



## Texture-driven parametric snakes for semi-automatic image segmentation

Anaïs Badoual<sup>a,\*\*</sup>, Michael Unser<sup>a</sup>, Adrien Depeursinge<sup>b</sup>

<sup>a</sup>Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland

<sup>b</sup>Institute of Information Systems, University of Applied Sciences Western Switzerland (HES-SO), 3960 Sierre, Switzerland

### ABSTRACT

We present a texture-driven parametric snake for semi-automatic segmentation of a single and closed structure in an image. We propose a new energy functional that combines intensity and texture information. The two types of image information are balanced using Fisher's linear discriminant analysis. The framework can be used with any filter-based texture features. The parametric representation of the snake allows for easy and friendly user interaction while the framework can be trained on-the-fly from pixel collections provided by the user. We demonstrate the efficiency of the snake through an extensive validation on synthetic as well as on real data. Additionally, we show that the proposed snake is robust to noise and that it improves the segmentation performance when compared to an intensity-only scheme.

© 2019 Elsevier Ltd. All rights reserved.

### 1. Introduction

In this work, our motivation is to develop a general and versatile framework for semi-automatic segmentation of a single structure of interest in an image, possibly under low-contrast conditions. We want the user to be able to easily specify the desired structure and to modify the outcome when needed.

Most often, structures cannot be fully characterized from their internal distribution of pixel values. Therefore, segmentation methods based on image intensity alone do not perform well on images where the contrast between the object of interest and the background is low. The incorporation of texture information is one complementary way to account for the spatial organization of the pixels inside the desired object (Hoogi et al., 2017; Reska et al., 2015; Sagiv et al., 2006). It allows one to capture the morphological structure of a tissue (Depeursinge and Fageot, 2017).

Active contours, also called “snakes,” are best suited to combine both efficient and well controlled image segmentation with extensive and easy user interaction. They were first proposed by Kass *et al.* (Kass et al., 1987) in 1987 and since became popular tools for segmentation (Delgado-Gonzalo et al., 2015; Blake and Isard, 1998; Bresson et al., 2007; Tang and Acton, 2004;

Heimann and Meinzer, 2009). They consist of an initial user-provided curve configuration that can automatically deform itself to delineate the boundary of the object of interest. The deformation is driven by the minimization of an energy functional. Currently, snakes are described either implicitly (*e.g.*, level sets (Caselles et al., 1997)) or explicitly with point-based and parametric snakes (Cardinale et al., 2012; Brigger et al., 2000; Uhlmann et al., 2016) or, more recently, by subdivision snakes (Badoual et al., 2017). For the energy term, the most common approaches are based on edge or intensity information aggregated from either inside or on the curve (Jacob et al., 2004).

Recent approaches were proposed to incorporate texture information into active contours and level sets (Boonnuk et al., 2015; Pons et al., 2008; Paragios and Deriche, 2002; Lu et al., 2017; Rousson et al., 2003; Wu et al., 2015; Awate et al., 2006). Among them, common characterizations of texture properties were gray-level co-occurrence matrices (Reska et al., 2015; Wu et al., 2015), Gabor filters (Sagiv et al., 2006; Sandberg et al., 2002; Moallem et al., 2016; Wu et al., 2015), sparse texture dictionaries (Gao et al., 2013), variational image decompositions (Wang et al., 2014), or deep learning based on convolutional neural networks (CNN) (Ronneberger et al., 2015; Rupprecht et al., 2016; Ngo et al., 2017; Hu et al., 2017; Hoogi et al., 2017; Yuan et al., 2017). Those methods can be categorized into supervised and unsupervised methods. A limitation of unsupervised approaches, such as in (Wang et al., 2014;

<sup>\*\*</sup>Corresponding author:

*e-mail:* [anaïs.badoual@epfl.ch](mailto:anaïs.badoual@epfl.ch) (Anaïs Badoual)

Rousson et al., 2003; Wu et al., 2015; Awate et al., 2006), is that the incorporation of prior knowledge is difficult. Meanwhile, a limitation of supervised approaches such as CNNs is that they cannot be trained on-the-fly with only a few labeled pixels, as required for natural interactions with snakes. For instance, in the interactive methods of (Reska et al., 2015) and (Gao et al., 2013), texture is learned from the pixels inside the manual initialization of the snake or by providing region boxes for the foreground and the background, respectively. Finally, most texture representations cannot combine local rotational invariance with directional sensitivity, which is sometimes required to accurately identify biomedical tissue structures (Depeursinge and Fageot, 2017). For instance, circular harmonic wavelets (CHWs) (Unser and Chenouard, 2013) provide a powerful tool to model local circular frequencies at multiple scales with invariance to local image rotations.

In this paper, we design a new texture-driven parametric snake for the supervised and semi-automatic segmentation of single and closed structures of interest in images. The choice of a parametric snake as model is motivated by the following: 1) they have a continuously defined spatial representation via the use of basis functions; 2) it is easy to introduce smoothness and shape constraints; 3) they require few parameters (*i.e.*, control points), which results in a faster optimization and better robustness; 4) they allow for friendly user interactions; 5) manual corrections are possible if desired. One known limitation of parametric snakes is their sensitivity to the initialization, which we further discuss in Section 4.2.3. The framework that we propose is based on a new energy term that combines image intensity and texture information. The method is valid for any texture representation yielding response maps and we compare Gabor filters (Bianconi and Fernández, 2007), CHWs (Unser and Chenouard, 2013) and deep feature maps from pre-trained UNets (Ronneberger et al., 2015).

The optimal balance between intensity and texture is learned using Fisher’s linear discriminant analysis (LDA). A very small number of samples provided by the user is sufficient to perform adequate on-the-fly training. We perform a comprehensive performance evaluation of the texture-driven parametric snake on both synthetic and natural images. We measure its robustness and accuracy with respect to noise and initialization, as well as to parameter sensitivity. In addition, we compare our model to supervised and semi-automatic segmentation methods. Regarding the advantages of our method, it is worth noting that a comparison to fully automatic approaches as CNN would not be relevant as they cannot be trained on-the-fly. Overall, our approach allows one to efficiently segment subtle structures in low-contrast images with only a few clicks while allowing a high level of interaction with the user. The method is implemented as a plugin for the platform Icy and is publicly available<sup>1</sup>.

The paper is structured as follows: we describe the proposed framework in Section 2. In Section 3, we discuss implementation details. Finally, Section 4 contains an extensive validation of the proposed snake based on synthetic data, where the

ground truth is known, as well as on real data. Conclusions are drawn in Section 5.

## 2. Framework

The flowchart of the proposed framework is depicted in Figure 1 and discussed next.

### 2.1. Texture Analysis with Filters

We use a set of  $N$  filters to extract texture properties at a given position of the input image  $f$ . We create a sequence  $\{f_k\}_{k \in [1, \dots, K]}$  of  $K = 1 + N$  intensity and texture channels defined by

$$f_k(\mathbf{x}) = \begin{cases} f(\mathbf{x}), & k = 1 \\ |(f * \phi_k)(\mathbf{x})|, & k \neq 1, \end{cases} \quad (1)$$

where  $\mathbf{x} = (x_1, x_2)$  is a coordinate position and  $\phi_k$  is a filter. For a color image in red-green-blue (RGB) representation, we compute the response maps of the red, green, and blue image components. In this case, we have that  $K = 3 \cdot (1 + N)$ .

The proposed method is valid for any collection of filters  $\{\phi_k\}_{k \in [2, \dots, K]}$  extracting texture information. Here after we describe three state-of-the-art filters: CHWs, Gabor filters, and deep feature maps from a UNet.

#### 2.1.1. Circular Harmonic Wavelets

CHWs provide an estimation of the local organization of image directions (LOID) in a rotation-invariant fashion and at a low computational price. The LOID was found to be a fundamental property of structures found in *e.g.* biomedical tissue (Depeursinge, 2017). It allows one to linearly characterize the local circular frequencies, which are at the origin of the success of texture approaches based on local binary patterns (Ojala et al., 2002).

In (1) let  $\phi_k = \phi_{(p,q)}$  be the CHWs of harmonic index  $p = 0, \dots, P - 1$  and scale  $q = 1, \dots, Q$  for  $k = 2, \dots, K$ . The  $N = P \cdot Q$  positive response maps  $|f * \phi_k|$  characterize local circular frequencies in  $f$  up to a maximum harmonic order  $(P - 1)$  and scale  $Q$  (Unser and Chenouard, 2013). They are also locally rotation invariant (Depeursinge et al., 2017). The CHWs are defined in the Fourier domain indexed with polar coordinates  $(\rho, \theta)$  as

$$\hat{\phi}_{(p,q)}(\rho, \theta) = \hat{h}(2^q \rho) \cdot e^{jp\theta}. \quad (2)$$

There,  $\hat{h}$  is a purely radial function that controls the scale profile of the wavelet. We use Simoncelli’s radial wavelet for  $\hat{h}$ , which is expressed by

$$\hat{h}(\rho) = \begin{cases} \cos\left(\frac{\pi}{2} \log_2\left(\frac{2\rho}{\pi}\right)\right), & \frac{\pi}{4} < \rho \leq \pi \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

#### 2.1.2. Gabor Filters

Gabor filter banks allow extracting multi-directional and multi-scale texture information via a systematic parcellation of the Fourier domain with elliptic Gaussian windows (Bianconi and Fernández, 2007). They are not rotation-invariant and are therefore best suited for application where the absolute feature orientation is meaningful. In the spatial domain, Gabor kernels

<sup>1</sup>After acceptance of the paper.

