Domain-Specific Seed Removal for Biomedical Image Segmentation using SICLE

Fábio F. Kochem*, Felipe Belém*, Alexandre X. Falcão†,
Zenilton K. G. do P. Jr.* and Silvio Jamil F. Guimarães*

* Image and Multimedia Data Science Laboratory (IMScience) – Pontifical Catholic University of Minas Gerais
Belo Horizonte, MG, 30535-610, Brazil
Email: {felipebelem,zenilton,sjamil}@pucminas.br, fabio.freire@sga.pucminas.br
† Laboratory of Image and Data Science (LIDS) – University of Campinas
Campinas, SP, 13083-970, Brazil
Email: afalcao@ic.unicamp.br

***Abstract*—Superpixels through Iterative CLEarcutting (SICLE) is an efficient framework for image segmentation that operates on an iterative principle of seed removal to refine the final result. While effective for general purposes, its reliance on standard seed removal criteria limits its performance in specialized domains. This is particularly evident in biomedical analysis, where the goal is often to isolate a single object of interest using a low number of superpixels, a task where generic criteria often fail. To address this limitation, this work proposes two seed removal criteria: (i) a Position-Based Criterion that leverages prior anatomical knowledge to guide segmentation in medical images; (ii) a Color-Based Criterion specialized for identifying targets in pathological images based on their distinct color signature. By replacing generic heuristics with these domain-specific functions, we demonstrate that SICLE can be transformed into a more robust and specialized tool for targeted biomedical image analysis, significantly improving object delineation accuracy, including situations with a small quantity of superpixels.**

I. Introduction

Superpixels are group of pixels that form perceptually meaningful, homogeneous regions, offering a richer representation of an image than individual pixels alone. For any superpixel method to be effective, it must demonstrate strong adherence to object boundaries, computational efficiency, and a controllable number of generated segments. Due to these valuable properties, superpixels serve as a fundamental pre-processing step in a wide array of computer vision tasks, including object tracking [1] and image classification [2]. Over the years, various strategies have been developed to generate them. Classical approaches like Simple Linear Iterative Clustering (SLIC) [3] employ clustering-based techniques , while others like Superpixel Hierarchy (SH) [4] are effective hierarchical graph-based examples. More recently, the Superpixels through Iterative CLEarcutting (SICLE) [5] framework was introduced as an efficient method that consists of three steps: seed oversampling, superpixel generation, and seed removal.

SICLE framework operates on a iterative refinement principle. Its pipeline begins by oversampling a large number of initial seeds in an image, a strategy that improves the probability of selecting seeds that will lead to an accurate object delineation. Following this, the framework repeatedly executes two steps: superpixel generation using the Image Foresting Transform (IFT) [6] and the removal of irrelevant seeds. This process continues until a final, desired number of superpixels is achieved. This methodology has proven to be highly effective, with results showing that SICLE surpasses state-of-the-art methods in both efficiency and delineation accuracy, and is significantly faster than its predecessor and the popular SLIC algorithm. Its successful application on challenging datasets, such as CT scans of the human liver where boundaries are smooth and difficult to detect, demonstrates its potential for biomedical image analysis.

The primary strength of the SICLE framework is its flexible seed removal step, which iteratively refines the segmentation. However, a limitation of the standard SICLE approach becomes apparent in scenarios that require a very low number of superpixels depending on the criterion used, as shown in Figure 1. Performance evaluations show that while the method is highly effective overall, metrics such as Boundary Recall tend to decrease while under-segmentation superpixel must cover a larger and more heterogeneous image region, making it difficult for criteria based on size or contrast to accurately preserve the boundaries of the specific object of interest. In biomedical analysis, the goal is often to isolate a single organ or a specific cellular structure, which is an ideal use case for segmentation with a low superpixel count. Therefore, optimizing SICLE's performance with a precise criterion is critical for its practical application in these fields.

To address this limitation and enhance SICLE's performance for specialized image analysis, this work proposes two novel seed removal criteria tailored for specific biomedical datasets. The **Position-Based Criteria** is designed for the segmentation of CT scans. Its purpose is to leverage prior anatomical knowledge by prioritizing superpixels based on both their size and their spatial proximity to the known center of the organ. This ensures the segmentation remains focused on the structure of interest. The **Color-Based Criteria** evaluates seed relevance by combining superpixel size with colorimetric similarity to the specific and distinct color signature of the target object.

