Segmentation of Skin Tumors in High-Frequency 3-D Ultrasound Images

Bruno Sciolla, Lester Cowell, Thibaut Dambry, Benoit Guibert, Philippe Delachartre

▶ To cite this version:

Bruno Sciolla, Lester Cowell, Thibaut Dambry, Benoit Guibert, Philippe Delachartre. Segmentation of Skin Tumors in High-Frequency 3-D Ultrasound Images. Ultrasound in Medicine & Biology, 2017, 43 (1), pp.227 - 238. 10.1016/j.ultrasmedbio.2016.08.029 . hal-01684323

HAL Id: hal-01684323
https://hal.science/hal-01684323

Submitted on 31 Jan 2024

Segmentation of Skin Tumors in High-Frequency 3D Ultrasound Images

Bruno Sciolla[a,*], Lester Cowell[b], Thibaut Dambry[c], Benoît Guibert[c], Philippe Delachartre[a]

[a] *CREATIS laboratory, INSA Lyon, France*
[b] *Melanoma Skin Cancer Clinic, Hamilton Hill, Australia*
[c] *Atys Medical, Soucieu-en-Jarrest, France*

Abstract

High-frequency 3D ultrasound imaging is an informative tool for diagnosis, surgery planning and skin lesions examination. The purpose of this article is to present a semi-automated segmentation tool, giving easy access to the extents, shape and volume of a lesion. We propose an adaptive log-likelihood level-set segmentation procedure using nonparametric estimates of the intensity distribution. The algorithm has a single parameter to control the smoothness of the contour, and we show that a fixed value yields satisfactory segmentation results with an average Dice coefficient of $D = 0.76$. The algorithm is implemented on a grid, which increases the speed by a factor of 100 compared with a standard pixel-wise segmentation. We compare the method with parametric methods making the hypothesis of Rayleigh or Nakagami distributed signals, and show that our method has greater robustness with similar computational speed. Benchmarks are made on realistic synthetic ultrasound images and a dataset of 9 clinical 3D images acquired with a 50 MHz imaging system. The proposed algorithm is suitable for use in a clinical context as a post-processing tool.

Keywords: 3D imaging, high-frequency imaging, segmentation, tumor segmentation, level-set, nonparametric method, grid, adaptive method

---

[*] Corresponding Author: Bruno Sciolla, CREATIS, Bat. 502, 7 av. Jean Capelle, 69100 Villeurbanne, France; Email, bruno.sciolla@creatis.insa-lyon.fr; Phone, +33472436469

Preprint submitted to Ultrasound in Medicine and Biology *September 7, 2016*

Introduction

High-frequency 3D ultrasound is a non-destructive, fast, readily available and affordable means to determine the size, structure and echoic properties of skin lesions. This imaging modality is useful to assist in the diagnosis, surgery planning and monitoring of benign or malignant tumors of the skin, such as naevus, melanoma and basal cell carcinoma (Schmid-Wendtner and Dill-Müller (2008); Machet et al. (2010)).

High-frequency ultrasound imaging systems are undergoing continuous development and have broader applications in dermatology, cosmetics and small-animal imaging (Vogt and Ermert (2013); Maev (2014)). In dermatology applications, there is a need for dedicated algorithms to perform automatic segmentation of a lesion, in a 3D ultrasound image, in order for the physician to quickly apprehend the extent of a lesion. Previous works related to dermatology applications include the segmentation of the epidermis tissues in histological images (Xu and Mandal (2015)) and of breast skin in conventional ultrasound (Gao et al. (2013)). A texture-based method to segment skin tumors in high-frequency ultrasound images has been proposed (Olivier and Paulhac (2011)), but has not been validated on clinical data, and the method has a significant computational load. In the following we propose an approach based on the intensity distribution only. The segmentation is performed in less than one minute, which is short enough for convenient use in clinical routine.

