



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/243152/>

Version: Accepted Version

---

### Book Section:

Martinelli, J.D., Rodrigues Filho, M.L., Salvagnini, F.C.D.R. et al. (2026) FLIM networks with bag of feature points. In: Chaves, D., Forero Vargas, M. and Rojas Camacho, O., (eds.) Lecture Notes in Computer Science. Lecture Notes in Computer Science, 16529. Springer Nature Switzerland, pp. 268-281. ISBN: 9783032231758.

[https://doi.org/10.1007/978-3-032-23176-5\\_19](https://doi.org/10.1007/978-3-032-23176-5_19)

---

© 2026 The Authors. Except as otherwise noted, this author-accepted version of a paper published in Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

### Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

### Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# FLIM Networks with Bag of Feature Points

João Deltregia Martinelli<sup>1</sup>[0000-0002-1249-5317], Marcelo Luis Rodrigues Filho<sup>1</sup>[0009-0002-3848-4795], Felipe Crispim da Rocha Salvagnini<sup>1</sup>[0000-0002-7896-0058], Gilson Junior Soares<sup>1</sup>[0000-0001-8805-5644], Jefersson A. dos Santos<sup>2</sup>[0000-0002-8889-1586], and Alexandre X. Falcão<sup>1</sup>[0000-0002-2914-5380]

<sup>1</sup> Institute of Computing  
UNICAMP  
Campinas, Brazil

<sup>2</sup> School of Computer Science  
University of Sheffield  
Sheffield, United Kingdom

**Abstract.** Convolutional networks require extensive image annotation, which can be costly and time-consuming. *Feature Learning from Image Markers* (FLIM) tackles this challenge by estimating encoder filters (*i.e.*, kernel weights) from user-drawn markers on discriminative regions of a few representative images without traditional optimization. Such an encoder combined with an *adaptive decoder* comprises a FLIM network fully trained without backpropagation. Prior research has demonstrated their effectiveness in Salient Object Detection (SOD), being significantly lighter than existing lightweight models. This study revisits FLIM SOD and introduces *FLIM-Bag of Feature Points* (FLIM-BoFP), a considerably faster filter estimation method. The previous approach, FLIM-Cluster, derives filters through patch clustering at each encoder’s block, leading to computational overhead and reduced control over filter locations. FLIM-BoFP streamlines this process by performing a single clustering at the input block, creating a bag of feature points, and defining filters directly from mapped feature points across all blocks. The paper evaluates the benefits in efficiency, effectiveness, and generalization of FLIM-BoFP compared to FLIM-Cluster and other state-of-the-art baselines for parasite detection in optical microscopy images.

**Keywords:** First keyword · Second keyword · Another keyword.

## 1 Introduction

Convolutional Neural Networks (CNNs) are essential for image analysis but increasingly rely on large annotated datasets. In domains requiring expert annotation, time constraints make manual labeling costly, tedious, and prone to errors, impacting CNN performance. Training complex CNNs from scratch also demands substantial computational resources, making it impractical in constrained

environments. Moreover, even running off-the-shelf CNNs in inference mode under restrictive scenarios (*e.g.*, mobile devices or low-powered computers in clinical environments of developing countries) might not be possible.

*Feature Learning from Image Markers* (FLIM) is a methodology for creating convolutional encoders without traditional optimization (*i.e.*, backpropagation), deriving filters from patches in discriminative regions marked by users on a few representative images [19]. This approach has led to competitive lightweight models in various applications [18, 4]. A FLIM encoder, combined with an *adaptive decoder*, enables flyweight networks [4]—tens to hundreds of times lighter than existing lightweight models—for *Salient Object Detection* (SOD). Since the adaptive decoder estimates its parameters for each image using a heuristic function, the entire CNN trains without backpropagation from minimal annotation requirements (Figure 1).




Fig. 1: Training pipeline of a FLIM SOD Network from user-drawn markers on three representative images. In this example, the markers are red disks on the object (*Schistosoma Mansoni* eggs) and white disks on the background (impurities and empty region).

*FLIM methodology learns convolutional filters from image patches through clustering.* For each marker, patches are extracted from regions centered at marker pixels, normalized by z-score, and grouped, such that the cluster centers derive the filters of the current block. For the first block, patches come from the input image, while subsequent blocks extract patches from feature maps generated by the previous convolutional block. The markers are mapped to the input of each block for patch extraction, normalization, and clustering.

Cluster centers represent key visual patterns, with filters activating similar patterns at the block’s output. Some approaches refine these clusters using labeled markers to assign filters per class [19]. While effective, per-block clustering increases computational overhead and limits filter placement control, affecting efficiency and explainability as network depth grows.

This paper introduces FLIM-BoFP, a faster and more effective filter estimation method that replaces per-block clustering with a single clustering process at the input layer. It groups patches from marker pixels, storing cluster centers in a "Bag of Feature Points" (BoFP). Feature points are then mapped at each encoder block’s input, where patches are extracted, and filters with biases are assigned per feature point, enabling precise placement and improved activation

control. As shown in Figure 2, FLIM-BoFP eliminates false positives more efficiently than the per-block clustering approach, hereforth called FLIM-Cluster, enhancing object detection in saliency maps across consecutive blocks. Previous work has shown that FLIM saliencies are suitable to initialize post-processing techniques, which take advantage of FLIM object detection capabilities and can improve initial saliency maps to fit the salient object’s border better [18, 17, 15]. Therefore, we also employ an object delineation algorithm that enhances post-processing for FLIM SOD models.




