



HAL
open science

Simultaneous Segmentation and Filtering Via Reduced Graph Cuts

Nicolas Lermé, François Malgouyres

► **To cite this version:**

Nicolas Lermé, François Malgouyres. Simultaneous Segmentation and Filtering Via Reduced Graph Cuts. 2012. hal-00624093v2

HAL Id: hal-00624093

<https://hal.science/hal-00624093v2>

Preprint submitted on 22 Nov 2011 (v2), last revised 19 Jun 2012 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Simultaneous Segmentation And Filtering Via Reduced Graph Cuts

Anonymous CVPR submission

Paper ID 984

Abstract

Graph cuts have now become a cornerstone in computer vision for efficiently solving numerous labeling problems. However, their large memory consumption make it hard to solve large-scale problems even when restricted to binary variables. Except some exact methods [14, 8, 4], the heuristics [10, 11, 13, 6] generally fail to fully capture details. In this paper, we first review a band-based strategy for reducing graphs in image segmentation [9]. It provides small graphs while accurately preserving thin structures but do not offer low memory usage when the amount of regularization is large. This is, for instance, the case when images are corrupted by an impulsive noise. In this paper, we overcome this situation by embedding a new parameter in [9] to both further reducing graphs and filtering the segmentation. This parameter appears to be generally less sensitive to noise variations and offers a good robustness against noise. We also provide an empirical way to automatically tune it and illustrate its behavior through some experiments.

Keywords: graph cuts, reduction, segmentation, filtering.

1. Motivation and scope

Graph cuts have become increasingly popular due to their ability to efficiently compute the Maximum A Posteriori of Markov Random Fields (MRF). This popularity is notably driven by the introduction of a fast maximum-flow (max-flow) algorithm [3] making near real-time performance possible for solving numerous labeling problems.

In parallel, technological advances in image acquisition have both increased the amount and the diversity of data to process. As an illustration, in the satellite SPOT-5 launched by Arianespace in 2002, high resolution sensors can capture multispectral and panchromatic images of about 1GB.

Processing this type of data amounts to solve large scale optimization problems. In the image segmentation context, almost all graph cuts-based methods are impractical to solve such problems due to the memory requirements. To overcome this situation, some amount of work has been done in this direction and a number of heuristics [10, 11, 13, 6] and

exact methods [8, 4, 14] have been proposed.

To our best knowledge, this problem seems to be first addressed in [10] where the underlying graph is built upon a pre-segmentation. Although this approach greatly reduce the computational burden of graph cuts, the results strongly depend on the algorithm used for computing the pre-segmentation. Also, better results are obtained when over-segmentation occurs, losing in this way the main benefit of such a reduction.

Others have also reported band-based methods [11, 13, 6]. A low-resolution of the image is first segmented and the solution is propagated to the finer level by only building the graph in a narrow band surrounding the interpolated foreground/background interface at that resolution. While such an approach drastically reduce time and memory consumption, it is limited to segment roundish objects. This problem is notably reduced in [13] but still present for low-contrasted details. In [6], finer bands are obtained using an uncertainty measure associated to each pixel.

Exact methods have been also investigated [8, 4, 14]. In [8], binary energy functions are minimized for the shape fitting problem with graph cuts in a narrow band while ensuring the optimality on the solution. One makes a band evolve around the object to delineate by expanding it when the minimum-cut (min-cut) touches its boundary. This process is iterated until the band no longer evolves. Although the algorithm generally converges in few iterations toward the optimal solution, an initialization is required and no bound on the band size is given.

A parallel max-flow algorithm yielding a near-linear speedup with the number of processors is described in [4]. Nevertheless, the algorithm is relatively sensitive to the available amount of physical memory and remains less efficient on small graphs.

The approach used in [14] is different: instead of reducing the graphs, the problem is decomposed into optimizable sub-problems, solved independently and updated according to the results of the adjacent problems. This process is iterated until convergence and optimality is guaranteed by Lagrangian decomposition.

Finally, another band-based method was proposed for re-

ducing graphs in binary image segmentation [9]. The graph is progressively built by only adding nodes which locally satisfy a condition. In the manner of [13, 6], the graph nodes are typically located in a narrow band surrounding the object edges to segment. Empirically, the authors show in [9] that the solutions obtained with and without reduction are identical and the time for reducing the graph is even compensated by the time for computing the min-cut in the reduced graph.