In ultrasound images, many efficient segmentation methods are based on the intensity distribution of the envelope, or the log-compressed envelope, to discriminate between the different regions of interest. Commonly used methods are based on a parametric estimate of the intensity distribution, assuming either a Rayleigh distribution (Sarti et al. (2005); Slabaugh et al. (2006); Pereyra et al. (2012, 2015)), or other distributions such as exponential and Gamma distributions (Anquez et al. (2013); Dahdouh et al. (2015)), 2-parameters Gamma distributions (Rahmati et al. (2012)) and log-compressed Gaussian distributions (Pons et al. (2011)).

However, in the context of high-frequency ultrasound, the high resolution implies a low number of scatterers within the typical resolved volume. As a result, complex distributions like the K-distributions (Destrempes and Cloutier (2010)) or mixtures of distributions are required. Yet it is technically difficult to introduce complex distributions in the level-set framework. To circumvent this problem, we propose using a non-parametric estimate of intensity distributions, making no assumption on the speckle pattern distribution.

Due to the variability of shapes in skin tumors, herein we consider only level-set segmentation methods for their powerful capacity to represent arbitrary smooth shapes. Additionally, the boundary of tissues must be regularized to remove structures due only to the speckle patterns, which level-set segmentation does well. On the other hand, the size of the images is of 300 × 300 × 832 voxels, or 74.88 million voxels. A standard level-set segmentation may take up to half an hour with such images; therefore we would like to propose a method to reduce the processing time.

To handle these different problems, we propose an Adaptive Log-Likelihood (AdLL) 3D level-set segmentation algorithm, which maximizes the log-likelihood of a contour. Our goal is to make a postprocessing tool to be used on a standard computer. To increase the speed of segmentation, we use a coarse-grid level-set implementation.

In a previous conference proceedings (Sciolla et al. (2015a)), we have presented early results on a non-parametric log-likelihood method. The method proposed here is different, because of the new adaptive level-set solver. Unlike the previous study, this article includes the Nakagami parametric segmentation, contains a systematic and a quantitative discussion on realistic synthetic images, and validates the segmentation on a dataset of 9 clinical lesions.

There are several non-parametric methods which have been proposed. One example is the Bhattacharyya distance, recently used in ventricle segmentation in brain ultrasound images (Qiu et al. (2015)). The method proposed here is based on the maximum log-likelihood estimate, in light of a previous study (Sciolla et al. (2015b)) which indicates a greater robustness of the log-likelihood method compared with the Bhattacharyya coefficient or the Hellinger distance.

The article is structured as follows. We first describe the statistics of high-frequency skin images and explain the choice of the nonparametric method. Then, we describe the proposed (AdLL) method, the non-parametric log-likelihood optimization problem and the level-set algorithm. In the Results section, the impact on speed and accuracy of using coarse grids is evaluated. The parametric and non-parametric methods are then compared on realistic synthetic images. Finally, the performance of the different algorithms are evaluated on the clinical data.

Segmentation method

Intensity distribution in high-frequency ultrasound images

Fig. 1(a) shows slices of a high-frequency ultrasound 3D image of the skin and a lesion. The epidermis, dermis and subcutaneous tissues are visible, as well as a hypoechoic tumor. Under the hypothesis of a homogeneous medium with a large density of scatterers, the envelope signal is expected to follow a Rayleigh distribution (Destrempes and Cloutier (2010)). This is the case of the subcutaneous tissue distribution shown in Fig. 1(b). However, this hypothesis is often not met in high-frequency images of the skin, as shown in Fig. 1(c). As a whole, the background region surrounding the lesion is composed of several tissues, including the dermis, epidermis and subcutaneous tissues. The global distribution of these regions, shown in Fig. 1(d), can be bimodal, and is therefore not well-fitted by any standard parametric distributions such as Rice, Nakagami or K-distributions (Destrempes and Cloutier (2010)).