Fig. 2: Progressive saliency maps by decoding consecutive blocks. Figures 2a and 2f show the original and groundtruth images, respectively. Saliency maps obtained by decoding the output of blocks 1-4 for FLIM-BoFP in 2b-2e and FLIM-Cluster in 2g-2j.

For evaluation, we consider a real-world problem that appears in the automated diagnosis of intestinal parasites – the detection of parasite eggs and cysts among impurities. The experiments adopt as baselines SOD models from the state-of-the-art, including FLIM-Cluster, and three parasite datasets – a challenging public dataset containing *Schistosoma Mansoni* eggs<sup>1</sup>, and two private datasets with *Entamoeba histolytica* cysts and *Ancylostoma spp.* eggs. The private datasets are used to evaluate the generalization ability of the models pretrained in the public dataset. Schistosomiasis is a deadly neglected tropical disease caused by *S. Mansoni*. Early diagnosis is crucial but typically relies on slow, manual stool examination [12]. Automated microscopy produces thousands of images per slide, demanding highly efficient solutions. The proposed FLIM-BoFP-based SOD model addresses this need effectively.

Section 2 reviews existing literature on SOD and lightweight networks, including FLIM models. In Section 3, we introduce FLIM-Cluster SOD networks.

<sup>1</sup> <https://github.com/LIDS-Datasets/schistossoma-eggs>

Section 4 presents the FLIM-BoFP method. Sections 5 and 6 showcase our experimental setup and results. We state the conclusion and discuss future work in Section 7.

## 2 Related work

Salient object detection (SOD) focuses on identifying visually prominent objects in images, aiding tasks like object recognition and scene analysis. Traditional methods, inspired by human vision, use local features like texture, color, and orientation in a bottom-up approach to generate saliency maps [20].

Modern SOD uses convolutional encoders to extract multi-level, multi-scale features. Feature maps are combined via Multilayer Perceptrons (MLPs) for classification or convolutional decoders for saliency maps, with FCNs leading the latter. PiCANet employs a U-Net-like structure with a ResNet or VGG encoder [7], BASNet refines boundaries using a ResNet-34 encoder and multi-layer supervision [14], and U2Net enhances contextual information without pretrained backbones [13].

Deep FCNs excel in accuracy but are computationally demanding, leading to lightweight SOD models that balance efficiency and performance. These models apply compression techniques [2], modify convolutions to reduce complexity, and leverage NAS (Neural Architecture Search) for optimal architecture search [3]. Many leverage the MobileNet-V2 encoder, including SeaNet [6] for remote sensing, HVPNet [8] for real-time performance via hierarchical visual perception, and SAMNet [9], which enhances efficiency with stereoscopic multi-scale fusion, avoiding pretrained backbones.

Feature Learning from Image Markers (FLIM) assumes that effective filters can be estimated from markers placed on discriminative regions of a few representative images, such that designing lightweight networks from scratch with minimal data requirements should be possible [19]. Furthermore, combining such an encoder with an adaptive decoder has shown effective for SOD, with flyweight networks considerably more efficient than existing lightweight models [17]. Inspired by these results, the present study improves the effectiveness, efficiency, and explainability of filter estimation from image markers.

## 3 FLIM-Cluster SOD networks

This section introduces the basic concepts of an FLIM-Cluster encoder and adaptive decoder used for the SOD models evaluated in this work. For the sake of space, and because this has already been formally explained in [19], we will briefly summarize the key concepts. We will start by introducing FLIM building blocks and how those concepts expand to define a convolutional encoder directly from an image, where extracted features may be adaptively decoded into saliency maps, eliminating the need for backpropagation.

### 3.1 Images, patches, and labeled markers

A 2D image  $\hat{I} = (D_I, \mathbf{I})$  is represented as a grid of pixels, where each pixel  $p = (x_p, y_p) \in D_I \subset \mathbb{Z}^2$  holds a feature vector with  $m$  channels (e.g., RGB or activation channels), as  $\mathbf{I}(p) = (I_1(p), I_2(p), \dots, I_m(p)) \in \mathbb{R}^m$ . Let  $\mathcal{A}(p)$  express the adjacent set of a pixel  $p$  – a window centered at  $p$  with  $k \times k$  adjacent pixels  $q$ .

$$\mathcal{A}(p) = \{q = (x_q, y_q) \in D_I \mid |x_q - x_p| \leq d \cdot \lfloor k/2 \rfloor, |y_q - y_p| \leq d \cdot \lfloor k/2 \rfloor\} \quad (1)$$

where  $k$  expresses the width and height of the neighborhood and  $d$  is a dilation factor expressing the spacing between adjacent pixels. We can then define a patch  $\mathbf{P}(p) \in \mathbb{R}^{k \times k \times m}$  as the concatenation of the feature vectors  $\mathbf{I}(q) \forall q \in \mathcal{A}(p)$ .