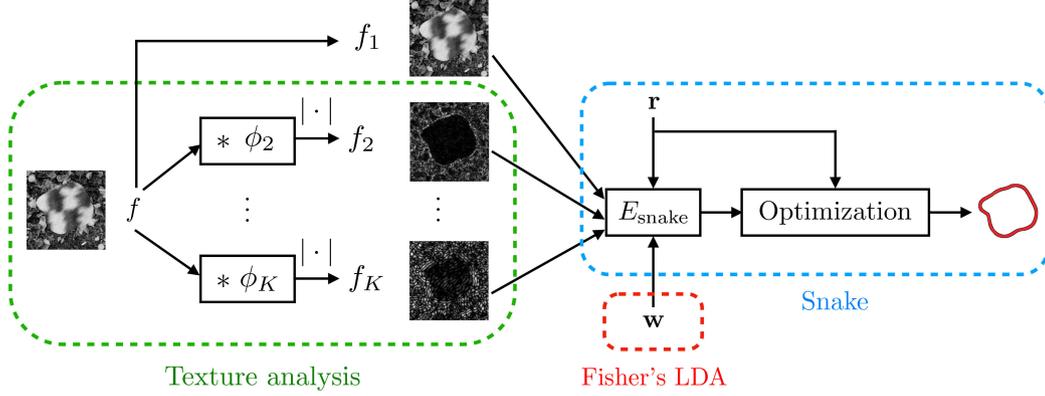


Fig. 1. Flowchart of the proposed framework: a texture analysis is first performed with a bank of filters  $\{\phi_k\}_{k \in [2, \dots, K]}$ . Note that the symbol  $*$  denotes a convolution. Then, the original image  $f_1 = f$  and the resulting positive response maps  $\{f_k = |f * \phi_k|\}_{k \in [2, \dots, K]}$  are balanced using Fisher's linear discriminant analysis. We thus obtain the vector of weights  $\mathbf{w} \in \mathbb{R}^K$ . Finally, the curve  $\mathbf{r}$  of the snake is deformed through the minimization of the region-based energy  $E_{\text{snake}}$ . This term allows for the distinction between homogeneous regions in each channel  $f_k$  weighted by  $w_k$ .

are complex Gaussian-windowed oscillatory functions defined as

$$\phi_{(p,q)}(\mathbf{x}) = \frac{F_q^2}{\pi\sigma_1\sigma_2} e^{-F_q^2\left(\left(\frac{\tilde{x}_{1p}}{\sigma_1}\right)^2 + \left(\frac{\tilde{x}_{2p}}{\sigma_2}\right)^2\right)} e^{j2\pi F_q \tilde{x}_{1p}}, \quad (4)$$

where  $(\tilde{x}_{1p}, \tilde{x}_{2p}) = \mathbf{R}_{\theta_p} \mathbf{x}$  defines the radial and orthoradial elliptic Gaussian axes at the orientation  $\theta_p$  via the rotation matrix  $\mathbf{R}_{\theta_p}$ . In polar Fourier,  $\sigma_1$  and  $\sigma_2$  are the radial and orthoradial standard deviations of the Gaussian window, respectively, and  $F_q$  is the radial position of its center.

We follow the procedure described in (Bianconi and Fernández, 2007) to extract response maps at multiple orientations  $\{\theta_p\}_{p \in [1, \dots, P]}$  and frequencies  $\{F_q\}_{q \in [1, \dots, Q]}$ , where  $\sigma_1$  and  $\sigma_2$  are defined to cover all directions and scales up to the maximum frequency  $F_Q$ .

### 2.1.3. Feature Maps of a Pre-trained UNet

Whereas CHWs and Gabor filters have well-controlled properties and invariances, they may have two limitations. First, they are handcrafted, meaning that they potentially extract image features that are not related to the segmentation task at hand. However, the latter has the advantage of not requiring training examples. Second, they can only characterize low-level features such as circular frequencies or edges and ridges. Since our approach can work with any response maps, it can leverage responses from higher-level deep filters such as used in UNet (Ronneberger et al., 2015). The resulting deep response maps can characterize higher-level features, including an increased complexity of the modeled texture primitives. UNet is the workhorse of deep segmentation methods and achieved state-of-the-art results in many image segmentation applications.

Let  $\{f_k\}_{k \in [2, \dots, 65]}$  be a collection of deep feature maps of a pre-trained UNet. The dimensionality of  $f_k$  is equivalent to the one of the input image. Figure 2 details the architecture of the UNet and the corresponding deep response maps used.

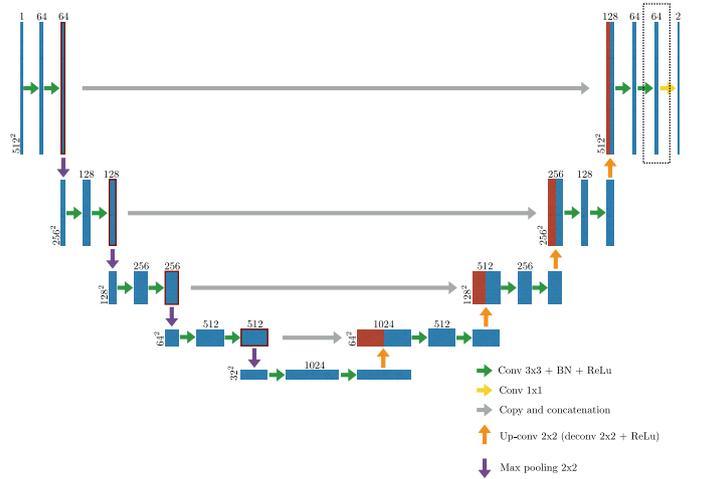


Fig. 2. UNet architecture. Dashed block: the feature maps  $\{f_k\}_{k \in [2, \dots, 65]}$ .

### 2.2. Parametric Snakes

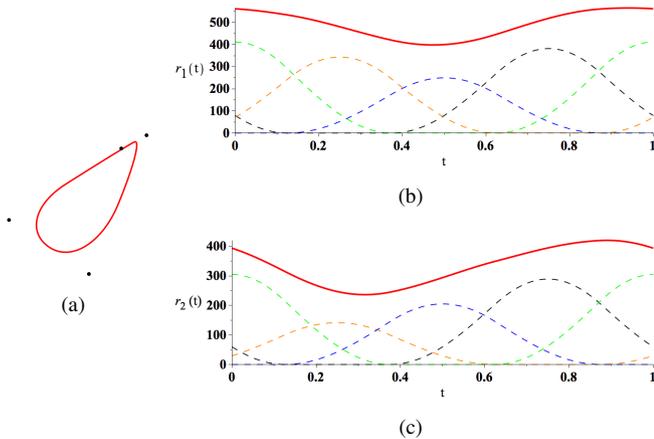
We describe our snake by the closed (*i.e.*, periodic) 2D parametric curve

$$\mathbf{r}(t) = \begin{pmatrix} r_1(t) \\ r_2(t) \end{pmatrix} = \sum_{m=0}^{M-1} \mathbf{c}[m] \varphi_m(t), \quad t \in [0, M], \quad (5)$$

where  $\mathbf{c}[m] = (c_1[m], c_2[m])_{m \in \mathbb{Z}}$  is an  $M$ -periodic sequence of control points such that  $\mathbf{c}[m] = \mathbf{c}[m + M]$ . The function  $\varphi_m$  is defined by

$$\varphi_m(t) = \sum_{n \in \mathbb{Z}} \varphi(t - m - nM), \quad (6)$$

where  $\varphi$  is the exponential B-spline given in (Delgado-Gonzalo et al., 2012) (Equation (8)). This basis function ensures that the snake can perfectly reproduce elliptical shapes using few control points, which is relevant to delineate blob-like objects. In addition, the snake is versatile enough to provide good approximations of any closed curves. The exponential B-spline has a small support, which is advantageous for both computational aspects and the user-interaction (moving one control



**Fig. 3. A parametric curve represented with the exponential B-spline and  $M = 4$  (a) and its coordinate functions (b) and (c). The dots are the control points and the dashed lines are the basis functions. For this plot, we normalized the period to one.**

point affects the structure of the snake locally only). Our model is invariant to affine transformations as  $\varphi$  verifies the partition-of-unity condition defined by  $\sum_{n=-\infty}^{+\infty} \varphi(t-n) = 1$  (Jacob et al., 2001; Unser, 2000). The number  $M$  of control points determines the degree of freedom of the model. A small  $M$  leads to smooth and constrained shapes, while increasing  $M$  brings additional flexibility to approximate arbitrarily complex curves. We show in Figure 3 a parametric curve and its corresponding coordinate functions.

