clustered together, being that a sign that any structure linking the two objects is preserved through the mask. A measure in percentage that quantifies the clustering will be given. We placed a δ -function marker in the core of one of the two galaxies undergoing merger, or alternatively, in one of two closely-spaced stars. The basic idea is shown in Figure 5. If no mask is used, reconstruction restores the adjacent source at the level of the saddle point between the two sources. When a mask is obtained from an extensive operator, it may restore more of the second source, with a maximum at the lower of the two peak levels.

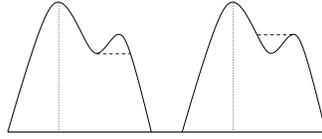


Fig. 5: Measuring level of connectivity in two closely-spaced sources. The solid curves indicate brightness profiles of two sources. The vertical fine dotted lines indicate the δ -function markers. On the left, the thick dashed line indicates the reconstruction from the marker using regular connectivity. On the right, the thick dashed line shows the worst-case merging height for clustering connections

4.1 Evaluation of the experiments

A percentage value is used to quantify how much the adjacent source is restored above the saddle point. A zero percent value is assigned to the intensity value of the adjacent component after a reconstruction in absence of clustering. It is not possible to reconstruct such a component at an intensity lower than the intensity value of the saddle point. A one hundred percent value is assigned to the intensity value of the adjacent component after a reconstruction that used a mask obtained from a standard dilation with an S.E. diameter of 20 pixels, corresponding to the maximum intensity that can be restored. Experiments were conducted varying the diameter of the structuring element used for the creation of the masks both for standard and viscous operators. The structuring element used in all the morphological operations has a circular shape, with a diameter ranging from 5 to 30 pixels. Viscous masks have been produced changing also the value of the parameter p that shapes the rate of decrease of the diameter length of the structuring element. The mask images were generated by standard dilations, closings and viscous closings. Ten 16-bit integer images taken in the R-band images of merging galaxies connected by filaments were selected from the Sloan Digital Sky Survey Data Release 7 [11] archives. Code was implemented to deal with 16-bit images. From the equations 3, a naive implementation of a viscous mask is obtained by dilating or eroding all the level sets and taking the supremum. To avoid looping over all the level sets in the intensity range, our implementation works by looping instead over all the possible different sizes of the structuring element. We use an adaptation of the multi-erode algorithm from [12] to perform this efficiently. The time complexity scales now with the number of possible different diameters that the structuring element can have. Hence, the code would work also on floating point images where the number of level sets is in principle infinite. Wall-clock time needed to generate a mask reduces by two orders of magnitude. For the max-tree computation the algorithm for high dynamic range images from [13] was used.

In every image, three positions for the markers were identified. A first marker was placed on the rightmost galaxy between the two connected by a filament in every image, the other two markers were placed in locations where two objects close to each other but not interacting were identified; either two stellar objects,

or two galaxies not showing signs of any interaction between them. The percentage of the reconstructed adjacent component quantifies how much these objects are considered bound together. The box diagrams in Figure 6 summarise the experiments. Figure 6a shows how the masks perform in clustering desired objects together, while Figure 6b shows how they perform in keeping separate the objects that do not interact. The numbers on the x-axis indicate the twenty-four connectivities that have been tested as listed in Table 1. Index 1 shows the case where no mask image is used to define an underlying connectivity: in fact, it shows a value of zero percent meaning no clustering above the level of the saddle point. The mask generated by a standard dilation with a diameter of 20 pixels is assigned to the index 4 and it has value of one hundred per cent. Parameter p used in equation 4 takes the value of 0.2, 0.3, 0.4 and 0.5 for every diameter of the structuring element tested. Percentage values after the reconstruction from masks generated through standard operators are shown in gray boxes, through viscous operators in black boxes.

Table 1: The table shows the mask images used in the diagrams in Figure 6.