The goal in the FLIM methodology is to learn convolutional filters directly from image patches extracted from marker pixels in discriminative regions. Given a set of representative training images  $\mathcal{D}_T$ , some pixels are manually annotated by "markers"  $\mathcal{M}$  (labeled disks with a small radius). The positions of these markers guide filter learning in each convolutional block, and their labels (background or foreground) define the label of the filter used to create each channel, guiding the parameter estimation of adaptive decoders, which yields a saliency map.

### 3.2 FLIM-Cluster encoder

The FLIM-Cluster encoder is a sequence of *convolutional blocks*, each applying marker-based normalization, convolution with learned filters, activation (ReLU), and pooling (average or max-pooling).

**Marker-based normalization** computes mean and standard deviation of each feature of marker pixels in  $\mathcal{M}$ . Those statistics are used to perform z-score normalization on input training images (first layer) or their feature maps (second layer onwards). The markers on input training images of the first block are mapped to the image domain of their feature maps at the input of the remaining blocks. Let  $\mathcal{P}_C$  be the set of patches  $\mathbf{P}(p)$  for all marker pixels  $p \in \mathcal{M}$  using features at the input of a given block. Marker-based normalization is essentially a z-score normalization of those features in  $\mathcal{P}_C$ . Such an operation, using the statistics of the marker pixels, must be applied to normalize patches of any input image in that block, eliminating the need for bias computation in the block [11].

**Convolution, activation, and pooling** Let  $\hat{I}$  be an image (or feature map) with  $m$  channels after marker-based normalization in a given block. The kernels in the block must have the same shape  $k \times k \times m$ . Hence, the convolution between image  $\hat{I} = (D_I, \mathbf{I})$  and each kernel  $\mathbf{K}_i \in \mathbb{R}^{k \times k \times m}$ ,  $i = 1, 2, \dots, m'$ , of a kernel bank with  $m'$  kernels defines a feature map  $\hat{J} = (D_I, \mathbf{J})$ , where  $\mathbf{J}(p) = (J_1(p), J_2(p), \dots, J_{m'}(p))$  and

$$J_i(p) = \langle \mathbf{P}(p), \mathbf{K}_i \rangle, \quad (2)$$

where  $\mathbf{P}(p)$  is the normalized patch at position  $p \in D_I$ .

Such kernels  $\mathbf{K}_i$  are determined from the training images in  $\mathcal{D}_T$  as follows. Let  $\mathcal{M}_I$  be a marker (disk) in a training image  $\hat{I}$  (or feature map) after marker-based normalization in a given block. Let  $\mathcal{P}_I$  be the set of normalized patches  $\mathbf{P}(p) \in \mathbb{R}^{k \times k \times m}$  extracted from marker pixels  $p \in \mathcal{M}_I$ . FLIM use the K-means algorithm to find  $c$  groups of patches in  $\mathcal{P}_I$ . The center of each group is forced to have a unit norm and used as a kernel  $\mathbf{K}_i$ . By that, each marker generates  $c$  kernels. The procedure is repeated by marker to obtain the total number  $m'$  of kernels in the block. The number  $c$  of kernels per marker is a hyperparameter of the model, together with  $k$ ,  $d$ , pooling type, adjacency, and stride, but the number  $m'$  of kernels in a block will depend on the number of markers drawn by the user. In this work, we fix the same number  $m'$  of kernels in each block.

Learning convolutional filters directly from image patches exploits the geometric properties of convolution operations to identify discriminative visual patterns. From a geometric perspective, each kernel  $\mathbf{K}_i$  is orthogonal to a hyperplane positioned at the origin of  $\mathbb{R}^{k \times k \times m}$ . Simultaneously, normalized image patches  $\mathbf{P}(p)$  correspond to points distributed around the origin of this local feature space. The inner product  $\langle \mathbf{P}(p), \mathbf{K}_i \rangle$  determines the signed distance between each normalized patch and the hyperplane associated with kernel  $\mathbf{K}_i$ . The resulting feature values  $J_i(p)$  are submitted to ReLU activation, eliminating patches on the negative side of the hyperplane. Patches exhibiting visual patterns similar to the learned kernel reside on the positive side of the hyperplane. By that, image regions with patterns of interest should be activated. The choice of kernels from cluster centers at each marker makes them representative visual patterns of interest. After each channel  $J_i(p)$  is submitted to ReLU activation, max/average pooling is applied to the resulting feature map, aggregating nearby activations.

A filter's label, derived from its corresponding marker, determines whether its activation channel highlights the foreground or background in the feature map. Foreground filters brighten objects, while background kernels do the opposite, though separation may be imperfect due to image variability. Adaptive decoders address this issue using a heuristic function to dynamically identify which channels in a feature map can be classified as foreground, background, or undefined channels.

### 3.3 Adaptive decoder

An adaptive decoder may apply point-wise convolution (weighted average) to combine feature map channels, followed by ReLU activation. Some estimate pixel-wise weights per channel before averaging. This work adopts an example of the latter approach, the mean-based decoder from [17].