The choice of the energy term is crucial because it drives the evolution of the snake and determines the quality of the segmentation. We propose

$$E_{\text{snake}} = \sum_{k=1}^K w_k E_k, \quad (7)$$

where  $w_k$  is a weight and  $E_k$  is a region-based energy term that allows for the distinction between homogeneous regions in the channel  $f_k$ .

For  $E_k$ , we adopt a strategy similar to (Thévenaz et al., 2011; Delgado-Gonzalo et al., 2012). To discriminate the object from its background, we build a curve  $\mathbf{r}_\lambda$  around the snake  $\mathbf{r}$ , obtained by dilating it by a factor  $\sqrt{2}$  with respect to its center of gravity. The surfaces enclosed by  $\mathbf{r}$  and  $\mathbf{r}_\lambda$ , denoted by  $\Omega$  and  $\Omega_\lambda$ , respectively, are such that  $\Omega \subset \Omega_\lambda$  and  $|\Omega_\lambda \setminus \Omega| = |\Omega|$ . The corresponding energy functional is given by

$$E_k = -\frac{1}{|\Sigma|} \left| \left( \iint_{\Omega} f_k(\mathbf{x}) dx_1 dx_2 - \iint_{\Omega_\lambda \setminus \Omega} f_k(\mathbf{x}) dx_1 dx_2 \right) \right|, \quad (8)$$

where  $\{f_k\}_{k \in [1, \dots, K]}$  is the sequence of images described in Section 2.1 and  $|\Sigma| = \iint_{\Omega} d\mathbf{x}$  is the area enclosed by  $\Omega$ . The global energy  $E_{\text{snake}} = E_{\text{snake}}(\mathbf{c}[m])$  is then minimized by iteratively updating the collection of control points  $\{\mathbf{c}[m]\}_{m \in [0, \dots, M-1]}$  from a starting position. The optimization of the snake is carried out by a Powell-like line-search method (Press et al., 1986).

### 2.3. Fisher's Linear Discriminant Analysis

To set the weights  $\{w_k\}_{k \in [1, \dots, K]}$  in (7), we use Fisher's LDA (Fisher, 1936), which is a supervised technique for dimensionality reduction and classification. We want to identify the classes C (core of the desired object) and B (background). For this purpose, we consider the two set of samples  $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_K(\mathbf{x}))$ , where  $f_k$  is given by (1), for  $\mathbf{x}$  a pixel belonging either to the region of interest (ROI)  $\Omega_C$  or to  $\Omega_B$ . These two ROIs are extracted once during the initialization of the snake, a step that is detailed in Section 3.2. Fisher's LDA seeks the most discriminant hyperplane, characterized by the normal vector  $\mathbf{w}$ , that maximizes the between-class variance while minimizing the within-class variance. The optimal solution is given by (Martínez and Kak, 2001)

$$\mathbf{w} \propto (S_C + S_B)^{-1} (\boldsymbol{\mu}_C - \boldsymbol{\mu}_B), \quad (9)$$

where  $S_C, S_B \in \mathbb{R}^{K \times K}$  are covariance matrices given by

$$S_n = \frac{1}{\#\Omega_n} \sum_{\mathbf{x} \in \Omega_n} (\mathbf{f}(\mathbf{x}) - \boldsymbol{\mu}_n)(\mathbf{f}(\mathbf{x}) - \boldsymbol{\mu}_n)^T, \quad n = B, C \quad (10)$$

and  $\boldsymbol{\mu}_n = \frac{1}{\#\Omega_n} \sum_{\mathbf{x} \in \Omega_n} \mathbf{f}(\mathbf{x})$  is the mean vector of size  $K$  of the class C or B. The number of pixels belonging to  $\Omega_n$  is denoted by  $\#\Omega_n$ . The implicit assumption here is that the texture classes are stationary which justifies the integration over ROI.

In the global framework, Fisher's LDA is trained on-the-fly once during the initialization of the snake. The resulting weights remain then unchanged during the entire optimization process.

## 3. Implementation Details

We implemented the proposed framework as a user-friendly open-source plugin<sup>2</sup> available for the platform Icy (De Chaumont et al., 2012).

### 3.1. Fast Implementation

The framework is implemented according to the theory presented in Section 2. The main computational bottleneck is the evaluation of the surface integrals in (8), which needs to be performed  $K$  times at each iteration of the optimization process. We use Green's theorem to efficiently implement (8) with line integrals as

$$E_k = -\frac{1}{|\Sigma|} \left| \left( 2 \oint_{\mathbf{r}} F_k(\mathbf{r}) dr_2 - \oint_{\mathbf{r}_\lambda} F_k(\mathbf{r}) dr_2 \right) \right|, \quad (11)$$

where

$$F_k(x_1, x_2) = \int_{-\infty}^{x_1} f_k(\tau, x_2) d\tau. \quad (12)$$

Similarly, we have that  $|\Sigma| = \oint_{\mathbf{r}} r_1 dr_2$ . The use of Green's theorem dramatically reduces the computational cost. Moreover, the images  $F_k$  and the weights  $w_k$ , for  $k \in [1, \dots, K]$ , are pre-computed and stored in lookup tables to further accelerate the computation.

<sup>2</sup>A demo of the plugin is available at <http://bigwww.epfl.ch/demo/texture-snake/>, as of August 2018.

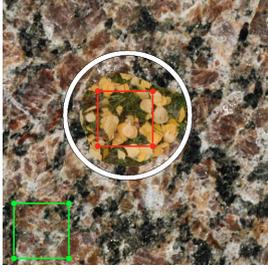


Fig. 4. Extraction of  $\Omega_C$  (red ROI) and  $\Omega_B$  (green ROI) for training Fisher’s LDA.

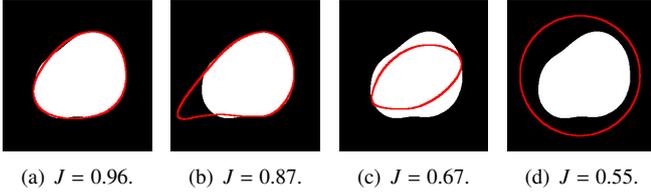


Fig. 5. Illustrations of Jaccard indices.

### 3.2. Supervision of Fisher’s LDA

Two rectangular ROIs  $\Omega_C$  and  $\Omega_B$  (see Section 2.3) are automatically extracted from the initialization of the snake. The first one is localized at the center of gravity of the initialization and the second one outside of the snake (Figure 4). For the LDA to be accurate enough, the number of texture channels should not exceed  $1/10^{\text{th}}$  of the number of pixels used for supervision, which is easily fulfilled in practice. A manual mode is also provided to adjust either one of the ROIs when needed.

## 4. Experiments and Validation

We proceed in three steps to evaluate the performance of the proposed texture-driven parametric snake. First, we test the effect of the parameters on the accuracy of the outcome and study the robustness with respect to initialization and noise. Second, we compare the proposed snake in term of accuracy against other segmentation methods. Third, we illustrate applications on real data. We use the Jaccard ( $J$ ) index to measure the overlap between a segmentation result  $\Omega$  and the corresponding ground truth  $\Omega_{GT}$ . It is defined as

$$J = \frac{|\Omega \cap \Omega_{GT}|}{|\Omega \cup \Omega_{GT}|}. \quad (13)$$

Clearly,  $0 \leq J \leq 1$ , and perfect overlap is described by  $J = 1$ . Texture boundaries are not well defined by definition, according to the uncertainty principle (Petrou and García Sevilla, 2006). This motivated our choice of relatively simple shapes to test the algorithm and the use of the Jaccard index to measure the segmentation accuracy as it does not focus on boundary similarities. Figure 5 illustrates various Jaccard indices to further ease result interpretation. In all following experiments we use CHWs in our method to extract texture information.



Fig. 6. Masks used for the evaluation.

### 4.1. Databases

To validate our model, we created three databases drawn from the real textures of the Prague Texture Segmentation Benchmark<sup>3</sup> (Haindl and Mikes, 2008). Each database was created based on the following pipeline. First, we selected a set of texture classes; then, each texture was combined in pairs using a binary mask of blob-like shape to create an image of  $(512 \times 512)$  pixels. The mask was obtained by thresholding a mixture of several Gaussians with random parameters. To vary the shape to be segmented, we used the five different masks of Figure 6 for each combination. For the first database, called Database 1, we used a set of ten textures of different classes (*e.g.*, wood, stone, flowers). The textures are shown in Figure 7 (a). Database 1, made of 450 images, allows us to test the snake on a diverse set of texture patterns. An image of this database is illustrated in Figure 17. The two other databases were constructed using five textures of the same class. Database 2 is made of the class “wood” and Database 3 of the class “flower”. The corresponding textures are shown in Figure 7 (b) and (c). We use those two databases, made of 100 images each, to study the efficiency of the snake when segmenting similar textures that differ only in subtle ways.

