

SCIENTIFIC REPORTS



OPEN

Unsupervised Medical Image Segmentation Based on the Local Center of Mass

Iman Aganj^{1,2}, Mukesh G. Harisinghani^{1,3}, Ralph Weissleder^{1,3} & Bruce Fischl^{1,2}

Image segmentation is a critical step in numerous medical imaging studies, which can be facilitated by automatic computational techniques. Supervised methods, although highly effective, require large training datasets of manually labeled images that are labor-intensive to produce. Unsupervised methods, on the contrary, can be used in the absence of training data to segment new images. We introduce a new approach to unsupervised image segmentation that is based on the computation of the local center of mass. We propose an efficient method to group the pixels of a one-dimensional signal, which we then use in an iterative algorithm for two- and three-dimensional image segmentation. We validate our method on a 2D X-ray image, a 3D abdominal magnetic resonance (MR) image and a dataset of 3D cardiovascular MR images.

Image segmentation is the process of partitioning the set of image pixels into subsets, where the pixels in each subset are related, e.g. with respect to their intensities and/or locations. Segmentation of biomedical images is a central step in many medical imaging studies. Automating segmentation can be highly beneficial, as it can enable the large-scale studies needed to find subtle changes and disease effects.

Numerous approaches to medical image segmentation have been proposed (see surveys)^{1,2}. Provided that a large training dataset of a specific class of images – with ground-truth labels – is available, *supervised* segmentation can be effective in producing accurate results. Such a dataset can be exploited to train a deep neural network³ or create probabilistic atlases⁴, which can in turn generate accurate segmentations of new images in the same image class. Creating training datasets, however, requires manual delineation of labels by experts, which is labor-intensive and costly for a dataset that is sufficiently large for training. Furthermore, new image types are constantly being developed, making the creation of the manual training sets needed for supervised segmentation an onerous and ongoing task. Finally, in the case of an uncommon type of data, providing a large-enough dataset to label may be challenging.

Unsupervised segmentation, in contrast, has the advantage of not requiring training data to segment images and therefore is helpful in the absence of a manually-labeled dataset. Unsupervised segmentation methods are more generally applicable and more robust to atypical or unseen situations^{5–10}. In addition, the results of such methods are potential starting points for manual segmentation, thereby expediting the creation of training datasets. Since even an unsupervised segmentation algorithm may be effective in segmenting some but not all classes of images, it is important to explore new methodologies to expand the suitable options for unsupervised segmentation of various classes of medical images.

In this work, we introduce a new approach to unsupervised medical image segmentation, which is based on the computation of the one-dimensional (1D) local center of mass (CM). We first propose an efficient method to compute, for each pixel in a 1D signal, the CM of the region containing the pixel. In other words, by using the information from the entire signal, we group the pixels based on the CM of the regions where they are located. Next, we exploit this method for 2D and 3D image segmentation, by computing the local CMs of image pixels in many different orientations and then iteratively updating the segmentation label of each pixel by choosing from the labels at its local CMs. (We use the word “pixel” interchangeably to mean the index of a sample, a pixel, or a voxel, in a discrete-domain 1D signal, a 2D image, or a 3D volume, respectively.) We evaluate our approach qualitatively on a 2D hand X-ray image and a 3D abdominal magnetic resonance (MR) image, as well as quantitatively on a 3D dataset of 10 cardiovascular MR images.

¹Department of Radiology, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA. ²Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA. ³Center for Systems Biology, Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA. Correspondence and requests for materials should be addressed to I.A. (email: iman@nmr.mgh.harvard.edu)

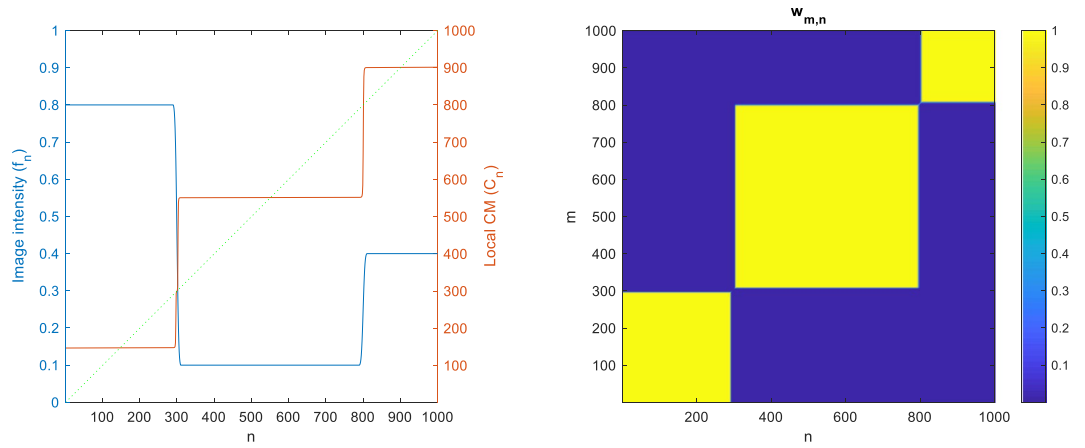


Figure 1. Left: Image intensity profile (f_n , blue), local CM (C_n , red) and the identity line (n , dotted green). Right: The weighting ($w_{m,n}$).