The mean-based decoder processes a feature map with  $m'$  channels by computing mean activations,  $\mu_F(p)$  and  $\mu_B(p)$ , for foreground and background channels within a small neighborhood of each pixel  $p$ . The method assigns positive weights  $\alpha_i(p)$  to pixels in foreground channels when  $\mu_F(p) > \mu_B(p)$ , as it is expected that object pixels satisfy this condition in foreground channels. Conversely, it assigns negative weights  $\alpha_i(p)$  to pixels in background channels when

$\mu_F(p) < \mu_B(p)$ , as it is expected that background pixels satisfy this condition in background channels. Otherwise, in the absence of evidence about the pixel's label, the weight of the pixel is assigned to zero. Averaging channels with pixel-wise weights and applying ReLU tends to isolate true positives in the saliency map (Figure 2). In Equation 3 [17], channel  $J_i$  belongs to the output feature map  $\hat{J}$  of the last block and  $\lambda(J_i(p)) \in \{1, 2\}$  is the label of kernel  $\mathbf{K}_i$ , where 1 is foreground and 2 is background, as derived from the corresponding labeled marker.

$$\alpha_i(p) = \begin{cases} +1, & \text{if } \lambda(J_i(p)) = 1 \text{ \& } \mu_F(p) > \mu_B(p), \\ -1, & \text{if } \lambda(J_i(p)) = 2 \text{ \& } \mu_F(p) < \mu_B(p), \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

#### 4 FLIM encoders with Bag of Feature Points

The training of a FLIM encoder with BoFP consists of two steps: (i) clustering-based feature point estimation from each marker to identify discriminative locations in each training image  $\hat{I} \in \mathcal{D}_T$ , storing the feature points of all markers in the image into a set  $\mathcal{B}_I$ , and (ii) direct filter estimation from mapped feature points at each block's input. The complete BoFP is the union set  $\mathcal{B}$ .

$$\mathcal{B} = \bigcup_{I \in \mathcal{D}_T} \mathcal{B}_I \quad (4)$$

Figure 3 shows each step in different colors. Step (i) runs once, while step (ii) repeats for each block  $i$ , using the previous block's output as input. The key difference from FLIM-Cluster lies in the control for filter estimation. As set  $\mathcal{B}$  is defined only once, all layers derive convolutional filters from the same spatial locations mapped across blocks. This choice is based on the assumption that if a patch  $\mathbf{P}(p)$  at location  $p$  of the input image defines a discriminative filter, the spatial location  $p$  remains important in subsequent feature maps.

In step (i), for a given number  $c$  of kernels per marker, patches are extracted from marker pixels and grouped into  $c$  clusters by the  $K$ -means algorithm, but with no need for marker-based normalization. Each cluster center identifies a discriminative location  $p$  in a training image  $\hat{I} \in \mathcal{D}_T$ , whose patch  $\mathbf{P}(p)$  is the closest to that center. The location  $p$  is then inserted into the set  $\mathcal{B}_I$ .

In step (ii), the points in  $\mathcal{B}_I$  are mapped to the input of each block for kernel estimation. Let  $\hat{I} = (D_I, \mathbf{I})$  be a training image with  $m$  channels (or a feature map) at the input of a given block. From the feature points  $p$  in  $\mathcal{B}_I$ , FLIM-BoFP extract patches  $\mathbf{P}(p) \in \mathbb{R}^{k \times k \times m}$  to compose a patch set  $\mathcal{P}_B$ . Let  $\mu_B$  and  $\sigma_B$  be the mean and standard deviation of the features in  $\mathcal{P}_B$ . They are used for z-score normalization of those patches. Afterwards, each normalized patch  $\mathbf{P}(p)$

generates a kernel  $\mathbf{K}_i$  with bias  $b_i$  as follows.

$$\mathbf{K}_i = \frac{\mathbf{P}(p)}{\|\mathbf{P}(p)\|} \oslash \sigma_B, \quad (5)$$

$$b_i = -\langle \mu_B, \mathbf{K}_i \rangle, \quad (6)$$

where  $\oslash$  is the element-wise division and  $\langle \cdot \rangle$  is the inner product. The kernel bank with  $m'$  kernels in the block is the union of all kernels  $\mathbf{K}_i$  from all training images  $\hat{I} \in \mathcal{D}_T$ , using all  $m'$  points in  $\mathcal{B}$ . Again, each block has a fixed number  $m'$  of kernels. This procedure eliminates the need for marker-based normalization in each block, since it uses the statistics derived from the patches in  $\mathcal{P}_B$  to create kernels and biases. After convolution with the kernel bank,  $J_i(p) + b_i$  is submitted to ReLU activation, and pooling is applied to obtain the output feature map of the block.




Fig. 3: Training of a FLIM encoder with BoFP begins with a single clustering-based feature point estimation to insert discriminative image locations in the BoFP. Filters for each block are then directly derived from mapped points in the feature map at the block’s input, ensuring precise filter placement and efficient processing.

## 5 Experimental Setup

This section describes the experimental setup to evaluate FLIM-BoFP for parasite detection in optical microscopy images. FLIM SOD networks consist of an encoder trained with FLIM-BoFP or FLIM-Cluster followed by the mean-based adaptive decoder. The encoders contain almost the same operations, except that FLIM-BoFP dismisses marker-based normalization per block since it estimates bias for each filter. Markers are usually strokes or disks. In this work, we use disks with a small radius of 3.0 pixels for both methods to improve control over filter location and, consequently, explainability.