By replacing generic heuristics with these domain-specific functions, our goal is to achieve more accurate and robust object delineation, making SICLE a more powerful tool for targeted medical and pathological analysis.

This work is organized as follows. Section II details our proposed methodology, reviewing core concepts of the SICLE framework, Image Foresting Transform (IFT) and describing the standard seed removal criteria before formally introducing our two novel contributions. Section 3 presents the experimental results, outlining the setup and providing a quantitative and qualitative analysis of our methods. Finally, Section 4 offers the conclusion of our work, summarizing and discussing potential avenues for future research.

II. Methodology

In this section, we present the SICLE framework and, then, we introduce our new seed relevance criteria for improving segmentations of computer tomography slices of the human liver, and of delineating colorized images of helminth eggs. SICLE is composed of three individual steps: (i) seed oversampling; (ii) superpixel generation; and (iii) seed removal. In this work, we focus only on step (iii).

In (i), N_0 initial points (*i.e.*, *seeds*) significantly greater than the final number of superpixels N_f are randomly selected in order to reduce the chances of not placing a seed within the objects of interest. As one may note, by imposing such high quantity (up to 3000 as recommended [5]), we assume that the resulting *seed set* $\mathcal{S} \subset \mathcal{V}$ contains the necessary seeds for accurate object delineation and, therefore, the goal is to sustain performance whilst removing seeds in step (iii).

From such seed set $\mathcal{S}$, the superpixels are generated using the seed-restricted version of the *Image Foresting Transform* (IFT) framework. First, let $I = (\mathcal{P}, \mathbf{F})$ denote an *bidimensional image*, where $\mathcal{P} \subset \mathbb{Z}^2$ represents the set of pixels whose features $\mathbf{F}(p) \in \mathbb{R}^m$, for $m > 0$, are mapped by $\mathbf{F}$. When $m = 1$, we say that I is a *grayscale image*; otherwise, it is a *colored* one. From I we may build an *undirected graph* $G = (\mathcal{V}, \mathcal{E})$ in $\mathcal{V} \subseteq \mathcal{P}$ represents the *vertex set*, and $\mathcal{E} \subseteq \mathcal{V}^2$ represents the set of *edges*. One may define $\mathcal{E}$ from the 4- or 8-neighborhood from each pixel.

A *path* is a sequence of distinct and *adjacent* vertices from a seed s to a non-seed vertex t, and it is denoted by $\rho_{s \to t}$. For simplicity, we may omit the *origin* s whenever it is irrelevant for the context. In the IFT, each path ρ_t is assigned to a *path-cost*, being the $\mathbf{f}_{\max}$ and $\mathbf{f}_{sum}$ (see Equation 1) being the most popular ones due to their efficacy.

$$\begin{aligned} \mathbf{f}_*(\rho_t = \langle t \rangle) &= 0 \\ \mathbf{f}_{\max}(\rho_{s \to t} \odot \langle t, v \rangle) &= \max\{\mathbf{f}_{\max}(\rho_{s \to t}), \|\mathbf{F}(s) - \mathbf{F}(v)||_2\} \\ \mathbf{f}_{\text{sum}}(\rho_{s \to t} \odot \langle t, v \rangle) &= \mathbf{f}_{\text{sum}}(\rho_{s \to t}) + \|t - v\|_2 \\ &\quad + (\iota \|\mathbf{F}(s) - \mathbf{F}(v)\|_2)^\beta \end{aligned} \tag{1}$$

where $\iota \geq 0$ and $\beta \geq 0$ controls the superpixels' *irregularity* and *boundary adherence*. A path ρ_t is said to be *optimum* if, for any other path $\tau_t \neq \rho_t$ in the set of all paths Π in G, we have that $\mathbf{f}(\rho_t) \leq \mathbf{f}(\tau_t)$. Even when certain conditions are not satisfied, the output still presents important properties for segmentation.