In the following, we describe the proposed method in detail (the Methods section), present experimental results (the Results section), discuss them (the Discussion section) and conclude the paper (the Conclusion section).

Methods

Computation of 1D Local Center of Mass. Let $f: \Omega \rightarrow \mathbb{R}$ be a 1D discrete-domain image-intensity signal of length N , with $\Omega: = \{1, \dots, N\}$. We group the pixels of f into disjoint regions based on the computed CM of each pixel’s putative region, which we shall refer to as the *local* CM, $C: \Omega \rightarrow \mathbb{R}$. In other words, the value of the local CM at n , C_n , is the CM of the region containing the n^{th} pixel. We calculate the local CM as:

$$C_n = \frac{\sum_{m=1}^N w_{m,n} m}{\sum_{m=1}^N w_{m,n}}, \tag{1}$$

where the nonnegative weighting $w_{m,n}$ must be computed from f such that it is large if the m^{th} and n^{th} pixels are in the same region and small otherwise. Consequently, for pixels that are in the same region, the function C will point to roughly the same location (the region’s CM), thereby assigning those pixels to the same cluster. A pixel is therefore clustered by exploiting the information provided by the entire signal, as opposed to only its neighboring pixels. Computing C (the entire set $\{C_n | n \in \Omega\}$) is, however, in general expensive, costing $\mathcal{O}(N^2)$.

Computational Complexity Reduction. To make the problem tractable, we must choose w such that, while it serves the abovementioned pixel-grouping purpose, the computational cost of C can also be reduced. Here, by choosing w as follows, we will significantly reduce the cost of computing C , to $\mathcal{O}(N)$:

$$w_{m,n} := e^{-|D_m - D_n|}, \tag{2}$$

where,

$$D_n := \alpha \sum_{i=1}^n |f_{i+1} - f_i|^p, \tag{3}$$

with manually chosen $\alpha, p > 0$. (We define $f_{N+1} := f_N$ for ease of notation.) One can see that the more and stronger the signal edges between the m^{th} and n^{th} pixels are, the larger $|D_m - D_n|$ and thus the smaller $w_{m,n}$ is, indicating that the two pixels are in different regions. On the other hand, if there are few edges between the m^{th} and n^{th} pixels, $w_{m,n}$ will be large, indicating that the two pixels are in the same region. Figure 1 shows an example of a 1D signal f_n (left, blue curve), along with its computed $w_{m,n}$ (right). Note that Eq. (3) is trivially extendable to multichannel (such as RGB) images, by computing the sum separately for each channel and adding the results leading to an aggregate D_n .

As for the computation of C , since D_n is non-decreasing with respect to n , $w_{m,n}$ can be rewritten as $e^{-D_n} e^{D_m}$ for $m \leq n$ and $e^{D_n} e^{-D_m}$ for $m > n$, leading to the following expansion of Eq. (1) to compute the elements of C :

$$C_n = \frac{e^{-D_n} \sum_{m=1}^n e^{D_m} m + e^{D_n} \sum_{m=n+1}^N e^{-D_m} m}{e^{-D_n} \sum_{m=1}^n e^{D_m} + e^{D_n} \sum_{m=n+1}^N e^{-D_m}}. \tag{4}$$

Given that all the sums in Eqs (3) and (4) can be pre-computed *recursively* in $\mathcal{O}(N)$ and stored for all n , the entire C can now be computed efficiently and non-iteratively in $\mathcal{O}(N)$. Figure 1 (left) shows the plot of C_n (red) for our example, in addition to the identity line (dotted green) for the reference. As expected, C_n is piecewise constant

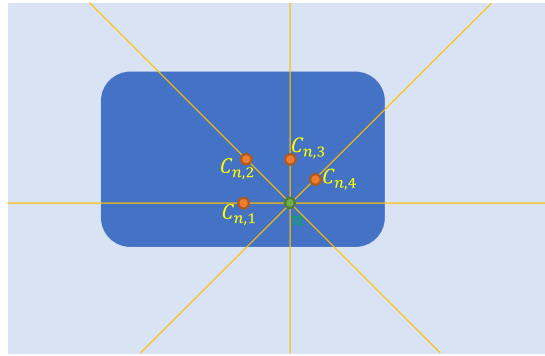


Figure 2. Local CMs, $C_n = \{C_{n,1}, \dots, C_{n,4}\}$, of the n^{th} pixel of the image, along $K=4$ orientations.

and its value points to the CM of each interval in the signal (e.g., it intersects the identity line at the centers of the intervals).

Image Segmentation. We now employ the aforementioned 1D local CM computation method to develop an iterative algorithm for unsupervised segmentation of a 2D or 3D image with N pixels. $\Omega = \{1, \dots, N\}$ now reflects the set of pixel indices of the image. Essentially, we assign the label of an estimated local CM to pixels pointing to that CM, causing the pixels in the same region to eventually share the same label. This is reminiscent of watershed methods⁵, where pixels corresponding to a common point form a catchment basin with a unique label. However, the proposed local-CM-based approach labels a pixel using the information from the entire image, i.e. beyond just the neighboring pixels, while avoiding a quadratic computational complexity.