*Parasite datasets* The experiments involve three datasets of parasites with food debris (impurities), such that the first dataset is used for training and testing, and the others are used to evaluate the generalization ability of the models.

1. *S. Mansoni*, a challenging dataset, used for training and evaluation, with 1,219 RGB images ( $400 \times 400$  pixels) containing none, one, or multiple eggs of *Schistosoma Mansoni* [16].
2. *Entamoeba*, a private dataset with 395 RGB images ( $256 \times 256$  pixels) containing a single *Entamoeba Hystolitica* cyst per image.
3. *Ancylostoma*, a private dataset with 320 RGB images ( $400 \times 400$  pixels) containing one or multiple eggs of *Ancylostoma* spp. per image.

The *S. Mansoni* dataset was randomly split into training and testing sets, with each comprising 50% of the data ( $\mathcal{X}_1$  and  $\mathcal{X}_2$ , respectively). The training set  $\mathcal{X}_1$  served multiple purposes: (i) a few representative images in  $\mathcal{X}_1$  were selected for the set  $\mathcal{D}_T \subset \mathcal{X}_1$  used to train and find the best architectures for FLIM-BoFP and FLIM-Cluster, with different images chosen across experimental runs to ensure robust marker-based training; (ii) set  $\mathcal{D}_T \subset \mathcal{X}_1$  is the same used to train/fine-tune the deep-learning models; (iii) the remaining images in  $\mathcal{X}_1 \setminus \mathcal{D}_T$  served as validation set for both FLIM-based and deep-learning models. The testing set  $\mathcal{X}_2$  served for the final performance evaluation of all methods. The remaining datasets were used to evaluate whether models pretrained on *S. Mansoni* can generalize to a different parasite detection task through zero-shot evaluation, representing a transfer learning scenario within the same parasitology application domain.

*Training image selection* The algorithm in [4] was used to iteratively select training examples from  $\mathcal{X}_1$  (*S. Mansoni*), refining the model with each addition while keeping only images that improved performance on the validation set. We start by adding one randomly selected image from  $\mathcal{X}_1$  to  $\mathcal{D}_T$ , then we assess the validation metrics on the images from  $\mathcal{X}_1 \setminus \mathcal{D}_T$ . Among the images for which the method obtained low performance, we select an additional image to compose  $\mathcal{D}_T$ . This procedure is repeated until the user verifies satisfactory performance. FLIM-Cluster was trained with 3-4 images per split, and FLIM-BoFP with 2-3 images per split.

*Compared Methods* FLIM-BoFP-based SOD networks were compared with FLIM-Cluster and deep-learning models: HVPNet [8], SAMNet [9], SeaNet [6], and U2-Net [13], known for strong performance with small training sets. FLIM models were trained from scratch, while the others used recommended hyperparameters, boundary-aware loss, and pretrained backbones (when applicable). Fine-tuning of the deep-learning models utilized the combined representative training images from FLIM-BoFP and FLIM-Cluster, providing 5-6 images in  $\mathcal{D}_T$  per split for training models for *S. Mansoni*. To improve SAMNet and HVPNet’s performance, training was extended to 400 epochs, while SeaNet remained with 50 epochs due to no observed benefit from longer training. U2-Net was trained for

1000 epochs, and the model with the minimum loss value on validation was selected. Lastly, we chose not to evaluate the state-of-the-art iterative segmentation methods, as SAM [5] and MedSAM [10], since those models are not suitable for execution under limited resource scenarios (*e.g.*, low-cost computers commonly found in clinical analysis laboratories of developing countries).

*Post-processing* FLIM SOD networks benefit from post-processing algorithms that enhance object boundary delineation, as they do not employ boundary-aware losses. We use Dynamic Trees (DT) [1] as a competitive solution for fair comparison to baseline models. In DT, saliency maps serve as seed initialization within an optimum-path forest framework. Root pixels are extracted through Otsu thresholding, morphological operations (erosion and dilation), and area filtering — components within the range [1000, 9000] — define internal seeds, while their complement forms external seeds. Such area filters were also used to reduce false positives from the deep-learning models.

**Performance Metrics** To compare the models, we measure their performance using popular metrics such as F-score, Mean Absolute Error (MAE), and weighted F-score (wF). Due to the unique properties of the dataset, namely the typically off-center object locations, some commonly used metrics, such as E-measure, which rely on centered objects, were not considered. Otsu thresholding was used to binarize the saliency maps necessary to calculate the F-score. Metrics were calculated using the three-split average for the *S. Mansoni* dataset for all models. For *Entamoeba* and *Ancylostoma*, we apply the three best split’s architectures found for *S. Mansoni* and average their results to present the metrics of the FLIM models.

## 6 Results and discussion

This section presents an ablation study to determine the optimal architecture for each FLIM model per split and analyze the impact of DT. It includes both quantitative and qualitative results, along with a comparative evaluation of the models.