Databases 1, 2, and 3 are made of color images in the RGB representation. One advantage of the proposed snake is that it can handle several channels. In the case of RGB images, it uses both texture and intensity information in every color channel. However, the multichannel information can be predominant over the texture. Typically, textures of Database 3 are very similar (*i.e.*, flowers) but often the color differs. Hence, in order to evaluate the ability of our snake to discriminate textures, as opposed to colors, the validation is performed on both RGB and grayscale versions of the three databases. The grayscale images are obtained by averaging the red, green, and blue channels of the RGB images.

For all experiments, Fisher’s LDA is trained using two fixed ROIs that contain the foreground and background in each mask of Figure 6.

### 4.2. Parameters and Validation of the Model

#### 4.2.1. Degrees of Freedom of the Curve

The number  $M$  of control points is an important parameter of the proposed snake. The choice of  $M$  depends on the application. A large value of  $M$  increases the ability of the snake to approximate intricate shapes but makes the optimization process more complex and penalizes robustness. To illustrate this, we segmented Database 1 for different values of  $M$ , for  $P = 5$

<sup>3</sup>The textures were taken from <http://mosaic.utia.cas.cz/index.php?act=intro>, as of August 2018.

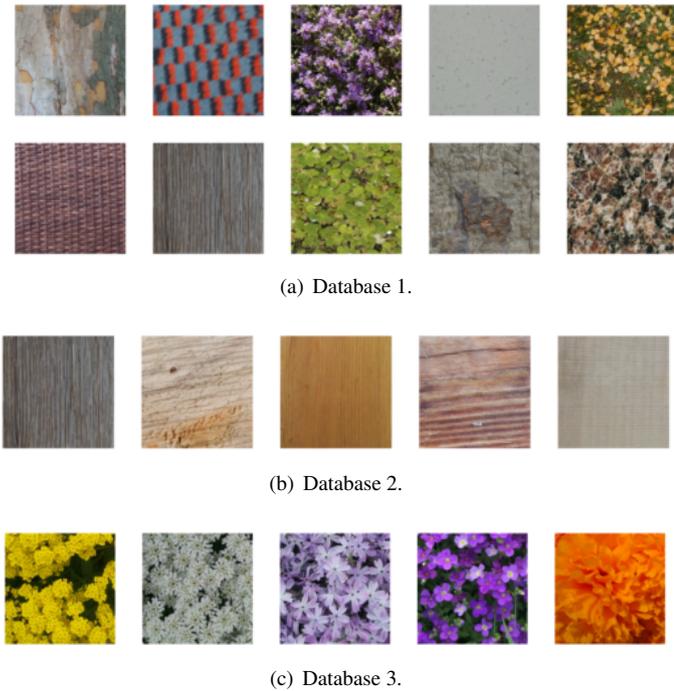


Fig. 7. Textures used for the evaluation.

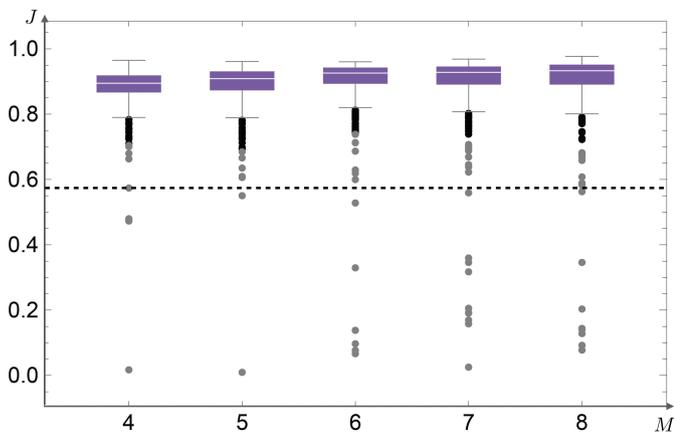


Fig. 8. Segmentation performance on Database 1 (RGB images) according to the number  $M$  of control points. We used 3 scales and 5 harmonics. Dashed line: average Jaccard of the snake’s initialization over the five masks of Figure 6.

and  $Q = 3$ . The corresponding Jaccard indices are reported in Figure 8. The default box spans from the 0.25 quantile to the 0.75 quantile. The dark (grey, respectively) dots are the outliers defined as points beyond 1.5 (3, respectively) times the interquartile range from the edge of the box. We observe that the median increases as  $M$  increases. However, the segmentation becomes less robust as the number of outliers increases. The best tradeoff between accuracy and robustness was found to be  $M = 6$  for Database 1. In fact, keeping  $M$  small acts as a regularizer for the curve.

#### 4.2.2. Influence of $P$ and $Q$ for the CHW Decomposition

We study the impact of the number of harmonics and scales on the accuracy of the segmentation outcome. For fixed  $P$

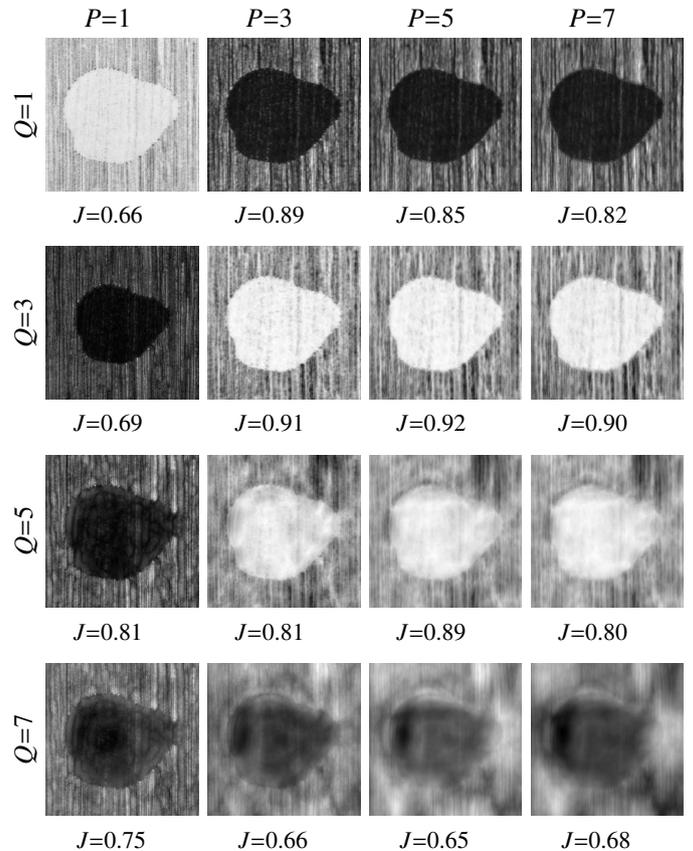


Fig. 9. Illustration of  $f_{\text{snake}}$  for  $P = 1, 3, 5, 7$ , and  $Q = 1, 3, 5, 7$ .

and  $Q$ , we can reconstruct the image  $f_{\text{snake}}$  to generate a two-dimensional projection that estimates what the snake “sees” using

$$f_{\text{snake}} = \sum_{k=1}^K w_k f_k, \quad (14)$$

where  $\{w_k\}_{k \in [1 \dots K]}$  are the weights in (7) estimated with Fisher’s LDA. In Figure 9,  $f_{\text{snake}}$  is shown for different values of  $P$  and  $Q$ , along with their Jaccard index. The original image is a grayscale image of Database 2. The initialization of the snake and the original image are depicted in Figure 10.

We observe that the wavelet scale acts as regularizer. It smooths the textures on  $f_{\text{snake}}$ . At high  $P$  and  $Q$ , the image is less detailed and the snake is less likely to be trapped in local minima but, when  $Q$  is too large, the boundary of the object is not well-defined, which results in an inaccurate segmentation. Increasing the number of harmonics leads to a better discrimination of the two textures. However, more than 5 harmonics yields no more relevant information, resulting in decreased segmentation performance because Fisher’s LDA fails to find adequate separating hyperplanes in spaces with too many dimensions. In addition, it is well-known from statistics of natural images that most information is contained in low frequencies (Hyvarinen et al., 2009).

In a second experiment, Database 2 was segmented using various values of  $P$  and a fixed  $Q$  equal to 3. The results are shown in Figure 11. It can be observed that the accuracy improves from  $P = 1$  to  $P = 5$ , which is even more remarkable on

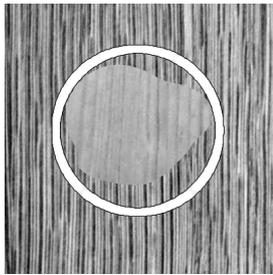


Fig. 10. Initialization on the original image,  $J = 0.62$ .