In order to circumvent this problem, we propose to use a method based on non-parametric intensity estimates, which overcomes this problem with little additional cost. Since no hypothesis has been made on the distribution of the tissue, the method is well-founded to segment inhomogeneous tissues with arbitrary distributions, which is suitable for our application.

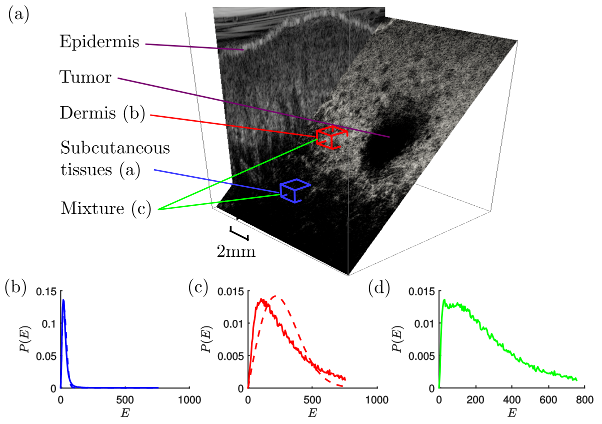

Figure 1: **(a)** 3D image of dermis with a tumor lesion, with vertical and diagonal slices. **(b-d)** Distribution $P(E)$ of the envelope E in different tissues surrounding a tumor. **(b)** In subcutaneous tissues. Dashed line: fit to a Rayleigh distribution. **(c)** In the dermis. Dashed line: fit to a Rayleigh distribution, showing large deviations from a Rayleigh law. **(d)** Mixture of the two regions: the distribution is bimodal and cannot be fitted to traditional parametric distributions proposed in ultrasound (Cardinal et al. (2006)).

Processing chain

The segmentation procedure takes as input the log-compressed envelope of the radiofrequency signal, as provided by standard ultrasound acquisition systems. In this study, the raw radiofrequency signal is also available, allowing us to compare the proposed method with parametric segmentation methods which rely on the envelope signal. The processing chain is sketched in Fig. 2. The segmentation is semi-automatic, since it requires the user to specify a volume of interest and a contour initialization with a box shape. The user is requested to enter 3 non-aligned points or more, and the smallest enclosing box is chosen as initialization. The next step is the segmentation itself, it is not performed at the level of individual voxels but rather on a coarser grid, allowing significant speedup of the process. Finally the segmentation is obtained at the original resolution after interpolation and smoothing.

Segmentation method

The image is split into two regions Ω_A and Ω_B, with boundary surface $\mathcal{S}$. In the level-set formalism (Chan and Shen (2005)), the region Ω_A is delimited by the zeros of a level-set function $\phi_{\mathbf{x}} = \phi(\mathbf{x})$, such that $\Omega_A = \{\mathbf{x}|\phi_{\mathbf{x}} \geq 0\}$ and $\Omega_B = \{\mathbf{x}|\phi_{\mathbf{x}} < 0\}$

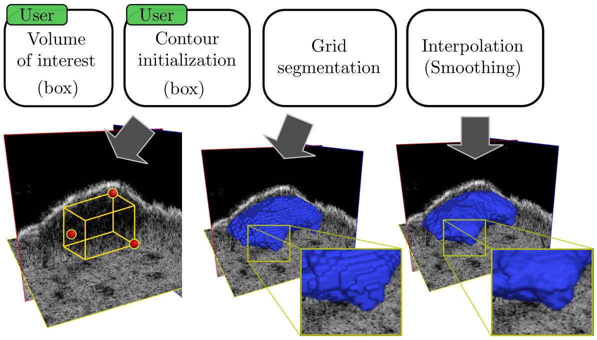

Figure 2: Processing chain for the tumor segmentation. A volume of interest is first selected by the user, and points of initialization (red spheres) are specified by the user. The segmentation is initialized as the smallest box enclosing all points (yellow frame box). Then segmentation is performed on a coarse grid. Finally, the contour is interpolated on the original image resolution and smoothed.