We have shown in the previous subsections how to efficiently compute the local CMs from a 1D signal. For 2D and 3D segmentation, however, we must exploit the spatial relations among pixels in all dimensions. Accordingly, we make use of the proposed 1D CM computation technique and find the 1D local CMs independently in K orientations uniformly distributed on the unit semicircle or hemisphere (for the 2D and 3D cases, respectively). For the n^{th} pixel, we obtain a set of K pixels, $C_n := \{C_{n,k} \in \Omega | k = 1, \dots, K\}$, where $C_{n,k}$ is (the index of) the pixel closest to the local CM of the n^{th} pixel along the k^{th} orientation (i.e., nearest-neighbor interpolation of the CM on the Cartesian grid). An example of a pixel with its local CMs computed in several orientations is illustrated in Fig. 2. Note that the discrete derivative of the image in Eq. (3) is computed along the chosen orientation. To ensure that the set C_n has a uniform spatial (rather than angular) distribution, we down-sample it nonuniformly with a rate proportional to $1/d_k$ or $1/d_k^2$ in the 2D and 3D cases, respectively, where d_k is the distance between the n^{th} pixel and $C_{n,k}$.

For a 3D image of size $N_1 \times N_2 \times N_3$ (where $N_1 N_2 N_3 = N$), let us consider the computation of the local CMs of all pixels in the orientation along the first image dimension, i.e. $C_{n,k}$ for all $n \in \Omega$, for the fixed orientation k that corresponds to the x -axis. Using the computational complexity reduction technique described in the Computational Complexity Reduction section, local CMs of each row will cost $\mathcal{O}(N_1)$. Given that there is a total of $N_2 N_3$ rows, the computation of the local CMs for all image pixels (for the fixed orientation) will cost $\mathcal{O}(N_2 N_3 \cdot N_1) = \mathcal{O}(N)$. Similarly, the computational complexity is the same for any other orientation. Therefore, we pre-compute the local CMs along all K orientations, i.e. the entire $\{C_n | n \in \Omega\}$, efficiently in $\mathcal{O}(KN)$ and store it in memory.

Next, we initialize a label image, $L^0: \Omega \rightarrow \mathbb{N}$, with N unique randomly-located labels. We then start an iterative loop, consisting of two phases, where we update the label image, L , at each iteration. For all $n \in \Omega$, we reassign the label L_n of the n^{th} pixel by making use of the pre-computed set C_n (that includes the local CMs of the n^{th} pixel in different orientations):

- In the first phase, we randomly choose a pixel $m_{\text{rand}} \in C_n$ and assign its label from the last iteration to L_n ; i.e., $L_n \leftarrow L_{m_{\text{rand}}}^0$.
- In the second phase, we instead find the most frequent label (from last iteration) corresponding to the pixels in C_n and assign it to L_n ; i.e., $L_n \leftarrow \text{mode}\{L_m^0 | m \in C_n\}$.

The generated L at each iteration becomes the new L^0 for the next iteration. We start with the first phase and after a predefined number of phase-1 iterations, t , switch to the second phase. Then, during the phase-2 iterations, we stop the loop if there is no more change in L , or when the iteration number reaches a predefined maximum. Finally, expecting the resulting L to have $l \ll N$ unique values, we map the labels in L to $\{1, \dots, l\}$, before returning L as the segmentation results. Note that, due to the randomness introduced in the first phase, one may obtain slightly different results by repeating a segmentation experiment.