The rest of this paper is organized as follows. First, we review the graph cuts framework in Section 2 as well as the strategy of [9] in Section 3 for reducing graphs. Afterwards, a new parameter is introduced in Section 4 for further reducing the graphs and removing small undesired segments in the segmentation when the image is corrupted by a noise behaving like an impulsive noise. The sensitivity of this parameter as well as its robustness against noise are also evaluated through experiments for segmenting images.

2. Preliminaries

Consider an image $I : \mathcal{P} \subset \mathbb{Z}^d \rightarrow [0, 1]^c$ as a function, mapping each pixel $p \in \mathcal{P}$ to a vector $I_p \in [0, 1]^c$ ¹. We define a binary segmentation as an application u affecting to each pixel $p \in \mathcal{P}$ either 0 (background) or 1 (object) and we write $u \in \{0, 1\}^{\mathcal{P}}$. A popular strategy to segment I is to minimize a MRF of the form [2]:

$$E(u) = \beta \cdot \sum_{p \in \mathcal{P}} E_p(u_p) + \sum_{(p,q) \in \mathcal{N}} E_{p,q}(u_p, u_q), \quad (1)$$

among $u \in \{0, 1\}^{\mathcal{P}}$ and for a fixed $\beta \in \mathbb{R}^+$. The neighborhood system $\mathcal{N} \subset \mathcal{P}^2$ is a subset of all pixel pairs $(p, q) \in \mathcal{P}^2$. In the sequel, "connectivity 0" will denote 4 and 6 neighbors in 2D and 3D images while "connectivity 1" will denote 8 and 26 neighbors for the same images. In (1), the data term $E_p(\cdot)$ is defined as the negative log-likelihood of a label being assigned to pixel p and is computed from its color and the appearance models of the object and background seeds²:

$$\begin{cases} E_p(1) = -\log \mathbb{P}(I_p | p \in \mathcal{O}) \\ E_p(0) = -\log \mathbb{P}(I_p | p \in \mathcal{B}) \end{cases}$$

For any pair $(p, q) \in \mathcal{N}$, the corresponding smoothness term in (1) is defined as a contrast-sensitive Ising model:

$$E_{p,q}(u_p, u_q) = \begin{cases} 0 & \text{if } u_p = u_q, \\ \frac{1}{\|p-q\|_2} \exp\left(-\frac{\|I_p - I_q\|_2^2}{2\sigma^2}\right) & \text{otherwise,} \end{cases}$$

where $\|\cdot\|_2$ is the Euclidean norm (either in \mathbb{R}^d or \mathbb{R}^c). When the smoothness terms are submodular [7], the minimizer of (1) can be efficiently obtained by computing a

¹Usually, \mathcal{P} corresponds to a rectangle.

²In this setting, the distributions are estimated using a Gaussian Mixtures Model. The number of gaussians is automatically computed using a Minimum Description Length criteria [1].

min-cut in a weighted digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, c)$ with a set of nodes $\mathcal{V} = \mathcal{P} \cup \{s, t\}$, a set of edges $\mathcal{E} \subset \mathcal{V}^2$ and capacities $c : \mathcal{E} \rightarrow \mathbb{R}^+$. The terminal nodes s and t are called the source and the sink, respectively. The set of edges \mathcal{E} is split into two disjoint sets \mathcal{E}_n and \mathcal{E}_t denoting respectively n-links (edges linking two nodes of \mathcal{P}) and t-links (edges linking a node of \mathcal{P} to s or t). Once the min-cut is computed in \mathcal{G} , we set $u_p = 1$ if a node p is connected to the source s and $u_p = 0$ if p is connected to the sink t .

3. Reduction

As said earlier, the memory consumption of graph cuts for segmenting high-resolution data can be very large. As an illustration, the max-flow algorithm of [3] v3.01 used in the experiments of Section 4, allocates $25\#\mathcal{P} + 16\#\mathcal{E}_n$ bytes³. For a fixed amount of RAM, one clearly see that the maximum image size quickly decreases as the dimensionality d of \mathcal{P} increases. As shown in [9], most of the nodes in the graph are however useless during the max-flow computation since they are not traversed by any flow. Ideally, one would like to extract the smallest possible graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ from $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ while keeping the max-flow value f'^* in \mathcal{G}' identical or very close to the max-flow value f^* in \mathcal{G} . This corresponds to an optimization problem which we will not try to solve since the method for determining \mathcal{G} also needs to be (very) fast.