Index	Operation	S.E. Diameter	p	Index	Operation	S.E. Diameter	p
1	-	-	-	13	Viscous Closing	30	0.2
2	Standard Dilation	5	-	14	Viscous Closing	30	0.3
3	Standard Dilation	10	-	15	Viscous Closing	30	0.4
4	Standard Dilation	20	-	16	Viscous Closing	30	0.5
5	Standard Closing	5	-	17	Viscous Closing	40	0.2
6	Standard Closing	10	-	18	Viscous Closing	40	0.3
7	Standard Closing	20	-	19	Viscous Closing	40	0.4
8	Standard Closing	30	-	20	Viscous Closing	40	0.5
9	Viscous Closing	20	0.2	21	Viscous Closing	50	0.2
10	Viscous Closing	20	0.3	22	Viscous Closing	50	0.3
11	Viscous Closing	20	0.4	23	Viscous Closing	50	0.4
12	Viscous Closing	20	0.5	24	Viscous Closing	50	0.5

The ideal situation expected for a given mask would show a high percentage for the corresponding box in Fig. 6a and a low percentage for the one in Fig. 6b. This means that the object that should be considered linked to another is restored more than an other object in the same image not linked to anything else. The best results observed in the experiments are obtained with the connectivity defined by a viscous mask generated by a structuring element whose diameter length is 50 pixels and parameter p equals to 0.4 (at index 23) and by a structuring element whose diameter length is 40 and p equals to 0.3 and 0.4 (at indexes 18 and 19). For larger values of p , too much clustering appeared, so these results are not shown in the diagram. Previous experiments conducted on 8-bit images taken from the same Sloan survey and derived from the 16-bit images showed that $p = 0.5$, meaning that the diameter of the structuring element decreases as the square root of the image intensity, turned out to be the best value. This

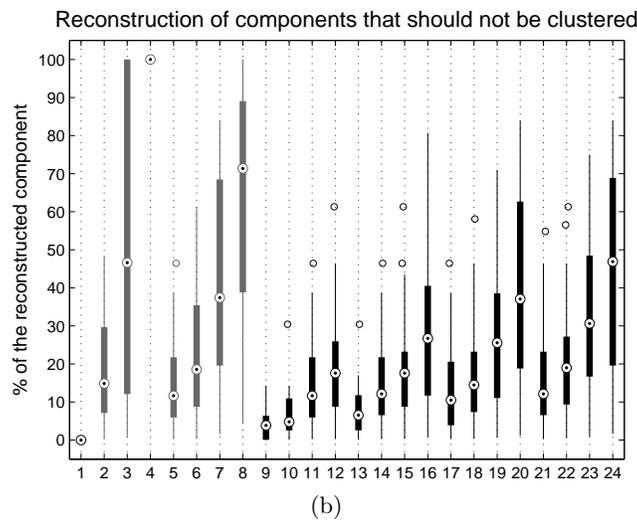
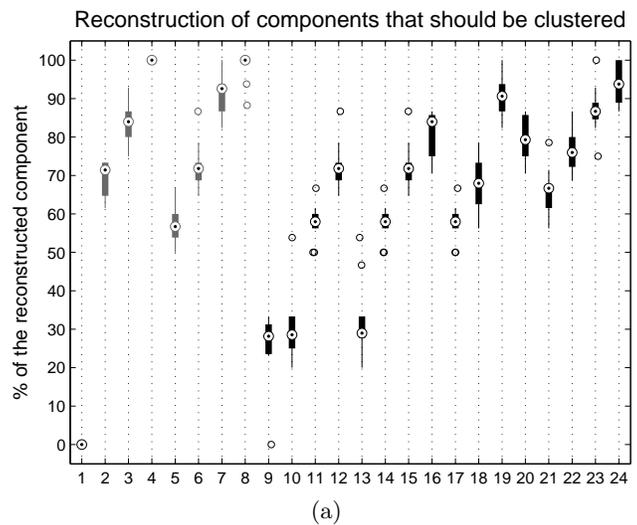


Fig. 6: The diagrams show the percentage value of the reconstructed level of the adjacent component. At index 1, no mask image is used at all. Indices from 2 to 8 illustrate the case of masks generated through standard operators. Indices from 9 to 24 illustrate the case of mask images obtained through viscous closings. The settings of every mask are given in Table 1.

difference is probably due to some dynamic range compression in the conversion to 8-bit per pixel. Using the actual 16-bit images, features like the filaments appear to be fainter with respect to the main objects than in the 8-bit images.