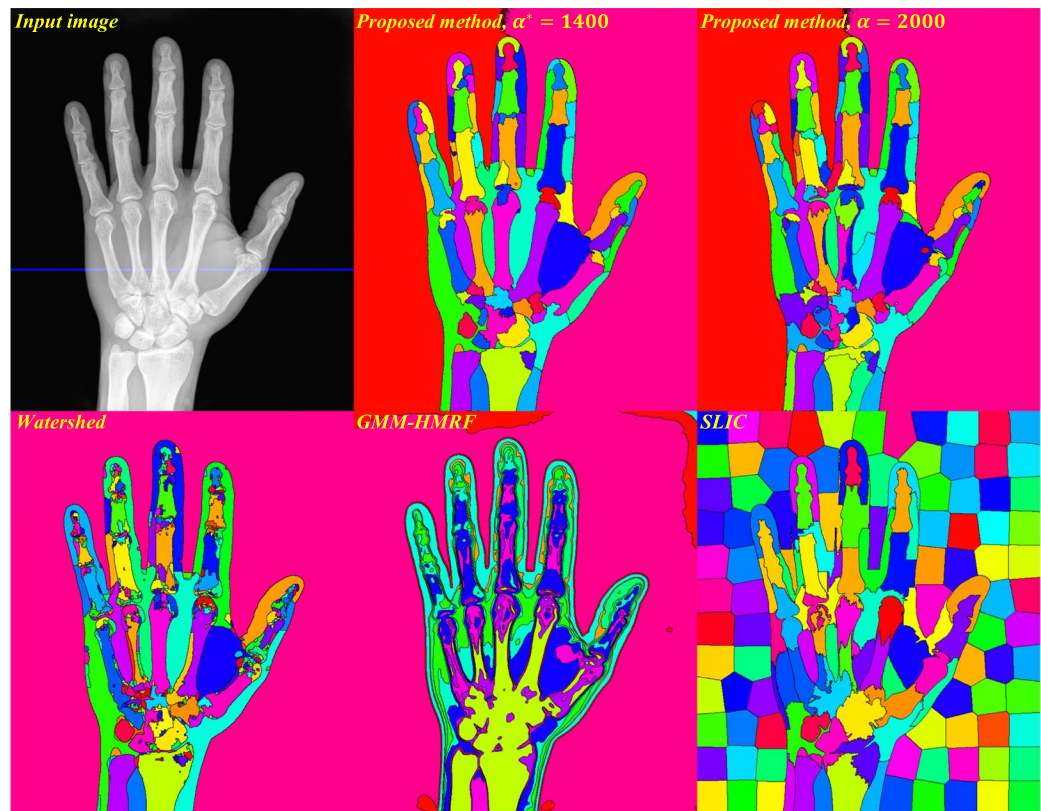


Figure 3. The input X-ray image (top, left), with segmentation results using the proposed local-CM-based (top, middle and right), the watershed (bottom, left), the GMM-HMRF (bottom, middle) and the SLIC (bottom, right) methods (X-Ray image courtesy of Trace Meek).

Results

We implemented our unsupervised segmentation algorithm in MATLAB and evaluated it on 2D and 3D medical images. We also compared it with three other unsupervised segmentation algorithms: the watershed segmentation⁵, a Gaussian-mixture-model-based hidden-Markov-random-field (GMM-HMRF) model⁷ initialized with k-means clustering and the simple linear iterative clustering (SLIC) superpixel algorithm⁸ (with constant compactness). We used the MATLAB Image Processing Toolbox™ for watershed and SLIC and a public MATLAB toolbox⁷ for GMM-HMRF. For watershed segmentation, we created the edge map of the image using the Sobel filter, computed its *h-minima* transform, i.e. removed its local minima that were shallower than the scalar *h* and then applied the watershed transform. For GMM-HMRF, we set the maximum number of iterations to the suggested value of 10 for both the EM algorithm and the MAP estimation.

After initially trying the values of $p = 1, 2, 3$ for the p -norm in Eq. (3), we heuristically concluded that our method produced the most reasonable results with $p^* = 2$. For each image, we ran our method with ranges of values for α (the Computational Complexity Reduction section) and for the number of phase-1 iterations t (the Image Segmentation section), the watershed algorithm with a range of values for h , the GMM-HMRF algorithm with ranges of values for the number of regions r (for both the GMM-HMRF and its k-means initializer) and the number of GMM components g and the SLIC algorithm with a range of values for the number of superpixels s . Next, for each method, we determined the optimal parameter values: in the 2D X-ray Image Segmentation and the 3D Abdominal MR Image Segmentation sections by heuristically choosing the least over-segmented result among those segmentation results that reflected all major anatomical regions and in the 3D Cardiovascular MR Image Segmentation section once by choosing the result with the highest mean Dice score and a second time by cross validation.

Label edges are shown in black for better contrast and clearer visualization in the figures. We also used a simple greedy algorithm to roughly match the colors of overlapping labels among subfigures for easier comparison.

2D X-ray Image Segmentation. We first evaluated the four (proposed local-CM-based, watershed, GMM-HMRF and SLIC) methods on a publicly available 2D hand X-ray image¹¹, with the dimensions of 750×880 pixels (Fig. 3, top left). Here and in the 3D Abdominal MR Image Segmentation section, images were initially normalized to have intensity values between 0 and 1. We chose an angular resolution of 1° and computed the local CMs in $K = 180$ orientations uniformly distributed on the semicircle. Figure 3 (top) shows segmentation results by our local-CM-based method, with optimal $t^* = 2000$ phase-1 iterations (10,000 total iterations) and optimal $\alpha^* = 1400$ (middle) and $\alpha = 2000$ (right). Figure 3 (bottom) depicts the results for the watershed method