where $\mathbf{x}$ spans the 3D image coordinates. The segmentation problem is reduced to an optimization problem, minimizing the cost function:

$$E[\phi] = \alpha E^{\mathrm{reg}}[\phi] + \beta E^{\mathrm{DA}}[\phi] \tag{1}$$

The cost function contains a data attachment term with coefficient β and a regularization term with coefficient α ensuring the smoothness of the contour and guaranteeing that the minimization is a well-posed problem. Only the ratio β/α is relevant for minimizing $E[\phi]$ but we keep the parameter α for future convenience. We choose a regularization term $E^{\mathrm{reg}}[\phi]$ equal to the area of the boundary surface $\mathrm{Surf}(\mathcal{S}) = \int_{\mathcal{S}} \mathrm{d}^2\mathbf{x}$.

Log-likelihood cost function

The data attachment term $E^{\mathrm{DA}}[\phi]$ is chosen to maximize the likelihood that the image is split into two regions with different distributions. The likelihood of the hypothesis $H_1 = \{$ the distribution of intensities I_x in the two regions Ω_A and Ω_B are i.i.d. random variables drawn from two distributions P_A and P_B $\}$ is compared to the likelihood of the null hypothesis $H_0 = \{$ all image intensities I_x are i.i.d. variable drawn from a single distribution $P_{\Omega_A \cup \Omega_B}(I)$ $\}$.

The data attachment term in (1) reads $E^{\mathrm{DA}}[\phi] = -\mathrm{LL}[\phi]$, minus the log-likelihood ratio of the two hypothesis, such that minimizing $E^{\mathrm{DA}}[\phi]$ maximizes the likelihood of a given contour (Sarti et al. (2005); Kim et al. (2005)):

$$\mathrm{LL} = \log\left(\frac{\prod_{\mathbf{x}\in\Omega_A} P_A(I_{\mathbf{x}}) \prod_{\mathbf{x}\in\Omega_B} P_B(I_{\mathbf{x}})}{\prod_{\mathbf{x}\in\Omega_A\cup\Omega_B} P_{\Omega_A\cup\Omega_B}(I_{\mathbf{x}})}\right). \tag{2}$$

$$= \sum_{\mathbf{x}\in\Omega_A} \log P_A(I_{\mathbf{x}}) + \sum_{\mathbf{x}\in\Omega_B} \log P_B(I_{\mathbf{x}}) + \mathrm{cst} \tag{3}$$

$$= |\Omega_A| \sum_I P_A(I) \log P_A(I) + (\cdots)_{A\leftrightarrow B} + \mathrm{cst} \tag{4}$$

$|\Omega_A|$ is the volume of Ω_A. The term in $P_{\Omega_A\cup\Omega_B}(I)$ is constant (cst) with respect to the contour and thus can be neglected in the minimization. In the last line, the sum over coordinates $\mathbf{x}$ is recast upon a sum over intensity I.

We estimate the distributions $\hat{P}_A(I)$ (and $\hat{P}_B(I)$) of the log-compressed envelope $I(\mathbf{x})$ using a Parzen estimate (Kim et al. (2005); Michailovich et al. (2007)):

$$\hat{P}_A(I) = \int_{\Omega_A} \mathrm{d}\mathbf{x}\, K_\lambda(I_{\mathbf{x}} - I)/|\Omega_A| \tag{5}$$

where K_λ is a centered Gaussian kernel of fixed width λ. In this context, Parzen estimates can be thought of as smoothed histograms. Parzen estimates have a better behavior than histograms to undersampling in the case where one region Ω_A or Ω_B is small.