The IFT recurs to a generalization of the Dijkstra's algorithm for minimizing a *cost map* $\mathbf{C}(t) = \min_{\rho_{s \to t}}\{\mathbf{f}(\rho_{s \to t})\}$ by assigning an optimum path to each vertex. From such assignment, we may map each t recursively to its *predecessor* $\mathbf{P}(t)$ in the optimum-path $\rho_{s \to t}$, or even to its seed $\mathbf{R}(t) = s$. From $\mathbf{R}$, we may also assign t to its *optimum-path tree* $T_s = \{v : \mathbf{R}(v) = \mathbf{R}(t) = s\}$. In this work each superpixel is a tree and the segmentation a *spanning forest*.

After generating the superpixels, SICLE ranks seeds based on some mathematical criterion in order to define those deemed as "more relevant" (*i.e.*,, best assists in the object delineation). At the end of each iteration $i < \Omega$, an amount of $\max\{N_0^{-i/(\Omega-1)}, N_f\}$ most irrelevant seeds are removed, where $\Omega > 1$ is the maximum number of iterations for segmentation.

The flexibility of the SICLE framework is exemplified by its support for various seed removal criteria. These criteria are functions that evaluate a superpixel based on its intrinsic properties, such as size, and its relationship with adjacent superpixels, such as color contrast or spatial distance. Among the standard criteria, one of the most common is **Minimum Size-Contrast (MINSC)**. This criterion calculates relevance by the minimum color gradient between it and any of its neighbors. This approach favors the preservation of seeds within large, low-contrast, and homogeneous areas. Conversely, the **Maximum Size-Contrast (MAXSC)** criterion multiplies the size by the maximum color gradient, prioritizing seeds that form superpixels along high-contrast edges. Other criteria can incorporate spatial information, such as the **Spread** criterion, which combines size with the minimum spatial distance to a neighboring superpixel's centroid to encourage a more uniform distribution. While these criteria are effective in many scenarios, they operate without any high-level understanding of the image content.

To overcome the limitations of general-purpose criteria, we introduce two novel seed removal criteria designed for a specific biomedical imaging task by incorporating prior knowledge about the image content. The Position-Based Criterion is designed for single-channel medical images, where the object of interest has a known general location. The relevance $V(s)$ for a given seed s is calculated by combining its superpixel's relative size with its spatial proximity to a predefined anatomical landmark. It is formally defined as:

$$V(s) = S_{rel}(s) \cdot \left(1 - \frac{||C(s) - C_{liver}||}{A_{img}}\right)$$

where $S_{rel}(s)$ is the relative size of the superpixel, $C(s)$ is its centroid, C_{liver} is the predefined coordinate of the anatomical landmark, and A_{img} is the total area of the image. This formula ensures that the relevance score is the highest for large superpixels located near the landmark and decreases quadratically as the distance from the center increases. The

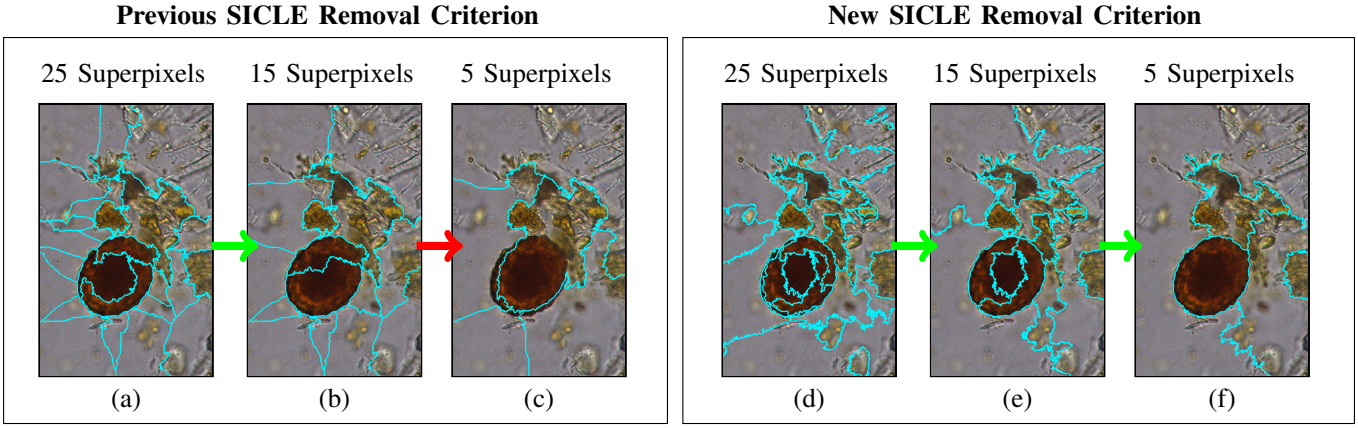

Fig. 1. Different segmentation outputs using two different seed removal criteria of an helminth egg image.