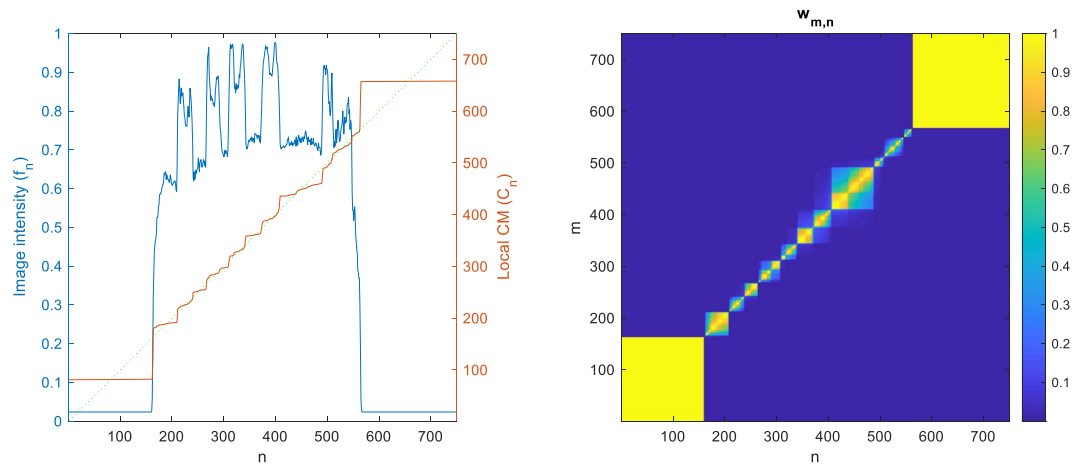


Figure 4. Left: Image intensity profile (blue) of the horizontal blue line in Fig. 3 (top, left), local CM (red) and the identity line (dotted green). Right: The weighting ($w_{m,n}$).

with optimal $h^* = 0.116$ (left), the GMM-HMRF model with optimal $r^* = 11$ and $g^* = 2$ (middle) and the SLIC algorithm with optimal $s^* = 130$ (right).

The 1D image intensity profile along the horizontal blue line in Fig. 3 (top, left) is plotted in Fig. 4 (blue curve), along with the local CM of this 1D signal computed from Eq. (4) with $\alpha = 200$ (red curve) and the identity function for the reference (green dotted line). The local CM curve (red) can be seen to be almost piecewise constant, with its values indicating the centers of the 1D intervals.

3D Abdominal MR Image Segmentation. Next, we applied the four methods to a 3D abdominal MR image acquired as part of a previous study¹², quoting from which: “All participants were 18 y of age or older. Informed consent was obtained after the nature and possible consequences of the studies were explained. [...] The Joslin Diabetes Center Committee on Human Studies and the Massachusetts General Hospital Institutional Review Board approved the protocols. FDA Investigational New Drug approval was obtained for the use of ferumoxytol for MRI.”

The T1-weighted image had been acquired with the voxel size of $1.4 \times 1.4 \times 3.5 \text{ mm}^3$, after the administration of a magnetic nanoparticle agent. We upsampled the volume in the slice-selection direction with linear interpolation to obtain a $(1.4 \text{ mm})^3$ isotropic-voxel volume of the size $224 \times 256 \times 98$ voxels. Figure 5 illustrates coronal, sagittal and axial slices of the image (top, left), along with those of the 3D segmentation results using the four methods. We chose an angular resolution of 10° and computed the local CMs in $K = 193$ orientations uniformly distributed on the hemisphere. For the proposed method, we show the results obtained with optimal $\alpha^* = 2700$ and $\alpha = 6000$ and optimal $t^* = 1500$ phase-1 iterations (5000 total iterations). Optimal parameter values for other methods were $h^* = 0.154$ (watershed), $r^* = 9$ and $g^* = 1$ (GMM-HMRF) and $s^* = 46$ (SLIC).

3D Cardiovascular MR Image Segmentation. Lastly, for a quantitative assessment of the performance of the four methods, we used a public dataset of 10 cardiovascular MR images¹³ acquired at 1.5 T, which also included manual segmentation of the ventricular myocardium and the blood pool. We used the available cropped images (“axial, cropped training data”) with varying sizes ranging from $97 \times 102 \times 164$ to $165 \times 239 \times 181$ voxels and the voxel size of $0.9 \times 0.9 \times 0.85 \text{ mm}^3$. We divided each image by its intensity standard deviation and then normalized the images by the same constant so their maxima were on average 1 (their minima were 0). We chose an angular resolution of 20° , resulting in $K = 54$ uniformly distributed orientations on the hemisphere and ran our algorithm for 5000 iterations. To segment each image, we used a range of values for each parameter of each method (see Table 1). For each of the two manual labels, we identified the corresponding label in each resulting segmentation as the one with the maximal Dice’s overlap coefficient with the manual label. Then, for each subject and each method, we chose the optimal segmentation result that produced the maximal Dice score (averaged across the two labels).

Cross-subject mean (\pm standard error) of the optimal Dice scores and optimal parameters are shown in Table 1. As can be seen, the proposed local-CM-based method resulted in a higher mean optimal Dice score than the competing methods did. Figure 6 illustrates, for all the 10 subjects, how robust the mean Dice score was with respect to varying α (left) and t (right). The vertical axis shows the best score achieved with respect to the remaining parameter. Alternatively, we also performed a leave-one-out cross validation (CV) by testing the methods on a left-out subject with the median of the optimal parameter values computed from the other 9 subjects, repeating it 10 times (each time leaving out a different subject) and averaging the Dice scores. As Table 1 shows, the proposed method achieved a higher CV Dice score than the watershed and SLIC methods did, but a lower one than the GMM-HMRF method did.