Parametric methods

For a parametric method, it is assumed that the envelope signal $E_\mathbf{x}$ follows a given distribution, a Rayleigh distribution (Sarti et al. (2005); Lecellier et al. (2010); Pereyra et al. (2015)) or a Nakagami distribution:

$$P_\sigma^{\text{Rayleigh}}(E) = \frac{E}{\sigma^2} \mathrm{e}^{-\frac{E^2}{2\sigma^2}} \tag{6}$$

$$P_{m,\omega}^{\text{Nakagami}}(E) = \frac{2m^m E^{2m-1}}{\Gamma(m)\,\omega^m} \mathrm{e}^{-\frac{mE^2}{\omega}} \tag{7}$$

where the parameters, σ_A or ω_A and m_A, are estimated in each region, according to:

$$\hat{\sigma}_A^2 = \frac{1}{2|\Omega_A|} \sum_{\mathbf{x}\in\Omega_A} E_\mathbf{x}^2, \quad \hat{\omega}_A = 2\hat{\sigma}_A^2, \tag{8}$$

$$\hat{m}_A = \frac{\hat{\omega}_A^2}{\frac{1}{|\Omega_A|}\sum_{\mathbf{x}\in\Omega_A} E_\mathbf{x}^4 - \hat{\omega}_A^2} \tag{9}$$

where $|\Omega_A|$ is the volume of the region of interest Ω_A. The probability distribution is then used as an input in the log-likelihood data attachment as in (4), replacing the intensity I with the envelope E. All details of the algorithm can be found in (Sarti et al. (2005)).

Non-adaptive gradient descent

For both the parametric and nonparametric methods, the energy is:

$$E[\phi] = \alpha \int_\mathcal{S} \mathrm{d}^2\mathbf{x} - \beta\,\mathrm{LL}[\phi] \tag{10}$$

It is minimized via gradient descent using standard techniques. A detailed derivation is given in the Appendix A. Following (Vese and Chan (2002)):

$$\partial_t \phi_\mathbf{x} = |\nabla\phi_\mathbf{x}| f_\mathbf{x} \tag{11}$$

$$f_\mathbf{x} = \alpha\kappa_\mathbf{x} + \beta\,\log\left(\frac{\hat{P}_A(I_\mathbf{x})}{\hat{P}_B(I_\mathbf{x})}\right) \tag{12}$$

$$f_\mathbf{x} = \alpha f_\mathbf{x}^{\text{reg}} + \beta f_\mathbf{x}^{\text{DA}} \tag{13}$$

In the above expression, the "time"t is the virtual parameter of the gradient descent, and $\kappa_\mathbf{x} = \mathrm{div}\left(\frac{\nabla\phi}{|\nabla\phi|}\right)$ is the curvature of the contour. The term $f_\mathbf{x}$ is split in two parts, a data part $f_\mathbf{x}^{\text{DA}}$ and a regularization part $f_\mathbf{x}^{\text{reg}}$. The descent is performed in constant steps $t = k\Delta t$ with $k = 0, 1, \ldots$ and fixed Δt. The gradient descent reads $\phi_\mathbf{x}^{k+1} = \phi_\mathbf{x}^k + \Delta t|\nabla\phi_\mathbf{x}|f_\mathbf{x}$. Since only $\alpha\Delta t$ and $\beta\Delta t$ appear, we set $\Delta t = 1$ in the rest of the article. In the non-adaptive algorithm, α, β are free parameters.

The main steps of the algorithm are summarized in Table 1. The gradient descent is performed with a robust and efficient solver based on Additive Operator Splitting (AOS) (Weickert et al. (1998)). In step 3, the level set function is reinitialized such that $|\nabla\phi_\mathbf{x}| = 1$, preserving the boundary surface $\mathcal{S}$, to improve the stability of the algorithm. It is performed with a fast-marching algorithm (Sethian (1996)).

```
1: User initializes $\phi^{k=1}$
2: for $k = 1, 2, \ldots$ until convergence do
3:     Reset $|\nabla \phi^k| = 1$
4:     if Adaptive then
5:         Set $\beta^k$ according to (14)
6:     end if
7:     Let $\phi^k_{\mathbf{x}} \leftarrow \phi^k_{\mathbf{x}} + \beta^k \, \Delta t |\nabla \phi^k| f^{\mathrm{DA}}_{\mathbf{x}}$
8:     Let $\phi^{k+1}_{\mathbf{x}} \leftarrow \phi^k_{\mathbf{x}} + \, \alpha \Delta t |\nabla \phi^k| f^{\mathrm{reg}}_{\mathbf{x}}$
9: end for
```

Table 1: Algorithm for the (AdLL) Gradient descent solver.