Color-Based criterion is tailored for multi-channel color images, where the target has a distinct color signature. Relevance is computed based on size and colorimetric similarity to a target color profile. The relevance $V(s)$ is given by:

$$V(s) = S_{rel}(s) \cdot \max\left(\frac{1}{d_{color}(s) + \epsilon}, G_{min}(s)\right)$$

where $G_{min}(s)$ is its minimum color gradient to a neighbor, ϵ is a small constant to prevent division by zero, and $d_{color}(s)$ is the minimum color distance from the superpixel (or its immediate neighbors) to a predefined target color. This formulation strongly prioritizes superpixels that closely match target color, as the inverse distance term becomes very large for small color differences, effectively guiding the segmentation to preserve objects with a specific color profile.

III. Experimental Results

In this section, we detail the experimental framework for evaluating our proposal and discuss the obtained results. To evaluate the criteria proposed, we selected two datasets, named LiverCT [7] and Parasites [8], that correspond to the purpose of each criterion. The Human liver dataset imposes a challenge in terms of delineation due to its smooth boundaries but always has a common region where the liver is located, making it an ideal test for the Position-Based Criterion. The Parasites dataset, conversely, contains colorized images of helminth eggs, presenting the challenge of isolating a specific object based on its distinct color signature against a complex background with impurities. To quantitatively assess perfomance, we employ two standard metrics: *Boundary Recall* (BR) measures the overlap between the superpixels' borders and the object boundaries (*i.e*, higher is better); *Under-segmentation Error* (UE) measures the error from multiple-object overlap by the superpixels (*i.e*, lower is better). Considering different segmentation approaches, we selected two variants of SICLE for criteria evaluation: (i) SICLE-COMP [5] and (ii) SICLE-IRREG [5]. SICLE-COMP favors spatial compactness, resulting in more regular, tile-like superpixels. SICLE-IRREG prioritizes adherence to object boundaries, allowing for the creation of irregular superpixels that can closely follow an object's true contour. The level of compactness in the SICLE-COMP variant is controlled by a irregularity parameter ι. While the standard SICLE-COMP method uses $\iota = 0.12$ for this parameter, our experiments determined that the optimal value for our proposed criteria is $\iota = 0.20$.