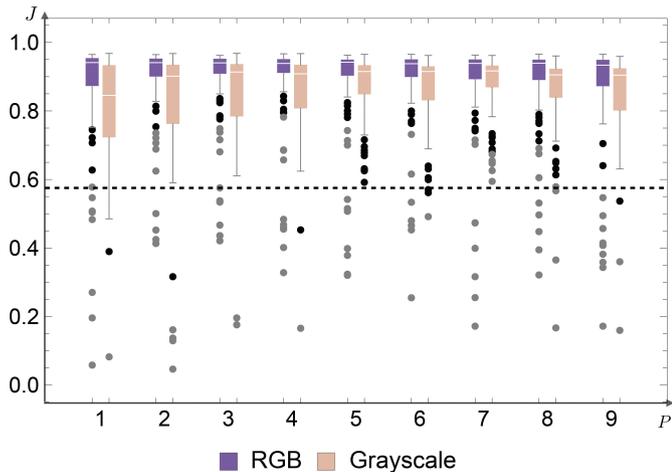


Fig. 11. Segmentation performance on Database 2 for various numbers  $P$  of harmonics. We used 3 scales and 6 control points. Dashed line: average Jaccard of the snake's initialization over the five masks of Figure 6.

grayscale images. Then, the accuracy plateaus and decreases. To conclude, the combination of 5 harmonics with 3 scales provides enough information to discriminate the textures while preserving an accurate segmentation. Hence,  $P = 5$  and  $Q = 3$  were fixed in all following experiments.

#### 4.2.3. Dependence on Initialization

An important aspect is the initial position from which the snake is optimized. Circular shapes for closed snakes are common initial contours. We segmented Database 1 using the five initializations shown in Figure 12. The corresponding Jaccard indices are reported in Figure 13. The best accuracy is obtained for the first two initializations. In fact, the energies  $E_k$ ,  $k \in [1 \dots K]$ , given in (8), are sensitive to the image contrast between the core and the shell of the snake. Hence, the snake should be initialized such that the core intersects the object of interest and the shell intersects the background.

#### 4.2.4. Robustness with Respect to Noise

We investigated the robustness of the texture-driven snake to noise in the image as a function of the number  $M$  of control points. We generated 100 realizations of noisy data for each one of five levels of additive white Gaussian noise. We ran the optimization process until convergence using the proposed texture-driven snake. Signal-to-noise ratio (SNR) corresponding to a given noise level and median Jaccard index were computed. We used a pixelwise SNR that compares the noisy

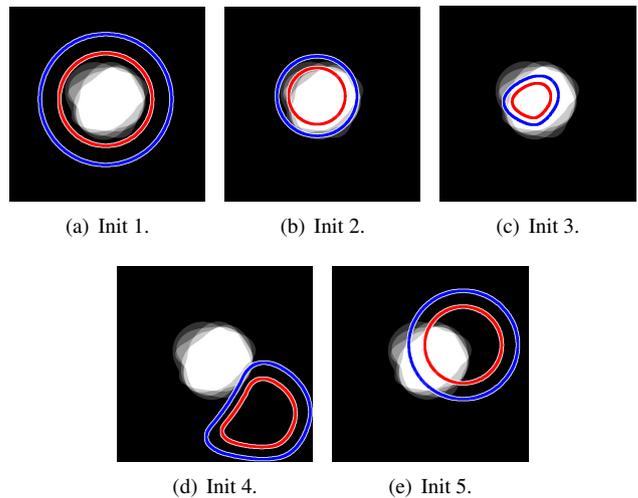


Fig. 12. Various initializations with 6 control points. Inner red circle: snake; Outer blue circle: shell. The initializations are superimposed on the image of the sum of the five masks given in Figure 6.

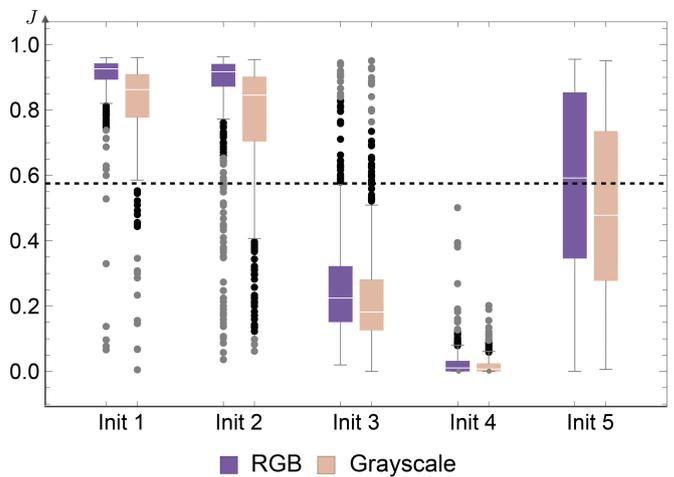


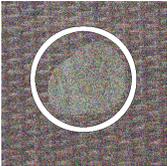
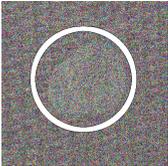
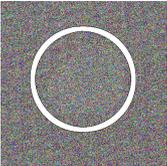
Fig. 13. Segmentation performance on Database 1 for the initializations depicted in Figure 12. We used 6 control points, 3 scales, and 5 harmonics. Dashed line: average Jaccard of the snake's initialization over the five masks of Figure 6.

image and the ground-truth image. The results are summarized in Table 1. The initialization of the snake is overlaid in the thumbnails which depict the noise-corrupted images. Its overlap with the ground truth corresponds to  $J = 0.55$ . From the results, we observe that the texture-driven snake is robust with respect to noise since it is able to give a proper segmentation outcome even for SNRs close to 0dB. This can be explained by the fact that each energy  $E_k$  in (7), for  $k = 1, \dots, K$ , estimates the mean intensity over regions, while Gaussian noise has zero mean. The performance of the snake decreases faster for numerous control points, where higher noise levels induce many local minima.

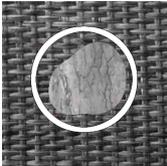
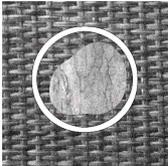
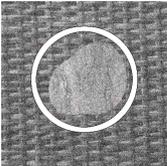
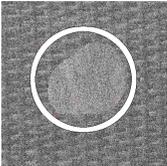
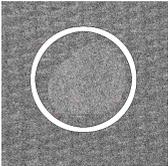
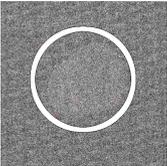
#### 4.2.5. Impact of the Texture Filtering Method

An advantage of our framework is that one can choose the filters that are best suited to his application. We compared the effectiveness of our method when using three different low- and

Table 1. Jaccard indices for the segmentation of noisy data on RGB (top) and grayscale (bottom) images.

						
$M$	SNR = $\infty$ dB	SNR = 14.11dB	SNR = 6.15dB	SNR = 0.13dB	SNR = -3.39dB	SNR = -5.89dB
4	0.90	0.88	0.88	0.81	0.68	0.62
5	0.90	0.91	0.87	0.80	0.71	0.65
6	0.96	0.92	0.88	0.78	0.69	0.64
7	0.95	0.92	0.87	0.78	0.69	0.64
8	0.95	0.90	0.87	0.78	0.68	0.65

						
$M$	SNR = $\infty$ dB	SNR = 14dB	SNR = 6.04dB	SNR = 0.02dB	SNR = -3.5dB	SNR = -6dB
4	0.76	0.87	0.84	0.69	0.64	0.57
5	0.84	0.90	0.81	0.70	0.63	0.59
6	0.96	0.90	0.82	0.68	0.61	0.56
7	0.95	0.90	0.81	0.70	0.63	0.58
8	0.96	0.89	0.84	0.69	0.59	0.57

high-level (*i.e.* deep) filtering methods in the texture analysis: 1) CHWs with 3 scales and 5 harmonics (see Section 2.1.1); 2) Gabor filters with 3 scales and 5 orientations (see Section 2.1.2); 3) 64 feature maps from a pre-trained UNet (see Section 2.1.3). In order to only compare the efficiency of the filters to extract texture information, we equalized mean and variance inside and outside the mask in the grayscale databases 1, 2, and 3. We compared the segmentation performance on the resulting databases for each filtering method. For Database 1, the UNet’s feature maps were obtained by training the network on 40 images made of the two other databases. We proceeded the same way for Database 2 and Database 3. The results of the segmentation are given in Figure 14. We also assessed the statistical significance of pairwise comparisons using a two-tailed paired Wilcoxon rank test. The results are summarized in Table 2. The performances of Gabor filters and CHWs are similar on Database 1, whereas UNet’s feature maps are slightly less discriminative. On Database 2, Gabor filters are significantly more efficient. This is due to the strong and consistent directionality of the textures in this database (Figure 7 (b)), which is efficiently captured by Gabor filters because they are not invariant to image rotations. However, this lack of rotation-invariance explains that their efficiency significantly decreases on Database 3 where the flower petals have different orientations within the same texture class (Figure 7 (c)).