### 6.1 Ablation study

For FLIM models using  $\mathcal{X}_1$  (*S. Mansoni*), a grid search was conducted across key hyperparameters — kernel size, number of kernels per marker, and number of blocks — to optimize architectures. Kernel sizes between  $3 \times 3$  and  $7 \times 7$  were tested, along with 1 to 4 kernels per marker and up to 4 blocks per split. The best architectures in each split were used for evaluation, with  $3 \times 3$  emerging as the optimal kernel size across all models. FLIM-Cluster performed best with 1 or 3 kernels per marker and up to 2 convolutional blocks. In comparison, FLIM-BoFP used 1 or 2 kernels per marker with 2 or 4 convolutional blocks, making it slightly heavier due to its deeper architecture. The same architectures were




Fig. 4: Average F-measure curves across splits on test sets for each model. Figure 4a compares FLIM models with and without DT post-processing for *S. Mansoni*. Figure 4b contrasts FLIM<sub>DT</sub> models against deep-learning models on *S. Mansoni*, with the FLIM models at the top.

applied to the three datasets. The specific architectures for each encoder and split can be seen in our GitHub (released after review).

The F-measure (F-score with  $\beta = 1$ ) was computed after post-processing in  $P_1$  to select the best FLIM model for each split, identifying BoFP1-BoFP3 and Cluster1-Cluster3 (with DT) as the top performers. Average curves of F-measure (F-score with  $\beta = 0.3$ ) across 255 thresholds on the saliency maps of the test sets show the DT’s positive impact on model robustness (Figure 4a). A significant variation in optimal binarization thresholds suggests that the adaptive decoder enhances saliency distribution across the object of interest. The curves show a superior performance of the FLIM models with DT under limited-training-data conditions compared to the deep-learning networks in both datasets, with the FLIM-BoFP models at the top in performance (Figure 4b).

## 6.2 Comparative analysis

The results in Table 1 reveal a clear contrast in model sizes, with FLIM models significantly smaller than the lightest deep-learning models. Despite more training images and ground-truth-mask-based fine-tuning, deep-learning models were outperformed, with FLIM-BoFP consistently achieving the best metrics (1st or 2nd place). Overfitting likely impacted deep-learning models due to limited training data—SEANet, for instance, generated random noise for images outside its training set. After applying the area filter, SEANet’s saliency maps were entirely nullified (Figure 5), explaining its poor performance across datasets. DT’s effect on FLIM models is evident in larger improvements on average metrics across all splits.

For the *Entamoeba* and *Ancylostoma* datasets, experiments highlight the higher resilience of the FLIM models, particularly FLIM-BoFP, when generalizing to new datasets of the same application domain (Table 2). The 0-shot

experiment shows FLIM-based models were more resilient than deep-learning models —SAMNet, a top performer in *S. mansoni* performed poorly. HVPNet maintained performance better, but still lagged behind FLIM-BoFP and -Cluster. FLIM-BoFP remained the best-performing model in all metrics. Figure 5 demonstrates the qualitative performance and shows the superior capacity of FLIM to generalize to unseen data, in the microscopy domain, even when related to other parasites.

Table 1: The average results from three *S. mansoni* test sets (one per split) for F-score ( $F_\beta$ ,  $\beta = 0.3$ ), Mean Absolute Error (MAE), and Weighted F-measure (wF). Performance is color-coded: best in green, second-best in blue, and worst in red. Arrows  $\uparrow$  and  $\downarrow$  indicate whether higher or lower values are preferable. The "M" and "K" notations represent millions and thousands of trainable parameters, respectively, with FLIM models showing the average number across splits of trainable parameters per model.

| Model                      | # Parameters         | <i>S. mansoni</i>                 |                                   |                                   |
|----------------------------|----------------------|-----------------------------------|-----------------------------------|-----------------------------------|
|                            |                      | $F_\beta\uparrow$                 | MAE $\downarrow$                  | wF $\uparrow$                     |
| SAMNet                     | 1.33 M               | 0.824 $\pm$ 0.120                 | <b>0.006<math>\pm</math>0.000</b> | 0.786 $\pm$ 0.120                 |
| HVPNet                     | 1.23 M               | 0.710 $\pm$ 0.041                 | <b>0.005<math>\pm</math>0.000</b> | 0.668 $\pm$ 0.049                 |
| SEANet                     | 2.76 M               | 0.494 $\pm$ 0.250                 | 0.013 $\pm$ 0.000                 | 0.494 $\pm$ 0.250                 |
| U2-Net                     | 44 M                 | 0.658 $\pm$ 0.151                 | 0.013 $\pm$ 0.005                 | 0.623 $\pm$ 0.148                 |
| FLIM-Cluster               | 5.677 $\pm$ 5.720 K  | <b>0.368<math>\pm</math>0.155</b> | <b>0.014<math>\pm</math>0.000</b> | <b>0.198<math>\pm</math>0.079</b> |
| FLIM-Cluster <sub>DT</sub> | 5.677 $\pm$ 5.720 K  | <b>0.844<math>\pm</math>0.017</b> | <b>0.006<math>\pm</math>0.000</b> | <b>0.837<math>\pm</math>0.015</b> |
| FLIM-BoFP                  | 31.345 $\pm$ 7.888 K | 0.497 $\pm$ 0.190                 | 0.012 $\pm$ 0.000                 | 0.327 $\pm$ 0.128                 |
| FLIM-BoFP <sub>DT</sub>    | 31.345 $\pm$ 7.888 K | <b>0.860<math>\pm</math>0.091</b> | <b>0.006<math>\pm</math>0.000</b> | <b>0.847<math>\pm</math>0.089</b> |