Adaptive gradient descent

In the adaptive algorithm, the coefficient β, instead of being fixed, is set in step 5:

$$\beta^k = \frac{T}{\Delta t \ \max_{\mathbf{x} \in \gamma_k} |f^{\mathrm{DA}}_{\mathbf{x}}|} \tag{14}$$

$$\gamma^k = \Big\{ x \in \Omega \text{ such that } \ |\phi_{\mathbf{x}}| < T \text{ and } \ \mathrm{sign}(\phi_{\mathbf{x}} f^{\mathrm{DA}}_{\mathbf{x}}) < 0 \Big\} \tag{15}$$

γ_k is a narrow-band region (Adalsteinsson and Sethian (1995)) within a distance T of the contour surface $\mathcal{S}$ (in addition, we constrain the region γ^k to contain only points where $\mathrm{sign}(\phi_{\overline{\mathbf{x}}} f^{\mathrm{DA}}_{\overline{\mathbf{x}}}) < 0$).

In the following, the threshold is set to $T = 2$. The parameter α is the only parameter of the adaptive method, which tunes the smoothness of the contour. For all cases considered in this article, α takes meaningful values in the range $[0.1, 1]$, changing between a rough and a smooth contour. The segmentation is stopped when the change in volume between two steps is less than 0.001% of the segmented volume (a few voxels for a tumor of ~ 20000 voxels). The algorithm is implemented in Matlab. The C++ subroutine for step 3 and AOS splitting in step 8 are inspired from (Zhang et al. (2008)).

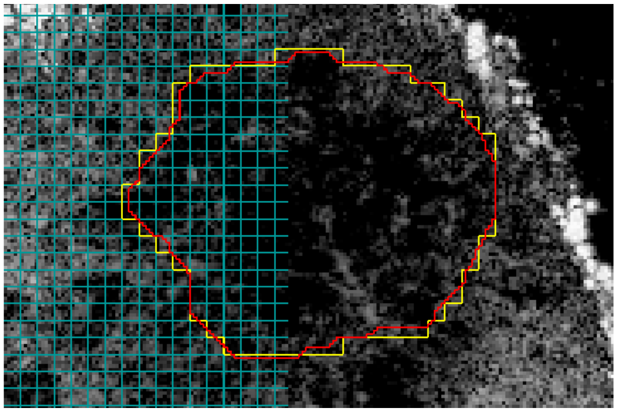

Figure 3: Slice of a 3D image with a superimposed grid of size $a = 5$. The red contour is a slice of the segmented 3D volume on the finest scale $a = 1$, the yellow contour is a slice of the volume on the $a = 4$ grid.

Grid solver

In order to reduce computation times, a crucial issue for practical use in a clinical context, we implement the level set on a coarse grid, as shown in Fig. 3. The details of the equations of evolution of the contour on a grid and their implementation can be found in Appendix B.

With the grid setup, a single segmentation takes approximately 20 s of loading and computing intensity histograms within the cells and 20 s for the segmentation. This performance is satisfactory for post-processing analysis, especially since the code is not optimized.