Let us first introduce some terminology before reviewing the method of [9] for building \mathcal{G}' . For the sake of clarity, the same notations are used as in [9]. In accordance with the construction given in [7], we consider (without loss of generality) that a node is connected to at most one terminal

$$(s, p) \in \mathcal{E}_t \Rightarrow (p, t) \notin \mathcal{E}_t, \quad \forall p \in \mathcal{P}.$$

We also summarize t-links capacities by

$$c(p) = c(s, p) - c(p, t), \quad \forall p \in \mathcal{P}.$$

For any $B \subset \mathbb{Z}^d$ and a node $p \in \mathcal{P}$, we denote by B_p the set translation of B at p

$$B_p = \{q + p \mid q \in B\}.$$

For $Z \subset \mathcal{P}$ and $B \subset \mathbb{Z}^d$, we denote by Z_B the dilation of Z by B as

$$Z_B = \{p + q \mid q \in B, p \in Z\} = \bigcup_{p \in Z} B_p.$$

From here, the idea developed in [9] for building \mathcal{G}' is to remove from the nodes of \mathcal{G} any $Z \subset \mathcal{P}$ where all nodes are linked to s (resp. to t) and such that all the flow that might

³The operator ' $\#\$ ' stands for the cardinality of a set.

⁴In practice, B is a square centered at the origin.

get in (resp. out of) the region Z_B does so by traversing its boundary and can be absorbed (resp. provided) by the band $Z_B \setminus Z$. Building such Z is done by testing each individual pixel $p \in \mathcal{P}$. In the manner of [13, 6], the remaining nodes are therefore located in a narrow band surrounding the object edges to segment. In practice, the authors of [9] use a more conservative test by testing each node $p \in Z$ in a square window B of size $(2r + 1)^d$ centered in p :

$$\begin{cases} \text{either} & \left(\forall q \in B_p, c(q) \geq +\delta_r \right), \\ \text{or} & \left(\forall q \in B_p, c(q) \leq -\delta_r \right). \end{cases} \quad (2)$$

where $\delta_r = \frac{P(B)}{(2r+1)^{d-1}}$ and $P(B)$ is defined as

$$P(B) = \max(\#\{(p, q) : p \in B, q \notin B \text{ and } (p, q) \in \mathcal{N}\}, \#\{(q, p) : p \in B, q \notin B \text{ and } (q, p) \in \mathcal{N}\}).$$

In words, for any node $p \in Z$ satisfying the first (resp. second) condition of (2), all its neighbors $q \in B_q$ are only linked to s (resp. t) and the flow that might get in (resp. out) through t-links in $B_p \setminus \{p\}$ suffices to saturate the n-links going out (resp. in) B_p . Thus, p becomes useless and need not to be added to \mathcal{G}' . An algorithm of complexity independent of r is also mentioned in [9] for computing (2). A key point of the method proposed in [9] is that the pixel error between segmentations obtained with and without reduction remains very low.

The experiments presented in [9] confirm the dependence between the size of \mathcal{G}' and model parameters. Indeed, when minimizing (1) by graph cuts, the t-links capacities are all multiplied by β . Thus, it is straightforward to observe that (2) is harder to satisfy as β decreases. In such a situation, we need a larger window radius for decreasing δ_r in order to reduce the size of \mathcal{G}' . This results in wider bands around the object contours. Conversely, when β is large, we can afford a large δ_r and therefore a small window radius. This time, \mathcal{G}' consists of narrow bands around the object contours. However, such an ideal situation cannot be guaranteed when segmenting images corrupted by an impulsive noise since β should not be too large to avoid noise in the solution. In the next section, we embed a new parameter in (2) for filtering the segmentation while keeping β large.