Table 2: The results in *Entamoeba* and *Ancylostoma* as test sets, using none training image. For the FLIM models, these are the average results of the best architectures of each split, as found in *S. Mansoni*.

| Model                      | <i>Entamoeba</i>                  |                                   |                                   | <i>Ancylostoma</i>                |                                   |                                   |
|----------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
|                            | $F_\beta\uparrow$                 | MAE $\downarrow$                  | wF $\uparrow$                     | $F_\beta\uparrow$                 | MAE $\downarrow$                  | wF $\uparrow$                     |
| SAMNet                     | 0.331 $\pm$ 0.176                 | 0.073 $\pm$ 0.001                 | 0.212 $\pm$ 0.094                 | 0.085 $\pm$ 0.079                 | 0.022 $\pm$ 0.002                 | 0.018 $\pm$ 0.015                 |
| HVPNet                     | <b>0.682<math>\pm</math>0.066</b> | 0.080 $\pm$ 0.002                 | 0.256 $\pm$ 0.050                 | 0.286 $\pm$ 0.155                 | 0.012 $\pm$ 0.004                 | 0.152 $\pm$ 0.093                 |
| SEANet                     | <b>0.000<math>\pm</math>0.000</b> | <b>0.086<math>\pm</math>0.001</b> | <b>0.000<math>\pm</math>0.000</b> | <b>0.007<math>\pm</math>0.000</b> | <b>0.506<math>\pm</math>0.038</b> | <b>0.007<math>\pm</math>0.000</b> |
| U2-Net                     | 0.362 $\pm$ 0.098                 | 0.077 $\pm$ 0.004                 | 0.210 $\pm$ 0.060                 | 0.082 $\pm$ 0.047                 | <b>0.013<math>\pm</math>0.008</b> | 0.036 $\pm$ 0.016                 |
| FLIM-Cluster               | 0.596 $\pm$ 0.093                 | 0.078 $\pm$ 0.003                 | 0.215 $\pm$ 0.055                 | 0.362 $\pm$ 0.073                 | <b>0.012<math>\pm</math>0.002</b> | 0.279 $\pm$ 0.031                 |
| FLIM-Cluster <sub>DT</sub> | 0.604 $\pm$ 0.218                 | <b>0.046<math>\pm</math>0.001</b> | <b>0.542<math>\pm</math>0.177</b> | <b>0.456<math>\pm</math>0.036</b> | 0.069 $\pm$ 0.000                 | <b>0.437<math>\pm</math>0.036</b> |
| FLIM-BoFP                  | 0.657 $\pm$ 0.046                 | 0.071 $\pm$ 0.002                 | 0.278 $\pm$ 0.043                 | 0.345 $\pm$ 0.023                 | 0.014 $\pm$ 0.001                 | 0.274 $\pm$ 0.014                 |
| FLIM-BoFP <sub>DT</sub>    | <b>0.792<math>\pm</math>0.125</b> | <b>0.036<math>\pm</math>0.001</b> | <b>0.697<math>\pm</math>0.105</b> | <b>0.517<math>\pm</math>0.010</b> | 0.069 $\pm$ 0.000                 | <b>0.492<math>\pm</math>0.021</b> |




Fig. 5: Qualitative detection results on test images of the *S. Mansoni*, *Entamoeba* and *Ancylostoma* datasets. *GT* is the ground-truth segmentation mask.

## 7 Conclusion

We introduced FLIM-BoFP, a novel kernel estimation technique for FLIM methodology, designed as a more efficient alternative to FLIM-Cluster. FLIM-BoFP demonstrated superior efficiency and effectiveness, outperforming FLIM-Cluster and several heavier deep-learning models. Despite having significantly fewer parameters, it excelled in salient object detection tasks involving microscopy images of parasite eggs and cysts. On the public *S. Mansoni* dataset, FLIM-BoFP surpassed the best deep-learning model, using less than 3% of the trainable parameters. Additionally, FLIM models demonstrated to be more resilient than deep-learning models, when generalizing to datasets of the same application domain, *Entamoeba* and *Ancylostoma*.

We may conclude that FLIM models, particularly the top performer, FLIM-BoFP, are valuable for medical object detection in data- and resource-limited environments. Future work will explore integrating shape-prior morphometric filters, leveraging the combination of multiple-block decoding to enhance saliency maps, and evaluating model generalization on more datasets of parasites. We also intend to examine how user-based image marking impacts performance. FLIM-BoFP’s efficiency also makes it promising for pseudo-labeling large datasets to aid in training heavyweight models and optimizing encoder architectures before training.