We distributed the above experiments to different processors with inhomogeneous hardware. Hence, to compare the runtime of the methods on the same hardware, we reran the optimal experiments (two experiments at a time) on a desktop PC with two Intel® Xeon® E5-2637 v3 3.50 GHz processors, each with four cores. We did not parallelize our codes; however, MATLAB may multithread some of its internal functions. The cross-subject

	Proposed	Watershed	GMM-HMRF	SLIC
Optimal Dice	0.60 ± 0.02	0.53 ± 0.04	0.55 ± 0.01	0.56 ± 0.01
CV Dice	0.51 ± 0.04	0.41 ± 0.05	0.54 ± 0.01	0.49 ± 0.02
Optimal parameters	$\alpha^* = 2100 \pm 300$ $t^* = 3050 \pm 400$	$h^* = 0.41 \pm 0.05$	$r^* = 4 \pm 0.5$ $g^* = 1 \pm 0$	$s^* = 8 \pm 2$
Tested range of parameters	α : 100~8000 t : 50~5000	h : 0~1.5	r : 2~20 g : 1~5	s : 2~500
CPU Runtime	37000 ± 11000 s (10 \pm 3 hr)	6 ± 1 s	9000 ± 2000 s (2.5 \pm 0.6 hr)	2.5 ± 0.4 s

Table 1. Summary of the quantitative results on cardiovascular MR images.

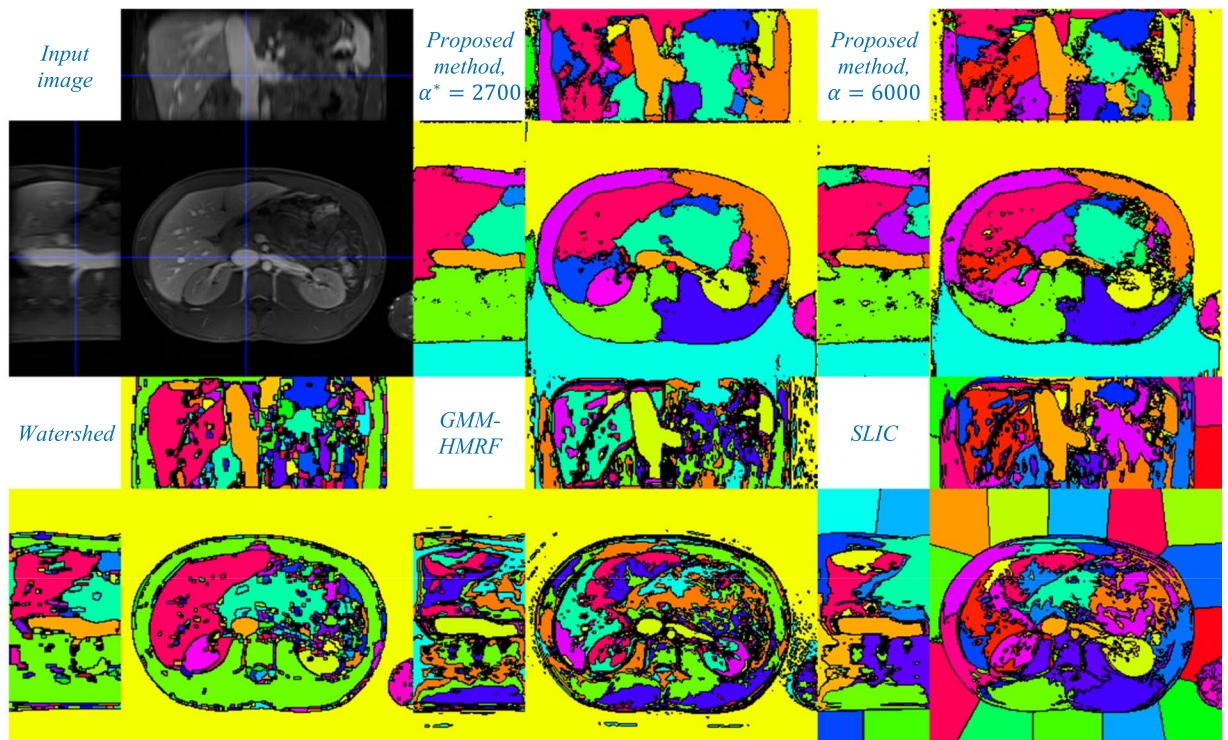


Figure 5. Input abdominal MR image (top, left), with segmentation results using the proposed local-CM-based (top, middle and right), the watershed (bottom, left), the GMM-HMRF (bottom, middle) and the SLIC (bottom, right) methods.