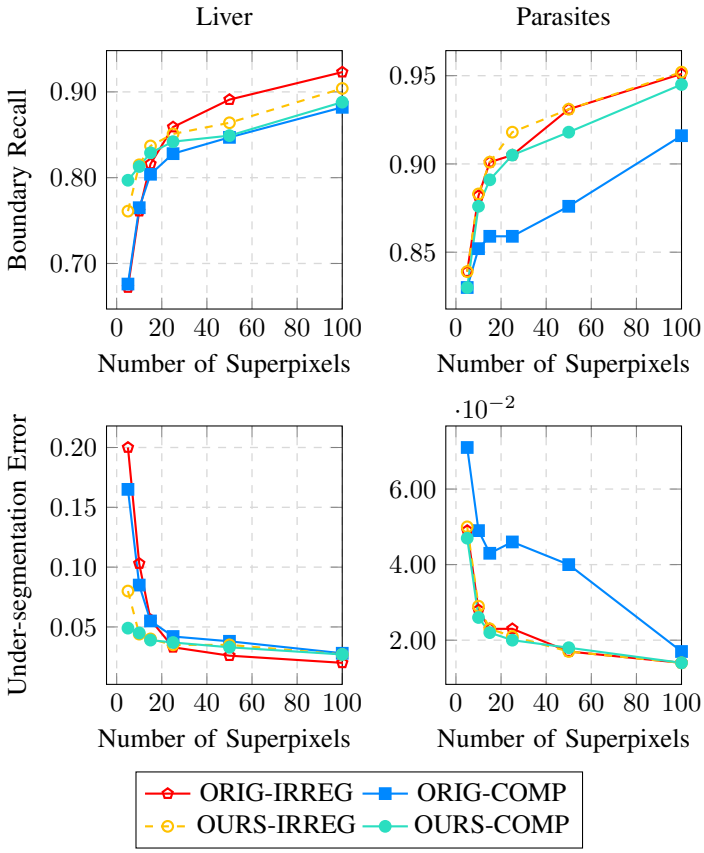

Fig. 2. Results of previous and new SICLE variants.

A. Analysis

As one may note from Figure 2, original SICLE variants significantly deteriorate their segmentation for quantities of

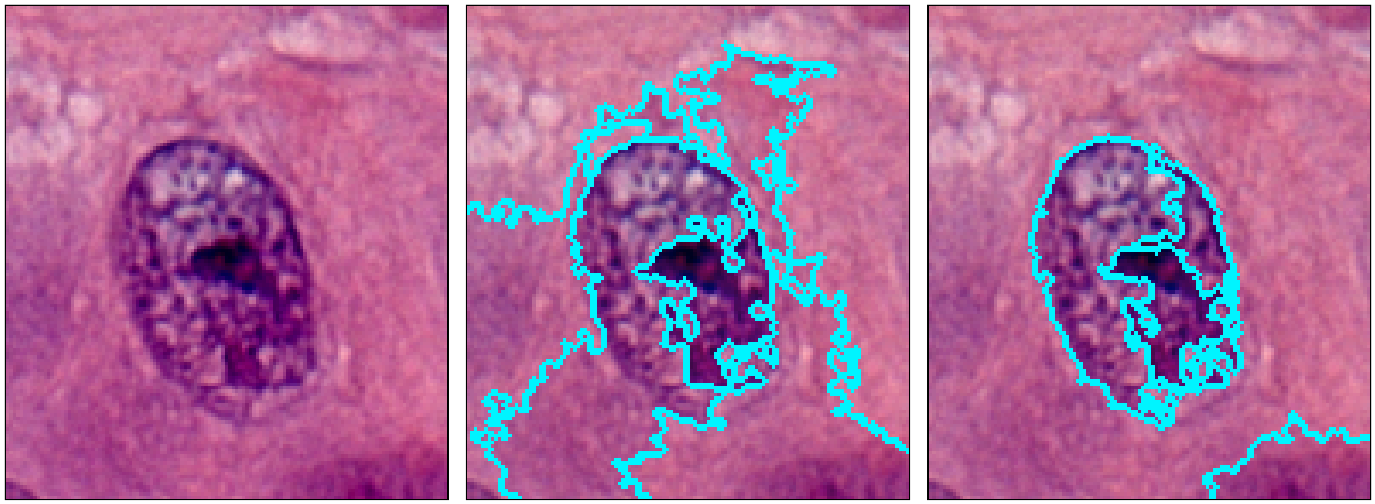

(a) Original (b) ORIG (c) OURS

Fig. 3. Qualitative analysis between ORIG-SICLECOMP and OURS-SICLECOMP when segmenting the nuclei of a cell.

$N_f < 20$, indicating that, although effective for superpixel segmentation, its default criteria ('ORIG') are not suitable for segmentation tasks with few regions. On the other hand, ('OURS') variants managed to improve the results from its predecessors in both datasets, especially concerning superpixel leaking prevention (*i.e.*, UE). In Parasites, BR and UE remained equivalent for both irregularity-driven variants. However, for the compact versions, OURS achieved a significant improvement for both metrics, presenting a performance on pair with the aforementioned methods. Concerning Liver, one can see another improvement achieved by our proposals by achieving almost 80% of delineation accuracy (BR) with only 5 superpixels, against *circa* 69% of accuracy from both original versions of SICLE. Similarly with Parasites, our contributions managed to significantly improve the superpixel leaking prevention especially with few superpixels (5 and 10), achieving almost 90% of BR in our proposals with 5 superpixels, against *circa* 85% from the original version of SICLE-COMP. Concerning Under-segmentation Error (UE), our proposal managed to improve around 2% for Liver in both of our proposals with 5 superpixels. Conversely, a huge improvement occurs in Parasites, mainly between ORIG-COMP and OURS-COMP, with a $2.00 \cdot 10^{-2}$ reduction in our proposal with 5 superpixels.