This means that more clustering is necessary in this case to preserve the fainter components, as reflected in p giving the best results for values lower than 0.5.

Regarding the behaviour of the masks generated by standard operators, they are obviously able to cluster merging galaxies at the same level as or higher than the viscous masks. The drawback is that they also cluster unwanted objects together. This is reflected in the diagram in the positions of indexes from 3 to 8. Note that at index 2 the percentages show a decent reconstruction of the galactic component while keeping other objects separate. However, it must be considered that a dilation with a structuring element with diameter 5, shown at index 3, also clusters the unwanted object. If masks generated from standard operators are used, the reconstruction looks very sensitive to the diameter: an increase of the length from 5 to 10 pixels modifies the output radically. In the viscous case, closer percentages are obtained when similar settings are used, like for example at indices 21, 22 and 23.

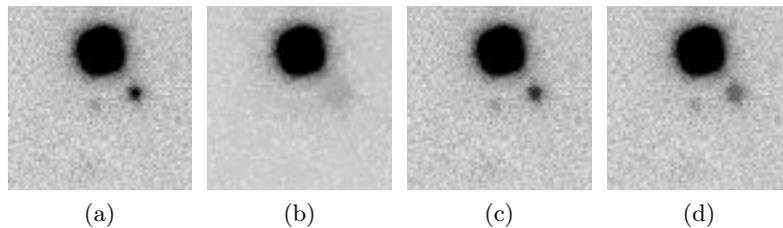


Fig. 7: (a) is a region of the original image; (b) shows the reconstruction when no mask is used; (c) and (d) show the reconstruction when a standard and a viscous closing were used as masks, respectively. The marker used for reconstruction is placed on the big star.

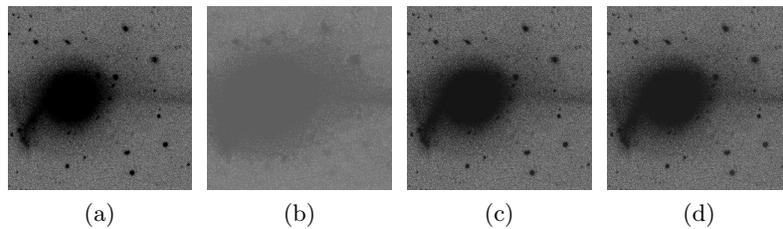


Fig. 8: (a) is a region of the original image; (b) shows the reconstruction when no mask is used; (c) and (d) show the reconstruction when a standard and viscous closing were used as masks, respectively. The marker used for reconstruction is placed on a galaxy on the right.

4.2 A visual evaluation

Figure 7 and Figure 8 represent two regions of Figure 1a, depicting a very critical situation composed of two stars very close each other and two merging galaxies. Image intensities have been clipped for the purpose of visualization. The markers were placed on the brighter star and on the rightmost galaxy. Two connectivity masks are compared here: the first is generated by a standard closing and the other by a viscous closing. Figure 7b shows that with no second-generation connectivity, the star component is weakly reconstructed: the small star is reconstructed up to the intensity value of the saddle point, or 0% in our graphs.

When a mask generated by a standard closing is applied, as in Figure 7c, the star component is reconstructed more than in the case of a mask generated by a viscous closing (73% vs. 38%). On the other hand, as in Figure 8, the galactic component reached thanks to a better preserved filament is reconstructed at the same level ($\sim 80\%$) both for a mask generated by a standard closing and by a viscous closing. Generally speaking, for the viscous mask case, it is then observed that stellar components that are not interacting are clustered less, while at the same time the interacting galaxies are reconstructed to a level comparable to that obtained in the standard closing or even higher. This depends on the maximum size of the structuring element and on the value of p used.

5 Conclusions

In this paper, an evaluation of connectivity masks generated through viscous mathematical morphology operators has been presented. These masks work well in keeping nearby stars separate, while clustering together the faint filaments typically linking merging galaxies. Viscous operators clearly work better than the regular closing used in this experiment. The tuning of the parameters driving the size reduction will be subject of further studies. More experiments are currently being conducted to determine the optimal settings for both this parameter and the maximum size of the structuring element in a wider range of test cases. While more fine-tuning is needed, it is evident from this initial work that a better object separation is already obtained by viscous masks. Other clustering masks are also under investigation.