4. Simultaneous segmentation and filtering

A simple way to relax (2) is to allow some nodes in B_p to fail complying the test. The proportion of nodes satisfying (2) can be controlled by a parameter $\eta \in [0, 1]$. As η decreases, the test (2) becomes easier to satisfy since a larger proportion of nodes can be connected to opposite terminals. Embedding η in (2) leads to

$$\begin{cases} \text{either} & \left(\#\{q \in B_p \mid c(q) \geq +\delta_r\} \geq \eta \#B_p \right), \\ \text{or} & \left(\#\{q \in B_p \mid c(q) \leq -\delta_r\} \geq \eta \#B_p \right). \end{cases} \quad (3)$$

4.1. Further reducing graphs

The parameter η can be used for decreasing the memory consumption of graph cuts. The Figure 3 illustrates how far the test (2) can be relaxed for further reducing graphs while getting nearly the same segmentation. In this experiment, the segmentation as well as the reduced graph are shown for segmenting a 2D noisy image. Since the test (3) is easier to satisfy as η decreases, the reduced graph \mathcal{G}' becomes thicker around the object contours.

4.2. Automatic tuning of η

For a fixed window radius, notice first that the value of η must be large enough to not increase the number of components in the reduced graph \mathcal{G}' (see Figure 2). Indeed, below some value (denoted by η_{min}), the reduced graph \mathcal{G}' is split into multiple pieces in areas with high-curvature and the min-cut is no longer ensured of being fully embedded into \mathcal{G}' . This implies that some voxels could be wrongly labeled in the segmentation.

The Figure 1 illustrates a situation where η_{min} can be easily computed with an image consisting of two highly-contrasted areas. Using (3) with a square window of radius r and $\eta = 1$, the reduced graph \mathcal{G}' corresponds to a thin band of size $2r$. An easy under-estimation of η_{min} is obtained by imposing that η_{min} permits to segment these two contrasted areas. In order to do so, we want the test (3) to be false for any pixel p at the boundary between these areas. For such a pixel, we have (assuming e.g. that $c(p) \geq +\delta_r$)

$$\#\{q \in B_p \mid c(q) \geq +\delta_r\} = (r + 1)(2r + 1)^{d-1}.$$

As a consequence, if

$$\eta \leq \frac{(r + 1)(2r + 1)^{d-1}}{(2r + 1)^d},$$

the node p does not belong to the reduced graph \mathcal{G}' . Since we want to avoid the situation, we must therefore have

$$\begin{aligned} \eta &> \frac{(r+1)(2r+1)^{d-1}}{(2r+1)^d} \\ &= 1 - \frac{r}{2r+1} = \eta_{min}. \end{aligned} \quad (4)$$

Remark that (4) does not depend on the dimensionality d of \mathcal{P} . By observing (4), it is straightforward to see that, as the window radius r tends to infinity, the proportion of nodes allowed to be connected to opposite terminals tends to $\frac{1}{2}$. In practice, we also observed that (4) is less accurate in connectivity 0 than in connectivity 1 (see Figure 2).

For a fixed window radius r and a positive amount of noise ξ , one can observe in Figure 3 that there exists a value of the parameter η for which most of the nodes in noisy regions are removed from the graph \mathcal{G} , leading to a diminution of the size of the reduced graph \mathcal{G}' .

The purpose of this paragraph is to identify, from a statistical point of view, a reliable value of the parameter η for which all nodes of \mathcal{P} are very likely to be removed from \mathcal{G} . For a fixed amount of noise ξ in the image I , we therefore want to find an upper bound on η by finding the maximum value of η in such a way that we control the proportion of nodes corresponding to noisy pixels in homogeneous areas.

Consider a noisy constant image I with a noise generated by a Bernoulli distribution of parameter $\xi \in]0, 1[$, corresponding to the amount of noise in I ⁵. The two cases where $\xi = 0$ and $\xi = 1$ are trivial and are not considered in our analysis. Assume now that the graph \mathcal{G} is defined as in Section 3 where the nodes corresponding to noise-free pixels are connected to the sink t with a capacity $c(q) \leq -\delta_r$ and the nodes corresponding to noisy pixels have a capacity $c(q) > -\delta_r$.