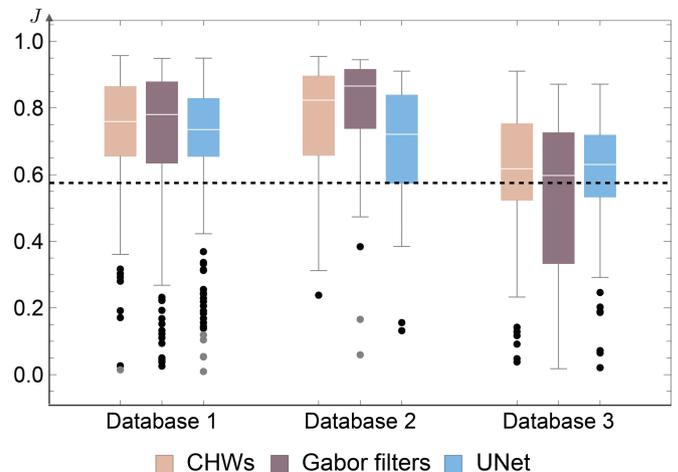


Fig. 14. Segmentation performance obtained with the proposed snake for the three databases when mean and variance were equalized inside and outside the mask. We used 3 scales and 5 harmonics for the CHW decomposition, 3 scales and 5 orientations for the Gabor filters, and 64 feature maps of a UNet. Dashed line: average Jaccard of the snake’s initialization over the five masks of Figure 6.

#### 4.3. Comparisons with Existing Approaches

We carry out two experiments in which we compare the proposed texture-driven snake in term of accuracy against two segmentation methods: 1) the exponential B-spline parametric snake described in (Delgado-Gonzalo et al., 2012). This

**Table 2.**  $p$ -values for pairwise comparisons of texture filtering methods (see Figure 14) assessed using a two-tailed paired Wilcoxon rank test.

	Database 1	Database 2	Database 3
CHW vs. Gabor	0.6397	<b>0.02576</b>	<b>0.00617</b>
CHW vs. UNet	<b>0.009096</b>	<b>3.775E-05</b>	0.8676
Gabor vs. UNet	<b>0.04813</b>	<b>1.06E-10</b>	<b>0.008927</b>

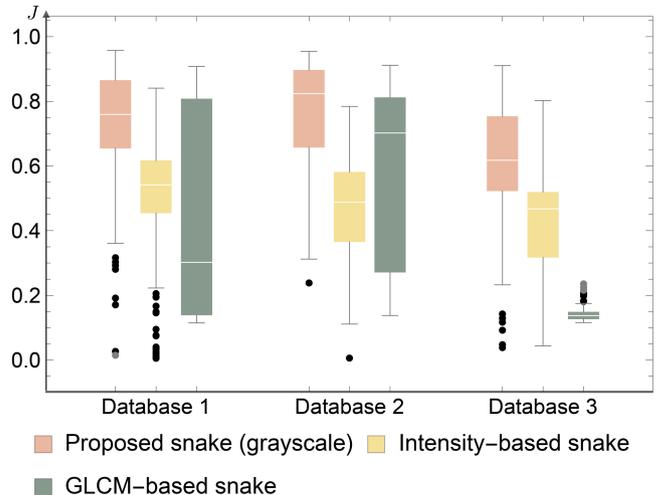
snake has the same reproduction properties and smoothness as the proposed snake but relies on a different region-based energy (intensity information only). The implementation of this method was taken from the free open-source image-processing package Icy (De Chaumont et al., 2012); 2) the texture-based discrete parametric snake described in (Reska et al., 2015). This algorithm generates texture feature maps from gray-level co-occurrence matrices (GLCM) and selects the features that are best suited using a relative standard deviation criteria. We used the implementation given in the platform MESA (Reska et al., 2014). In the following, we refer to those methods as “intensity-based snake” and “GLCM-based snake”, respectively. Similarly to our framework, those two snakes allow for user-interaction and can be trained on-the-fly. We recall that a comparison to fully automatic approaches would not be appropriate since we focus on methods that can be trained on-the-fly with one image. It is worth noting that the two competing methods assume grayscale images.

For all experiments in this section, we initialized the GLCM-based snake with a circle inside the texture of interest, and we set the sensitivity parameter to 3 and enable the option “All angles”. Its overlap with the five masks of Figure 6 corresponds to  $J = 0.53$ . We initialized the other methods with a circle centered<sup>4</sup> on the image and let them evolve automatically until convergence using 6 control points. The Jaccard of this initialization corresponds to  $J = 0.58$ . For our snake we used 3 scales and 5 harmonics.

#### 4.3.1. Mean and Variance Equalization

The goal of this experiment is to emphasize the importance of textural information by illustrating the limitations of the intensity-based snake, and to evaluate the efficiency of our method to detect the desired texture. To only have texture information in the grayscale databases 1, 2, and 3, we equalized mean and variance inside and outside the mask. We optimized the intensity-based, GLCM-based and proposed snakes on each resulting database. The corresponding Jaccard indices are reported in Figure 15. Note that the energy of the GLCM-based snake is based on a sensitivity parameter. Therefore, if the algorithm does not sufficiently discriminates the texture of interest, the snake will spread over the entire image yielding low Jac-

<sup>4</sup>The energy of the GLCM-based snake makes it expand into a region with an uniform texture that is similar to the initial region of the contour. The initialization of this method thus has to be inside the object of interest, whereas for our snake, the initialization has to include the object of interest with a substantial amount of background.



**Fig. 15.** Segmentation performance for the three databases when mean and variance were equalized inside and outside the mask. For the proposed snake we used 6 control points, 3 scales, and 5 harmonics.

card indices. This explains the high standard deviations for this method in Figure 15. We observe that, in each case, the proposed texture-driven snake achieved an adequate segmentation of the object of interest, whereas the intensity-based snake got trapped in local energy minima due to the presence of inhomogeneous regions. Thus, the additional value of texture information is clearly observed. This is reinforced by the GLCM-based snake that yields to a higher maximum Jaccard index than the intensity-based snake on Databases 1 and 2. However, in each database, the GLCM-based snake is less accurate and robust than the proposed snake. The bad result on Database 3 could be explained by the fact that the feature selection algorithm in (Reska et al., 2014) penalizes feature maps with high relative standard deviation, which is not a true discriminative criteria when compared to Fisher’s LDA.

#### 4.3.2. Original Data

In this experiment, we evaluate the segmentation performance on the original databases described in Section 4.1. For comparison purposes, we also provide results obtained with the proposed snake when  $\mathbf{w} = \mathbf{1}$  to investigate the influence of Fisher’s LDA. We compared the final segmentation result to the corresponding ground truth of the synthetic data. The associated Jaccard indices are reported in Figure 16. Illustrations of the segmentation results are shown in Figure 17. We observe that, for each database, we obtain more accurate segmentation outcomes with the proposed texture-driven snake, either on RGB or grayscale images. We also remark that removing Fisher’s LDA from the proposed method (*i.e.*, when  $\mathbf{w} = \mathbf{1}$ ) significantly decreases the performances. This shows the importance of the weights  $\mathbf{w}$  and that Fisher’s LDA is an adequate method to choose them. Finally, the proposed method get better results and robustness when it is applied on RGB images rather than on grayscale images. This is striking for Database 3 and highlights the advantage of our method to be able to deal with different channels.

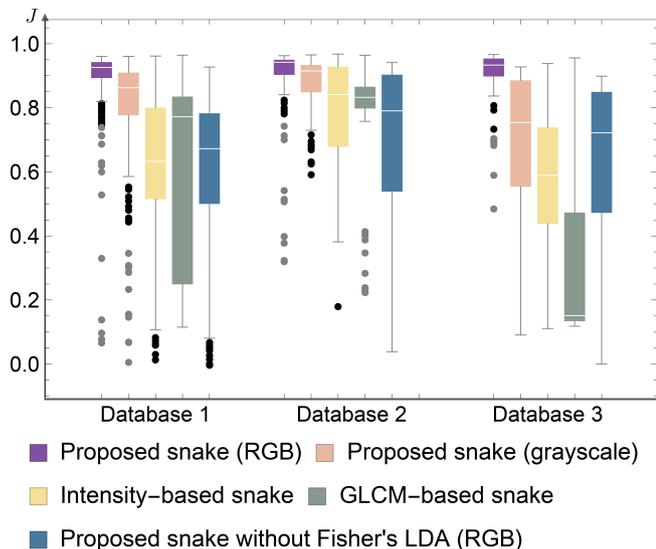


Fig. 16. Segmentation performance for the three databases. Database 1: 450 images; Database 2 and Database 3: 100 images. For the proposed snake we used 6 control points, 3 scales, and 5 harmonics.