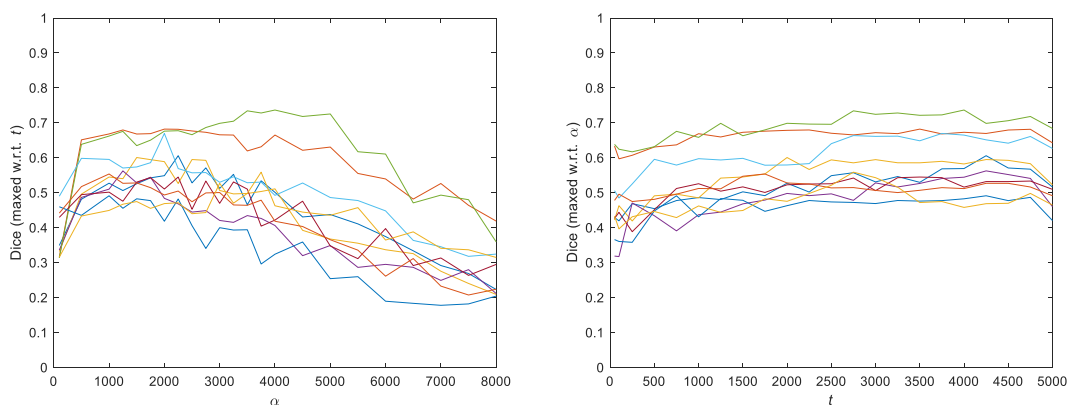


Figure 6. Dice scores for all 10 subjects, for a fixed α maximized over t (left) and for a fixed t maximized over α (right).

means of the runtimes are shown in Table 1. Due to the complexity of the mode function, the bulk of our prototype's runtime is spent in phase-2 iterations, with a phase-2 iteration being on average 7 times slower than a phase-1 iteration. Note that, to ensure convergence for all tests, we ran our method for a high total of 5000 iterations; however, one could reduce the number of phase-2 iterations and still obtain reasonable results. Further actions to gain speed would be to use a smaller number of orientations, K and to run the algorithm on a graphics processing unit (GPU).

Discussion

As the figures demonstrate, the proposed local-CM-based method generally produced less over-segmented results than the rest of the methods did (except in the background). This may be because our method assigns each pixel label by considering – not only the neighboring pixels, but – the entire image. Increasing h improved the over-segmentation by the watershed method, but at the price of not capturing some boundaries. The watershed method, however, resulted in better-defined borders between articular cartilages (Fig. 3) and a single-segment liver (Fig. 5). The GMM-HMRF model seems to emphasize more the image intensity than the geometry of a segment, hence leading to rather tissue- than organ-specific segmentation. This model is not so successful in segmenting individual bones (Fig. 3) and some organs (Fig. 5), possibly because it assigns a single gaussian to each identified connected tissue type. Regarding the SLIC segmentation, the borders of its individual bone and organ labels do not seem as well-defined as the segmentation by the proposed method. Also, as expected, larger regions such as the background are divided into many superpixels. To accurately evaluate the performance of the segmentation methods, we have avoided any postprocessing, such as region-merging techniques⁶ that could potentially alleviate the over-segmentation issue.

In our quantitative validation (the 3D Cardiovascular MR Image Segmentation section), the proposed local-CM-based method achieved the highest optimal Dice overlap score among the four methods, but came in second place with regard to the CV Dice score, after the GMM-HMRF model. The latter may be because the ventricular myocardium and the blood pool are rather tissue types, whose image intensities are suitable for mixture models (see above). Note that the Dice scores were produced without any shape prior or training. Nonetheless, given that we tuned one or two parameters for each method, our *evaluation* of the four methods can be considered as mostly but not totally unsupervised. We have empirically computed the optimal parameter values in our experiments and provided them throughout the Results section. In practice, these values can be a good starting point for the interested reader who wants to try the proposed CM-based segmentation on a new image (that has intensities normalized between 0 and 1). Automatic optimization of the parameters of our method (α and t) is part of our future research.

The often-improved results of the proposed method came at the price of a higher computational cost (Table 1). To reduce the runtime, the algorithm can be run with considerably lower numbers of orientations (K) and iterations, while still producing reasonable results. Moreover, our publicly available codes (the Data Availability section) are GPU-compatible, making it possible to significantly speed up the computation.

Our algorithm uses the CM of a region as a reference point, whose label it iteratively propagates to the rest of the pixels in that region. If such a reference point is located inside the region, it will be unique to the region. For an interval in a 1D signal (which is always convex), the CM is trivially located inside the region (Figs 1 and 4). This is, yet, not guaranteed to be the case in higher dimensions, unless the region is convex. That is why we did not use the direct definition of the 2D/3D CM (with double/triple sums); instead, we exploited 1D CMs in different orientations to approximate local 2D/3D CMs. Being always inside the region, 1D CMs can be used as reference points with labels that are unique to the region. A drawback of this approximation though is that, in a highly nonconvex region, more iterations may be needed for the label to propagate entirely. In some cases, the algorithm may even converge to a locally optimal result where a non-convex region is over-segmented. The two labels computed for the background in Fig. 3 (top) and for the liver in Fig. 5 (top) are examples of such a phenomenon. Nevertheless, choosing a large enough number of phase-1 iterations (t) can alleviate this issue and, in fact, many nonconvex regions can be seen to have been successfully segmented by the proposed method.

Lastly, if an initial segmentation for some parts of the image is available, it can be used to initialize the labels (instead of random assignment). The labels in those regions can be chosen to either remain fixed or get updated in the iterations. This is especially useful for semi-automatic segmentation, where the user provides some information to the algorithm on which pixels should be grouped together.