Data acquisition

We evaluate the performance of the segmentation algorithm on 9 lesion cases, comparing the results with manual expert segmentation. The skin tumor images have been acquired at the Melanoma Skin Cancer Clinic (Hamilton Hill, Australia) on a Dermcup 50MHz ultrasound imaging system (Atys Medical, Soucieu-En-Jarrest, France). The Dermcup 50 MHz imaging system from the company Atys Medical has a mechanical probe with a single transducer element. The transducer position is controlled along two orthogonal axis, so that a 3D field of view is acquired, of size 16 × 16 mm and 4 mm in depth. It has an axial resolution of 25 μm and a lateral resolution of 50 μm. This study is retrospective chart research on images for which written consent was obtained. All images have been de-identified. In accordance with clause (5.1.22) of the National Statement on Ethical Conduct in Human Research, self-approval has been deemed appropriate at the Melanoma Skin Cancer Clinic.

For all cases, reference volumes are drawn, based on an expert segmentation in vertical slices in the $x-z$ orientation (z is the vertical axis). The contours are interpolated between slices, to produce a 3D volume R_{m1}. To evaluate the dependency of this procedure on the slice orientation, a manual segmentation is also performed by the expert on vertical $y-z$ planes and volumes are reconstructed in a similar fashion, giving a second reference R_{m2}. A comparison of the R_{m1} and R_{m2} volumes gives an estimate of the manual segmentation errors related to the choice of the slice orientation and of the intra-operator variability.

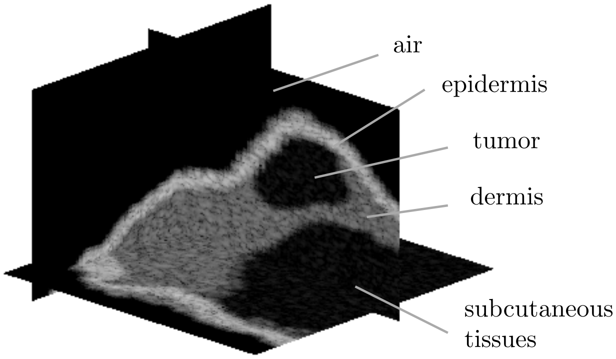

Figure 4: 3D view of a synthetic image representing an hypo-echoic tumor in a skin tissue.

Synthetic images of the skin

In order to provide a quantitative comparison of the parametric segmentation methods with the log-likelihood nonparametric segmentation method, we produce 3D synthetic ultrasound images of the skin. The images are obtained using a simple model of the imaging system, assuming that the received signal is a convolution by constant a point-spread function of the map of scatterers in the medium responsible for the back-scattering of the ultrasound (Sciolla et al. (2015a)). This is equivalent to a traditional simulation such as Field II (Jensen (1996)) with the additional assumption that the constant point spread function is independent of the depth. The scatterers map of the medium is generated at random using realistic maps of the tumor, epidermis and dermis shapes. Fig. 4 shows an example of a synthetic image, measuring 252 × 216 × 168.

We generate 13 different phantom volumes with different scatterers density in the dermis and subcutaneous tissue layers. In Fig. 5(a), both have the same density, whereas in Fig. 5(b), (c) the epidermis and subcutaneous tissue are more contrasted with respect to the dermis. In case (a), both the tumor and background region (epidermis, dermis and subcutaneous tissue, excluding the black air area) fit a Rayleigh distribution, as shown in the fit in Fig. 5(d). On the contrary, in case (b) and (c) the distribution in the background region deviates from a Rayleigh distribution, see Fig. 5(e) and (f).

Results

Metrics

In the next sections, the accuracy of the segmentation is measured with the Dice coefficient (Sørensen-Dice coefficient), an indicator of the accuracy of a segmented volume Ω with respect to a reference volume R, computed as $D(\Omega, R) = 2|\Omega \cap R|/(|\Omega| + |R|)$. We also compute the Mean Absolute Distance, MAD, which measures the accuracy of the boundary. For any voxel $\mathbf{x}$ in the boundary $\partial\Omega$ of Ω, we call $d^R(\mathbf{x})$ the distance of $\mathbf{x}$ to the closest point in R. $N_{\partial\Omega}$ is the number of voxels in $\partial\Omega$. The mean distance is $\mathrm{MAD} = \frac{1}{N_{\partial\Omega}} \sum_{\mathbf{x}\in\partial\Omega} d^R(\mathbf{x})$.