First, let X be a discrete random variable counting degraded pixels in a square window B of size $n = (2r + 1)^d$ in the image I . Then, the probability that at least k pixels are corrupted in B is given by

$$\mathbb{P}(X > k) = \sum_{i=k+1}^n \binom{n}{i} \xi^i (1-\xi)^{n-i}, \quad (5)$$

where $\binom{n}{i} = \frac{n!}{i!(n-i)!}$. For a fixed window radius r , it is straightforward to see that (5) is decreasing in k and tends to ξ^n if we impose that $\xi \in]0, 1[$. According to the test (3) and the hypothesis on \mathcal{G} : a node $p \in \mathcal{P}$ can be removed from \mathcal{G} if and only if

$$\#\{q \in B_p \mid c(q) \leq -\delta_r\} = \#\{q \in B_p \mid q \text{ is noise-free}\} \geq \eta n \quad (6)$$

Moreover, we assumed

$$\#\{q \in B_p \mid q \text{ is noise-free}\} \sim (n - X).$$

Therefore, we have

$$\mathbb{P}(p \text{ is not removed}) = \mathbb{P}((n-X) < \eta n) = \mathbb{P}(X > (1-\eta)n).$$

Fixing a proportion $\varepsilon \simeq 0$ of wrongly constructed nodes, we choose

$$\eta^+ = \max \{\eta \in [0, 1] \mid \mathbb{P}(X > (1-\eta)n) \geq \varepsilon\}, \quad (7)$$

Considering the lower bound η_{min} defined in (4), we set

$$\eta_{max} = \max \{\eta_{min}, \eta^+\}. \quad (8)$$

Combining the definitions of the lower and upper bounds (see (4) and (8)), it now becomes easy to get an estimation of the parameter η^* for a fixed window radius by setting

$$\eta^* = \left(\frac{\eta_{min} + \eta_{max}}{2} \right). \quad (9)$$

⁵Simple histogram-based techniques can be for instance used to estimate the amount of noise ξ in the image.

Let us now analyze the joint behavior of the lower and the upper bounds. When the amount of noise ξ is fixed, one can easily observe that the gap $\Delta\eta = (\eta_{max} - \eta_{min})$ grows as the window radius r increases. Indeed, we have previously seen that the lower bound η_{min} tends to $\frac{1}{2}$ as the window radius r increases (see (4)). The previous observation is also due to the fact that the upper bound η_{max} grows as the window radius r increases.

Similarly, when the window radius r is fixed, remark that $\Delta\eta$ decreases when the amount of noise ξ increases. This situation is consistent because η_{min} remains the same but η_{max} tends to $\frac{1}{2}$ since it is more likely that the number of degraded pixels increase in the window B . Notice that increasing the window radius r can compensate the augmentation of the amount of noise only up to $\xi = 0.5$.

4.3. Filtering

The parameter η can also serve to filter the segmentation. This behavior is illustrated in Figure 4 for segmenting a 3D noisy image acquired from a confocal microscope. White spots correspond to cell nuclei in a mouse cerebellum. Observe how far the filtering acts for small values of η : small regions in the reduced graph \mathcal{G}' as well as in the segmentation are progressively removed as η decreases.

The robustness to noise (see Figure 6 and 5) and the sensitivity of the parameter η (see Figure 7) have been also analyzed. Let us describe the experimental procedure for segmenting four grayscale and five color 2D images with an increasing noise level ranging from 4 to 48%. For each image, we compute a reference segmentation on the noise-free image by placing the seeds by hand. We set $\beta = +\infty$ and the σ parameter is automatically estimated as explained in [12]. Then, for each impulsive noise level, we select the segmentation maximizing the Dice Similarity Coefficient (DSC) [5] between the reference image and all segmentations obtained through a fixed range of r and η values. Each segmentation is computed using the same seeds as those used for the computing the reference segmentation. Again, the σ parameter is automatically estimated using [12].

As shown in Figure 6, for an impulsive noise level up to 45%, the parameter η appears to be reasonably robust with a DSC always greater than 94% for all images, except for the image "rice". However, such high and stable noise robustness can only be reached by increasing the amount of seeds (see Figure 6). The reason why the algorithm behaves poorly on the "rice" image is the following. As said earlier, r must be large enough when ξ increases for removing a maximum number of segments due to noise. This implies wider bands in \mathcal{G}' around the object contours. However, the object contours further oscillate as ξ increases. Another reason is due to the proximity of the objects to segment. As an illustration, consider two circles over a uniform background, separated by a distance $d_0 > 0$. We clearly see




Figure 1: Toy example for computing η_{min} .