Figure 3 illustrates a qualitative comparison between the original SICLE variants and our proposed criteria on both datasets. As observed, the original methods tend to generate segmentations with superpixel leakage and imprecise boundaries when constrained to very few superpixels, as shown in Figure 3 (b), where the cell delineation was completely lost. Our approach achieves tighter delineation of the target structures, as shown in Figure 3 (c), where the object was almost perfectly delineated. This visual evidence complements the quantitative results, confirming that the incorporation of domain-specific priors enables SICLE to preserve object boundaries and reduce under-segmentation even in challenging conditions.

IV. Conclusion and Future Works

This work addressed the limitation of standard SICLE criteria in biomedical imaging, in which performance degrades when using a low number of superpixels to isolate a single object. We introduced two novel, domain-specific seed removal criteria: a Position-Based Criterion for medical images and a Color-Based Criterion for pathological images. Our experiments demonstrated that these criteria significantly improve delineation accuracy and reduce under-segmentation error, particularly in situations with low quantity of superpixels. We conclude that incorporating simple domain knowledge into the seed removal step substantially enhances SICLE's utility for targeted biomedical analysis. Future work will focus on designing new criteria for other biomedical tasks that are not highly domain-specific, as in the proposed work, and on investigating hybrid criteria that combine multiple domain-specific features.

Acknowledgment

The authors thank the Pontifícia Universidade Católica de Minas Gerais – PUC-Minas, Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – CAPES – (CAPES STIC-AMSUD 88887.878869/2023-00, 23-STIC-10 and Finance Code 001), the Conselho Nacional de Desenvolvimento Científico e Tecnológico – CNPq (Grants 407242/2021-0, 306573/2022-9, 442950/2023-3) and Fundação de Apoio à Pesquisa do Estado de Minas Gerais – FAPEMIG (Grant APQ-01079-23, APQ-05058-23 and PCE-00301-25).

References

[1] G.-C. Xu, P.-J. Lee, T.-A. Bui, B.-H. Chang, and K.-M. Lee, "Superpixel algorithm for objects tracking in satellite video," in *2021 IEEE International Conference on Consumer Electronics-Taiwan (ICCE-TW)*. IEEE, 2021, pp. 1–2.

[2] J. Long *et al.*, "A graph neural network for superpixel image classification," in *Journal of Physics: Conference Series*, vol. 1871, no. 1. IOP Publishing, 2021, p. 012071.
[3] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "Slic superpixels compared to state-of-the-art superpixel methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 11, pp. 2274–2282, 2012.
[4] X. Wei, Q. Yang, Y. Gong, M.-H. Yang, and N. Ahuja, "Superpixel hierarchy," *IEEE Transactions on Image Processing*, vol. PP, 05 2016.
[5] F. C. Belém, I. B. Barcelos, L. M. João, B. Perret, J. Cousty, S. J. F. Guimarães, and A. X. Falcão, "Novel arc-cost functions and seed relevance estimations for compact and accurate superpixels," *Journal of Mathematical Imaging and Vision*, vol. 65, no. 5, pp. 770–786, 2023.
[6] A. Falcao, J. Stolfi, and R. de Alencar Lotufo, "The image foresting transform: theory, algorithms, and applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 1, pp. 19–29, 2004.
[7] J. E. Vargas-Muñoz, A. S. Chowdhury, E. B. Alexandre, F. L. Galvão, P. A. Vechiatto Miranda, and A. X. Falcão, "An iterative spanning forest framework for superpixel segmentation," *IEEE Transactions on Image Processing*, vol. 28, no. 7, pp. 3477–3489, 2019.
[8] F. Belém, S. J. F. Guimarães, and A. X. Falcão, "Superpixel segmentation by object-based iterative spanning forest," in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*, R. Vera-Rodriguez, J. Fierrez, and A. Morales, Eds. Cham: Springer International Publishing, 2019, pp. 334–341.