Conclusion

We have introduced a new unsupervised medical image segmentation approach that groups the pixels in each region based on the *local* CMs of the region. We proposed an efficient scheme for the 1D case, in $\mathcal{O}(N)$, which we extended to higher dimensions via an iterative algorithm. Through qualitative and quantitative validation, we have shown the proposed method to often outperform three existing unsupervised segmentation methods. Automatic optimization of the two parameters of the method, α and the number of phase-1 iterations t , is part of the future work.

Data Availability

Our MATLAB codes for the proposed method are publicly available at: www.nitrc.org/projects/seg. The hand X-ray image¹¹ (<https://cnx.org/contents/FPtK1zmh@9.1:vaCcDeiY@4/Medical-Imaging>) in the 2D X-ray Image Segmentation section and the de-identified cardiovascular MR images¹³ (<http://segchd.csail.mit.edu>) in the 3D Cardiovascular MR Image Segmentation section are also publicly available. The abdominal MR image in the 3D Abdominal MR Image Segmentation section was secondarily used from a different study¹² and – at least for now – is not available to the public due to institutional policies.

References

1. Elnakib, A., Gimel'farb, G., Suri, J. S. & El-Baz, A. In *multi modality state-of-the-art medical image segmentation and registration methodologies: volume ii* (eds Ayman S. El-Baz, Rajendra Acharya U andrew F. Laine, & Jasjit S. Suri) 1–39 (Springer New York, 2011).
2. Setarehdan, S. K. & Singh, S. *Advanced algorithmic approaches to medical image segmentation: state-of-the-art applications in cardiology, neurology, mammography and pathology*. (Springer London, 2012).
3. Litjens, G. *et al.* A survey on deep learning in medical image analysis. *Medical Image Analysis* **42**, 60–88, <https://doi.org/10.1016/j.media.2017.07.005> (2017).
4. Cabezas, M., Oliver, A., Lladó, X., Freixenet, J. & Bach Cuadra, M. A review of atlas-based segmentation for magnetic resonance brain images. *Computer Methods and Programs in Biomedicine* **104**, e158–e177, <https://doi.org/10.1016/j.cmpb.2011.07.015> (2011).
5. Beucher, S. & Meyer, F. In *Optical Engineering* Vol. 34 433–433 (New York-Marcel Dekker Incorporated, 1992).
6. Haris, K., Efstratiadis, S. N., Maglaveras, N. & Katsaggelos, A. K. Hybrid image segmentation using watersheds and fast region merging. *IEEE Transactions on image processing* **7**, 1684–1699 (1998).
7. Wang, Q. GMM-based hidden markov random field for color image and 3d volume segmentation. *CoRR* **abs/1212.4527** (2012).
8. Achanta, R. *et al.* SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34**, 2274–2282, <https://doi.org/10.1109/TPAMI.2012.120> (2012).
9. Jianbo, S. & Malik, J. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**, 888–905, <https://doi.org/10.1109/34.868688> (2000).
10. Chen, C.-C., Juan, H.-H., Tsai, M.-Y. & Lu, H. H.-S. Unsupervised learning and pattern recognition of biological data structures with density functional theory and machine learning. *Scientific Reports* **8**, 557, <https://doi.org/10.1038/s41598-017-18931-5> (2018).
11. *Anatomy and Physiology*. (OpenStax CNX, 2017).
12. Gaglia, J. L. *et al.* Noninvasive mapping of pancreatic inflammation in recent-onset type-1 diabetes patients. *Proceedings of the National Academy of Sciences* **112**, 2139–2144, <https://doi.org/10.1073/pnas.1424993112> (2015).
13. Pace, D. F. *et al.* *Medical image computing and computer-assisted intervention – MICCAI 2015* 80–88 (Springer International Publishing, Cham, 2015).

Acknowledgements

Support for this research was provided by the National Institutes of Health (NIH), specifically the National Institute of Diabetes and Digestive and Kidney Diseases (K01DK101631, R21DK108277), the National Institute for Biomedical Imaging and Bioengineering (P41EB015896, R01EB006758, R21EB018907, R01EB019956), the National Institute on Aging (AG022381, 5R01AG008122-22, R01AG016495-11, R01AG016495), the National Center for Alternative Medicine (RC1AT005728-01), the National Institute for Neurological Disorders and Stroke (R01NS052585, R21NS072652, R01NS070963, R01NS083534, U01NS086625) and the NIH Blueprint for Neuroscience Research (U01MH093765), part of the multi-institutional Human Connectome Project. Additional support was provided by the BrightFocus Foundation (A2016172S). Computational resources were provided through NIH Shared Instrumentation Grants (S10RR023401, S10RR019307, S10RR023043, S10RR028832) and an NVIDIA GPU grant.

Author Contributions

I.A. developed the methods, ran the experiments and wrote the manuscript. M.H., R.W. and B.F. reviewed the manuscript multiple times and provided substantial feedback. M.H. and R.W. provided the abdominal MR image used in the 3D Abdominal MR Image Segmentation section.

Additional Information

Competing Interests: B.F. has a financial interest in CorticoMetrics, a company whose medical pursuits focus on brain imaging and measurement technologies. B.F.'s interests were reviewed and are managed by Massachusetts General Hospital and Partners HealthCare in accordance with their conflict of interest policies. The rest of the coauthors declare no potential competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018