Figure 5: Qualitative analysis of the robustness to noise for segmenting the images "f117" (left-most column), "black-cat" (left column), "pyramid" (right column) and "flamingo" (right-most column) with a fixed impulsive noise level of 36%. The seeds and the model parameters are the same than those used in Figure 6 (top row).




Figure 6: Quantitative analysis of the robustness to noise for segmenting four 2D grayscale images (top-most curves in the list) and five 2D color images with an impulsive noise level ranging from 4 to 48%.

that the test (3) becomes more and more difficult to satisfy when the window radius r increases. When $(2r + 1) \geq d_0$, the reduced graphs of both circles fuse into one component. This is the case in the image "rice" because this photo consists of small assembled rice grains. Finally, the Figure 7 also illustrates that the parameter η is not very sensitive to the variations of r and η . The DSC does not vary much with respect to noise, except for the image "rice". The latter problem can be explained for the same reasons as before.




Figure 7: Sensitivity of η for segmenting the images in Figure 6 with an impulsive noise level of 36%. The seeds and model parameters are the same than those used in Figure 6.

References

- [1] C. A. Bouman. Cluster: An unsupervised algorithm for modeling Gaussian mixtures. April 1997. 2
- [2] Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In *ICCV*, volume 1, pages 105–112, 2001. 2
- [3] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on PAMI*, 26(9):1124–1137, 2004. 1, 2
- [4] A. Delong and Y. Boykov. A scalable graph-cut algorithm for N-D grids. In *CVPR*, pages 1–8, 2008. 1
- [5] L. Dice. Measure of the amount of ecological association between species. *Ecology*, 26(3):297–302, 1945. 4
- [6] P. Kohli, V. Lempitsky, and C. Rother. Uncertainty driven multi-scale energy optimization. In *DAGM*, pages 242–251, 2010. 1, 2, 3
- [7] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions on PAMI*, 26(2):147–159, 2004. 2
- [8] V. Lempitsky and Y. Boykov. Global optimization for shape fitting. In *CVPR*, pages 1–8, 2007. 1
- [9] N. Lerm, F. Malgouyres, and L. Ltcart. Reducing graphs in graph cut segmentation. In *ICIP*, pages 3045–3048, 2010. 1, 2, 3
- [10] Y. Li, J. Sun, C. Tang, and H. Shum. Lazy Snapping. *ACM Transactions on Graphics*, 23(3):303–308, 2004. 1
- [11] H. Lombaert, Y. Sun, L. Grady, and C. Xu. A multilevel banded graph cuts method for fast image segmentation. In *ICCV*, volume 1, pages 259–265, 2005. 1
- [12] C. Rother, V. Kolmogorov, and A. Blake. "GrabCut": Interactive foreground extraction using iterated graph cuts. In *SIGGRAPH*, pages 309–314, 2004. 4
- [13] A. Sinop and L. Grady. Accurate banded graph cut segmentation of thin structures using laplacian pyramids. In *MIC-CAI*, volume 9, pages 896–903, 2006. 1, 2, 3
- [14] P. Strandmark and F. Kahl. Parallel and distributed graph cuts by dual decomposition. In *CVPR*, pages 2085–2092, 2010. 1




Figure 2: Illustration of the lower bound η_{min} for segmenting a 2D synthetic image. In this experiment, $\eta_{min} \simeq 0.523$ and we set $r = 10$ using connectivity 1. On all images, the pixels belonging to \mathcal{G}' are superimposed in yellow to the original image by transparency. The bottom row correspond to a close-up of the box in purple color. Observe how the reduced graph \mathcal{G}' splits into multiple pieces as soon as $\eta \leq \eta_{min}$.




Figure 3: Memory gain when segmenting a 2D synthetic image corrupted by 10% of impulsive noise (left). Top row shows the nodes of the reduced graph in light gray while bottom row shows the corresponding segmentation. In this experiment, we set $r = 3$ and use connectivity 1.




Figure 4: Simultaneous segmentation and filtering of a 3D noisy image (left). In this picture, the white spots correspond to cell nuclei in a mouse cerebellum. Top row shows the nodes of the reduced graph in light gray while bottom row shows the corresponding segmentation. In this experiment, we set $r = 5$ and use connectivity 1.