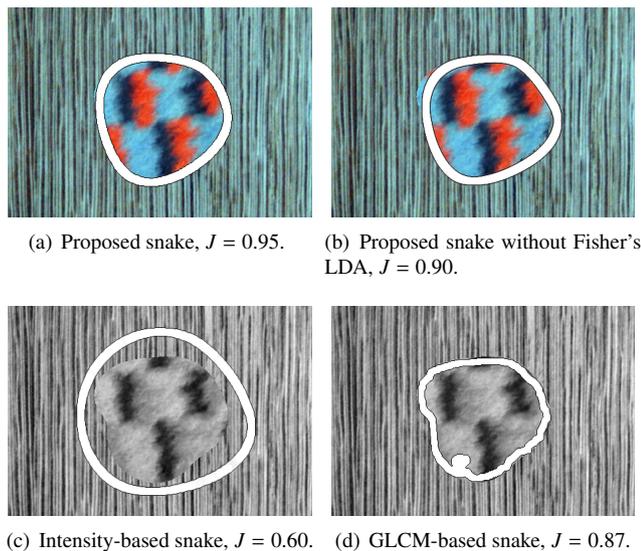


Fig. 17. Segmentation of an image of Database 1. The intensity-based snake and GLCM-based snake are optimized on the grayscale version of the image as they can not handle several channels.

#### 4.4. Real Data Scenarios

We illustrate the behavior of the proposed snake on real data scenarios. For each experiment we manually initialized the snake and let the optimization evolves until convergence for  $P = 5$  and  $Q = 3$ . As user-interaction is one of the main asset of our framework, we locally refined some segmentation outcome by manually moving one or several control points. The average time to segment each image was less than 1.2 seconds on a 1.7 GHz CPU and 8 GB RAM.

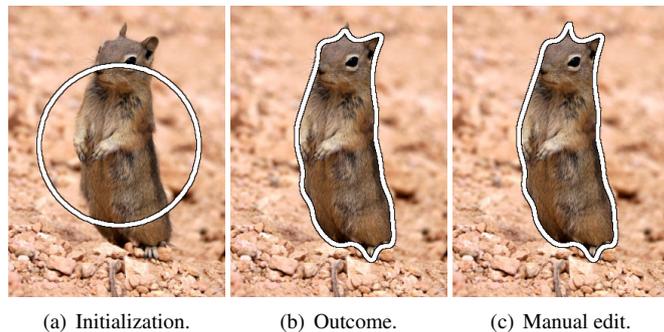


Fig. 18. Segmentation of a squirrel. Source: J de Gier.

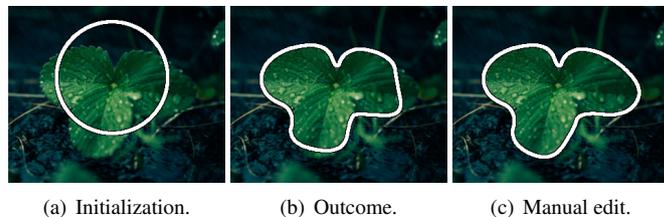


Fig. 19. Segmentation of leaves. Source: Mikael Kristenson.



Fig. 20. Segmentation of a mushroom. Source: Nancy Newton.

##### 4.4.1. Photographic Images

We applied our snake on 4 natural photographs taken from Unsplash<sup>5</sup>, a website dedicated to sharing copyright-free photography. The initializations, segmentation outcomes and manual edits are shown in Figures 18, 19, 20, and 21. The goal of Figure 18 is to show that our method can efficiently segment intricate shapes. Figures 19 and 20 are challenging as the background and the object of interest have similar color. Moreover, the illumination is not uniform which makes the texture more difficult to extract. In Figure 21, small dark areas around the boundaries of the leaf, with a color similar to the background, make the segmentation challenging. As a point of a comparison, we provide in Figure 21 (d) the segmentation outcome obtained with the intensity-based snake.

##### 4.4.2. Biological Images

Texture information is also widely used to characterized biological tissues. We applied our snake to 2 microscopy images

<sup>5</sup>The images were taken from <https://unsplash.com/>, as of September 2018.

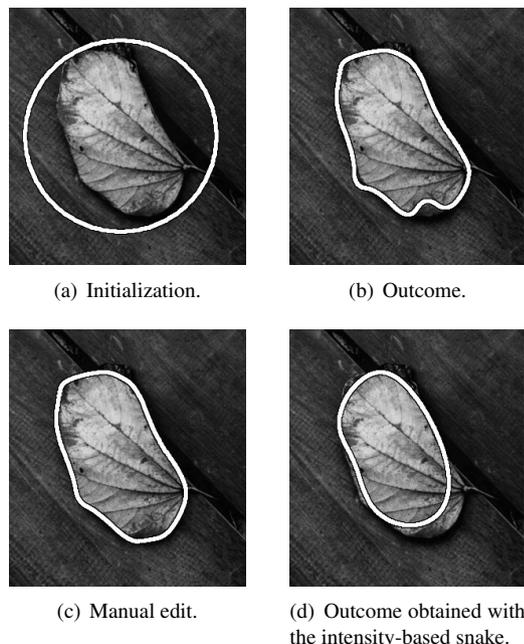


Fig. 21. Segmentation of a leaf. (a)-(c): proposed snake; (d): intensity-based snake. Source: Joshua Newton.

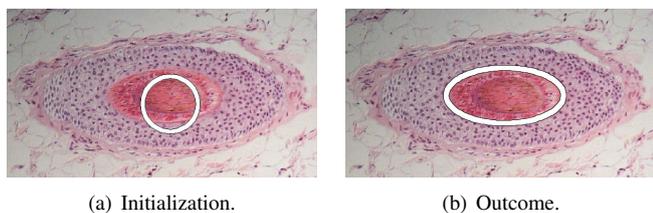


Fig. 22. Segmentation of a hair follicle on a light micrograph. Source: Ivor Mason, 2012, CIL:39094.

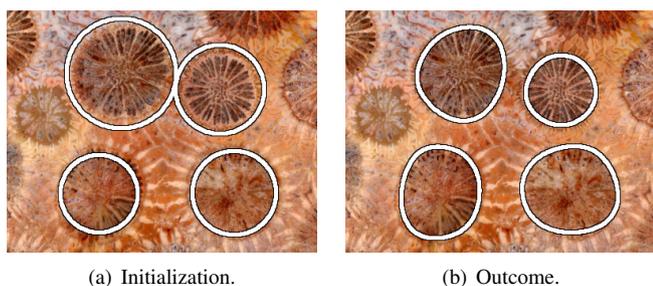


Fig. 23. Segmentation of a fossil of red sponge coral on a microscopy image. Source: Norm Barker, 2009, CIL:41842.

from the Cell Image Library <sup>6</sup>. Those images are challenging as the color inside and outside the structure to segment are similar, and they contain several textures. The initializations and segmentation outcomes are shown in Figures 22 and 23. The qualitative assessment of the segmentation yields satisfactory results.

<sup>6</sup>The images were taken from <http://www.cellimagelibrary.org/>, as of September 2018.

## 5. Conclusions

We have presented a new parametric snake that efficiently allows one to segment structures with similar intensity distribution and low contrast with the background. Our main contribution is the derivation of a new energy that combines intensity and texture information. The contribution of the two types of information is balanced using Fisher's LDA. The method is general and any suited filter banks can be used to extract texture features. This framework is trained on-the-fly from small collections of pixels provided by the user. One main advantage of this method is that one can easily interact with the snake to edit the segmentation outcome when required.

We have compared the performance of our snake to existing ones. In particular, we have observed that the texture-driven snake always performs better than classical parametric snakes that rely on intensity information only. This improvement was even more substantial when the intensity distributions are similar over the background and the object of interest. We have studied the parameter sensitivity of our proposed method as well as its robustness to noise. Finally, we have shown its practical usefulness on real images.

As future work, we plan to extend this framework for the segmentation of 3D and multi-modal data. For instance, we know that a combination of medical image modalities (*e.g.* MRI and CT with various contrasts) can provide complementary information about the texture of a specific tissue (*e.g.* organ, tumor). Hence, as the proposed snake can handle several channels, it could be of interest to use images of different modalities as inputs of the framework.

## Acknowledgment

The authors thank Yann Perret for his contribution. The research leading to these results has received funding from the Swiss National Science Foundation under Grants 200020-162343 and PZ00P2.154891.