**Acknowledgments.** The authors acknowledge grants from FAPESP (2023/14427-8, 2024/08332-7, and 2024/23772-3), CAPES (88887.820891/2023-00), and CNPq (304711/2023-3). The authors also acknowledge grants from the Eldorado Research Institute.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Bragantini, J., Martins, S., Castelo-Fernandez, C., Falcão, A.: Graph-based image segmentation using dynamic trees. In: Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications: 23rd Iberoamerican Congress, CIARP 2018. p. 470–478. Springer-Verlag (2018). [https://doi.org/10.1007/978-3-030-13469-3\\_55](https://doi.org/10.1007/978-3-030-13469-3_55)
2. Chen, F., Li, S., Han, J., Ren, F., Yang, Z.: Review of lightweight deep convolutional neural networks. Archives of Computational Methods in Engineering **31**(4), 1915–1937 (2024). <https://doi.org/10.1007/s11831-023-10032-z>
3. He, X., Zhao, K., Chu, X.: Automl: A survey of the state-of-the-art. Knowledge-based systems **212**, 106622 (2021). <https://doi.org/10.1016/j.knosys.2020.106622>
4. João, L., e Sousa, A.M., dos Santos, B., Guimarães, S., Kijak, E., Gomes, J., Falcão, A.: Building flyweight flim-based cnns with adaptive decoding for object detection (2023), <https://arxiv.org/abs/2306.14840>
5. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3992–4003 (2023). <https://doi.org/10.1109/ICCV51070.2023.00371>
6. Li, G., Liu, Z., Zhang, X., Lin, W.: Lightweight salient object detection in optical remote-sensing images via semantic matching and edge alignment. IEEE Transactions on Geoscience and Remote Sensing **61**, 1–11 (2023). <https://doi.org/10.1109/TGRS.2023.3235717>
7. Liu, N., Han, J., Yang, M.: Picanet: Learning pixel-wise contextual attention for saliency detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3089–3098 (2018)
8. Liu, Y., Gu, Y., Zhang, X., Wang, W., Cheng, M.: Lightweight salient object detection via hierarchical visual perception learning. IEEE Transactions on Cybernetics **51**(9), 4439–4449 (2021). <https://doi.org/10.1109/TCYB.2020.3035613>
9. Liu, Y., Zhang, X., Bian, J., Zhang, L., Cheng, M.: SAMNet: Stereoscopically attentive multi-scale network for lightweight salient object detection. IEEE Transactions on Image Processing **30**, 3804–3814 (2021)
10. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. Nature Communications **15**(1), 654 (1 2024). <https://doi.org/10.1038/s41467-024-44824-z>, <https://doi.org/10.1038/s41467-024-44824-z>
11. de Melo Joao, L., Cerqueira, M.A., Benato, B.C., Falcão, A.X.: Understanding marker-based normalization for flim networks. In: VISIGRAPP : VISAPP (2024), <https://api.semanticscholar.org/CorpusID:268234049>
12. Ogongo, P., Nyakundi, R.K., Chege, G.K., Ochola, L.: The road to elimination: Current state of schistosomiasis research and progress towards the end game. Frontiers in Immunology **Volume 13 - 2022** (2022). <https://doi.org/10.3389/fimmu.2022.846108>
13. Qin, X., Zhang, Z., Huang, C., Dehghan, M., Zaiane, O., Jagersand, M.: U2-net: Going deeper with nested u-structure for salient object detection. Pattern recognition **106**, 107404 (2020). <https://doi.org/https://doi.org/10.1016/j.patcog.2020.107404>
14. Qin, X., Zhang, Z., Huang, C., Gao, C., Dehghan, M., Jagersand, M.: Basnet: Boundary-aware salient object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7479–7489 (2019). <https://doi.org/10.1109/CVPR.2019.00766>

15. Salvagnini, F.C., Gomes, J.F., Santos, C.A.N., Guimarães, S.J.F., Falcão, A.X.: Multi-level cellular automata for flim networks (2025), <https://arxiv.org/abs/2504.11406>
16. Santos, B., Soares, F., Laryssa, S.R.S., Gomes, D., Oliveira, B., Peixinho, A., Suzuki, C., Bresciani, K., Falcão, A., Gomes, J.: TF-Test quantified: a new technique for diagnosis of schistosoma mansoni eggs. *Trop. Med. Int. Health* **24**(5), 586–595 (May 2019)
17. Soares, G., Cerqueira, M., Gomes, J., Najman, L., Guimarães, S., Falcão, A.: Flim-based salient object detection networks with adaptive decoders (2025), <https://arxiv.org/abs/2504.20872>
18. de Souza, I., Benato, B., Falcão, A.: Feature learning from image markers for object delineation. In: 2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI). pp. 116–123 (2020). <https://doi.org/10.1109/SIBGRAPI51738.2020.00024>
19. de Souza, I., Falcão, A.: Learning cnn filters from user-drawn image markers for coconut-tree image classification. *IEEE Geoscience and Remote Sensing Letters* **PP**, 1–5 (09 2020). <https://doi.org/10.1109/LGRS.2020.3020098>
20. Ullah, I., Jian, M., Hussain, S., Guo, J., Yu, H., Wang, X., Yin, Y.: A brief survey of visual saliency detection. *Multimedia Tools and Applications* **79**, 34605–34645 (2020). <https://doi.org/10.1007/s11042-020-08849-y>