Mask-Edge connectivity [9] looks to be a promising tool to generate mask images that do not allow clustering in presence of sharp edges. K -flat filters [14] might be used as well to preserve faint components up to k intensity values more, but only when they show certain characteristic or shape features, such as elongation, for example, that can suggest that it is a filament. Moving away from mathematical morphology, first experiments have been started with PDE-based methods relying on Perona-Malik [15] diffusion, or variants that are driven more by image intensities than by just image gradients. These might be suitable to be used in astronomical imagery in some cases. Softwares performing object segmentation like SExtractor would benefit from the use of connectivity mask for a more reliable deblending of the objects of interest.

Acknowledgement

This work was funded by the Netherlands Organisation for Scientific Research (NWO) under project number 612.001.110

References

1. Serra, J.: Connectivity on complete lattices. *J. Math. Imag. Vis.* **9**(3) (November 1998) 231–251
2. Braga-Neto, U., Goutsias, J.: Connectivity on complete lattices: new results. *Comput. Vis. Image Underst.* **85**(1) (January 2002) 22–53
3. Ouzounis, G.K., Wilkinson, M.H.F.: Mask-based second generation connectivity and attribute filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(6) (2007) 990–1004
4. Ouzounis, G.K., Wilkinson, M.H.F.: A parallel implementation of the dual-input max-tree algorithm for attribute filtering. In Banon, G.J., Barrera, J., Braga-Neto, U.M., Hirata, N.S., eds.: *Proc. Int. Symp. Math. Morphology (ISMM) 2007*. Volume 1., São José dos Campos, Universidade de São Paulo (USP), Instituto Nacional de Pesquisas Espaciais (INPE) (Oct. 2007) 449–460
5. Bertin, E., Arnouts, B.: Sextractor: software for source extraction. *Astronomy and Astrophysics, Suppl. Ser* **117** (1996) 393–404
6. Vachier, C., Meyer, F.: The viscous watershed transform. *J. Math. Imag. Vis.* **22**(2-3) (May 2005) 251–267
7. Vachier, C., Meyer, F.: News from viscousland. In Banon, G.J., Barrera, J., Braga-Neto, U.M., Hirata, N.S., eds.: *Proc. Int. Symp. Math. Morphology (ISMM) 2007*. Volume 1., São José dos Campos, Universidade de São Paulo (USP), Instituto Nacional de Pesquisas Espaciais (INPE) (Oct. 2007) 189–200
8. Maragos, P., Vachier, C.: A pde formulation for viscous morphological operators with extensions to intensity-adaptive operators. In: *Proc. Int. Conf. Image Proc. 2008*, IEEE (2008) 2200–2203
9. Oosterbroek, J., Wilkinson, M.H.F.: Mask-edge connectivity: Theory, computation, and application to historical document analysis. In: *icpr2012*. (2012)
10. Vincent, L.: Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms. *IEEE Transactions on Image Processing* **2** (1993) 176–201
11. Abazajian, K.N., Adelman-McCarthy, J.K., Agüeros, M.A., Allam, S.S., Allende Prieto, C., An, D., Anderson, K.S.J., Anderson, S.F., Annis, J., Bahcall, N.A., et al.: The Seventh Data Release of the Sloan Digital Sky Survey. *The Astrophysical Journal Supplement Series* **182** (June 2009) 543–558
12. Urbach, E.R., Wilkinson, M.H.F.: Efficient 2-D gray-scale morphological transformations with arbitrary flat structuring elements. *IEEE Trans. Image Proc.* **17** (2008) 1–8
13. Wilkinson, M.H.F.: A fast component-tree algorithm for high dynamic-range images and second generation connectivity. In: *Proc. Int. Conf. Image Proc. 2011*. (2011) 1041–1044
14. Ouzounis, G.K., Wilkinson, M.H.: Hyperconnected attribute filters based on k-flat zones. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **33**(2) (2011) 224–239
15. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **12**(7) (1990) 629–639