## References

- Awate, S.P., Tasdizen, T., Whitaker, R.T., 2006. Unsupervised Texture Segmentation with Nonparametric Neighborhood Statistics, in: Computer Vision – ECCV 2006. Springer, Berlin, Heidelberg, pp. 494–507.
- Badoual, A., Schmitter, D., Uhlmann, V., Unser, M., 2017. Multiresolution subdivision snakes. *IEEE Transactions on Image Processing* 26, 1188–1201.
- Bianconi, F., Fernández, A., 2007. Evaluation of the effects of Gabor filter parameters on texture classification. *Pattern Recognition* 40, 3325–3335.
- Blake, A., Isard, M., 1998. *Active Contours: The Application of Techniques from Graphics, Vision, Control Theory and Statistics to Visual Tracking of Shapes in Motion*. 1st ed., Springer-Verlag New York, Inc.
- Boonnuk, T., Srisuk, S., Sripramong, T., 2015. Texture segmentation using active contour model with edge flow vector. *International Journal of Information and Electronics Engineering* 5, 107–111.
- Bresson, X., Esedoglu, S., Vanderghyest, P., Thiran, J.P., Osher, S., 2007. Fast global minimization of the active contour/snake model. *Journal of Mathematical Imaging and Vision* 28, 151–167.
- Brigger, P., Hoeg, J., Unser, M., 2000. B-Spline snakes: A flexible tool for parametric contour detection. *IEEE Transactions on Image Processing* 9, 1484–1496.
- Cardinale, J., Paul, G., Sbalzarini, I., 2012. Discrete region competition for unknown numbers of connected regions. *IEEE Transactions on Image Processing* 21, 3531–3545.

- Caselles, V., Kimmel, R., Sapiro, G., 1997. Geodesic active contours. *International Journal of Computer Vision* 22, 61–79.
- De Chaumont, F., Dallongeville, S., Chenouard, N., Hervé, N., Pop, S., Provoost, T., Meas-Yedid, V., Pankajakshan, P., Lecomte, T., Le Montagner, Y., Lagache, T., Dufour, A., Olivo-Marin, J.C., 2012. Icy: An open bioimage informatics platform for extended reproducible research. *Nature Methods* 9, 690–696.
- Delgado-Gonzalo, R., Thévenaz, P., Seelamantula, C., Unser, M., 2012. Snakes with an ellipse-reproducing property. *IEEE Transactions on Image Processing* 21, 1258–1271.
- Delgado-Gonzalo, R., Uhlmann, V., Schmitter, D., Unser, M., 2015. Snakes on a plane: A perfect snap for bioimage analysis. *IEEE Signal Processing Magazine* 32, 41–48.
- Depeursing, A., 2017. Multi-scale and multi-directional biomedical texture analysis: Finding the needle in the haystack, in: *Biomedical Texture Analysis: Fundamentals, Applications and Tools*. Elsevier. Elsevier-MICCAI Society Book series.
- Depeursing, A., Fageot, J., 2017. Biomedical texture operators and aggregation functions: A methodological review and user’s guide, in: *Biomedical Texture Analysis: Fundamentals, Applications and Tools*. Elsevier. Elsevier-MICCAI Society Book series, pp. 55–94.
- Depeursing, A., Püspöki, Z., Ward, J.P., Unser, M., 2017. Steerable wavelet machines (SWM): Learning moving frames for texture classification. *IEEE Transactions on Image Processing* 26, 1626–1636.
- Fisher, R.A., 1936. The use of multiple measurements in taxonomic problems. *Annals of Human Genetics* 7, 179–188.
- Gao, Y., Bouix, S., Shenton, M., Tannenbaum, A., 2013. Sparse texture active contour. *IEEE Transactions on Image Processing* 22, 3866–3878.
- Haindl, M., Mikes, S., 2008. Texture segmentation benchmark, in: 2008 19th International Conference on Pattern Recognition, IEEE. pp. 1–4.
- Heimann, T., Meinzer, H.P., 2009. Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis* 13, 543–563.
- Hoogi, A., Subramaniam, A., Veerapaneni, R., Rubin, D.L., 2017. Adaptive estimation of active contour parameters using convolutional neural networks and texture analysis. *IEEE Transactions on Medical Imaging* 36, 781–791.
- Hu, P., Shuai, B., Liu, J., Wang, G., 2017. Deep level sets for salient object detection, in: *Proceedings of the IEEE Conference of Computer Vision and Pattern Recognition (CVPR 2017)*, pp. 2300–2309.
- Hyvarinen, A., Hurri, J., Hoyer, P.O., 2009. *Natural image statistics: a probabilistic approach to early computational vision*. Springer.
- Jacob, M., Blu, T., Unser, M., 2001. An exact method for computing the area moments of wavelet and spline curves, in: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 633–642.
- Jacob, M., Blu, T., Unser, M., 2004. Efficient energies and algorithms for parametric snakes. *IEEE Transactions on Image Processing* 13, 1231–1244.
- Kass, M., Witkin, A., Terzopoulos, D., 1987. Snakes: Active contour models. *International Journal of Computer Vision* 1, 321–331.
- Lu, J., Wang, G., Pan, Z., 2017. Nonlocal active contour model for texture segmentation. *Multimedia Tools and Applications* 76, 10991–11001.
- Martínez, A.M., Kak, A.C., 2001. PCA versus LDA. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 228–233.
- Moallem, P., Tahvilian, H., Monadjemi, S.A., 2016. Parametric active contour model using Gabor balloon energy for texture segmentation. *Signal, Image and Video Processing* 10, 351–358.
- Ngo, T.A., Lu, Z., Carneiro, G., 2017. Combining deep learning and level set for the automated segmentation of the left ventricle of the heart from cardiac cine magnetic resonance. *Medical Image Analysis* 35, 159–171.
- Ojala, T., Pietikäinen, M., Mäenpää, T., 2002. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 971–987.
- Paragios, N., Deriche, R., 2002. Geodesic active regions and level set methods for supervised texture segmentation. *International Journal of Computer Vision* 46, 223–247.
- Petrou, M., García Sevilla, P., 2006. *Non-stationary grey texture images, in: Image processing: dealing with texture*. John Wiley & Sons, pp. 297–606.
- Pons, S.V., Rodríguez, J.L.G., Pérez, O.L.V., 2008. Active contour algorithm for texture segmentation using a texture feature set, in: 19th IEEE International Conference on Pattern Recognition (ICPR 2008), pp. 1–4.
- Press, W., Teukolsky, S., Vetterling, W., Flannery, B., 1986. *Numerical Recipes: The Art of Scientific Computing*. Third ed., Cambridge University Press.
- Reska, D., Boldak, C., Kretowski, M., 2015. A texture-based energy for active contour image segmentation, in: *Image Processing & Communications Challenges* 6, pp. 187–194.
- Reska, D., Jurczuk, K., Boldak, C., Kretowski, M., 2014. Mesa: Complete approach for design and evaluation of segmentation methods using real and simulated tomographic images. *Biocybernetics and Biomedical Engineering* 34, 146–158.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Springer International Publishing. volume 9351 of *Lecture Notes in Computer Science*, pp. 234–241. doi:10.1007/978-3-319-24574-4\_28, arXiv:1505.04597.
- Rousson, M., Brox, T., Deriche, R., 2003. Active unsupervised texture segmentation on a diffusion based feature space, in: 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings., IEEE Comput. Soc. pp. II–699–704.
- Rupprecht, C., Huaroc, E., Baust, M., Navab, N., 2016. Deep active contours. arXiv preprint arXiv:1607.05074 .
- Sagiv, C., Sochen, N.A., Zeevi, Y.Y., 2006. Integrated active contours for texture segmentation. *IEEE Transactions on Image Processing* 15, 1633–1646.
- Sandberg, B., Chan, T., Vese, L., 2002. A level-set and Gabor-based active contour algorithm for segmenting textured images, in: *UCLA Department of Mathematics CAM report*, pp. 1–2.
- Tang, J., Acton, S., 2004. Vessel boundary tracking for intravital microscopy via multiscale gradient vector flow snakes. *IEEE Transactions on Biomedical Engineering* 51, 316–324.
- Thévenaz, P., Delgado-Gonzalo, R., Unser, M., 2011. The ovusculc. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 382–393.
- Uhlmann, V., Fageot, J., Unser, M., 2016. Hermite snakes with control of tangents. *IEEE Transactions on Image Processing* 25, 2803–2816.
- Unser, M., 2000. Sampling—50 Years after Shannon. *Proceedings of the IEEE* 88, 569–587.
- Unser, M., Chenouard, N., 2013. A unifying parametric framework for 2D steerable wavelet transforms. *SIAM Journal on Imaging Sciences* 6, 102–135.
- Wang, G., Pan, Z., Dong, Q., Zhao, X., Zhang, Z., Duan, J., 2014. Unsupervised texture segmentation using active contour model and oscillating information. *Journal of Applied Mathematics* 2014.
- Wu, Q., Gan, Y., Lin, B., Zhang, Q., Chang, H., 2015. An active contour model based on fused texture features for image segmentation. *Neurocomputing* 151, 1133–1141.
- Yuan, Y., Chao, M., Lo, Y.C., 2017. Automatic skin lesion segmentation using deep fully convolutional networks with Jaccard distance. *IEEE Transactions on Medical Imaging* 36, 1876–1886.