Accuracy of the grid solver

We evaluate the accuracy of the segmentation on a grid as a function of the grid size a. The grid size is $a \times a \times 2a$ (in pixels) with $a = \{1, 2, 3, 4, 5\}$. For this test, we use one image from the clinical dataset, and compare the segmentation on the grid to two reference contours drawn manually by the expert R_{m1}, R_{m2}. As shown in Fig. 6(a), the accuracy of the segmented volumes Ω with respect to the two references is almost constant with respect to a. The time of the segmentation decreases rapidly with increasing grid size $a \geq 3$; as shown in Fig. 6(b). The segmentation in the following is done with a grid of size $a = 3$. This yields a speedup of a factor of 100 compared with the finest grid of size $a = 1$, with almost no loss in precision. The good visual agreement between the segmentation on grid size $a = 3$ is shown in Fig. 6(c).

Segmentation of synthetic images

We make a comparison of the proposed adaptive nonparametric method (AdLL) with a parametric Rayleigh segmentation (Ra) widely used on ultrasound images and the more advanced parametric Nakagami segmentation (Na). The comparison is made on the synthetic 3D ultrasound images described above.

The results of (Ra) and (AdLL) segmentation are shown in Fig. 5(a-c). The (Ra) method diverges on case (c), the contour encloses the tumor, dermis and subcutaneous area. This is explained by the nature of the intensity distribution outside of the tumor, which is very different from a Rayleigh distribution, see Fig. 5(d-f). In order to measure this effect quantitatively, we introduce the Kolmogorov-Smirnov statistic to measure the difference between the distribution $P_B(I)$ and a fit to a Rayleigh distribution $P_B^R(I)$ (with σ_B computed from (7)), $KS = \sup_I \left|C_B(I) - C_B^R(I)\right|$. Here $C_B(I)$ $(C_B^R(I))$

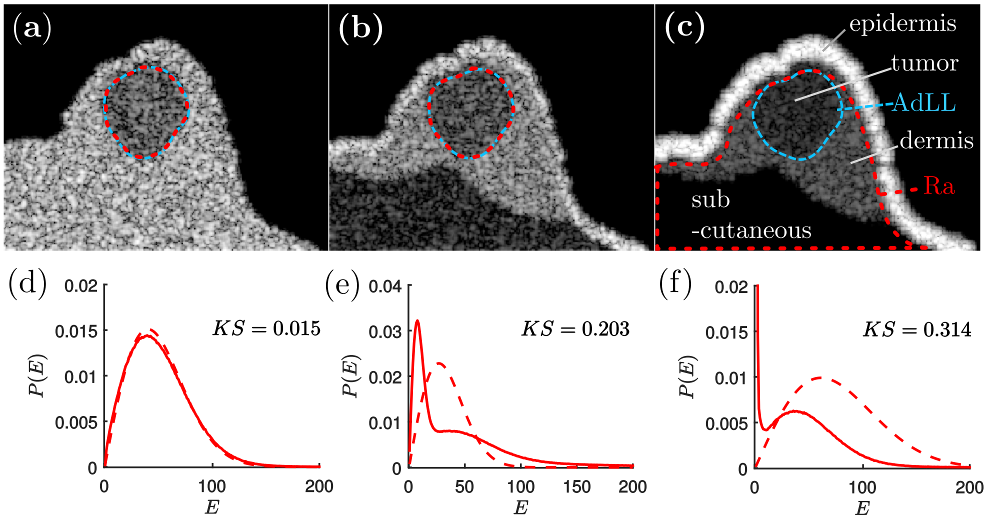

Figure 5: **(a,b,c)** Slice of the synthetic tumor image. The Rayleigh (Ra) segmentation is illustrated by the red dotted line, the (AdLL) segmentation by the blue dashed line. The (AdLL) segmentation correctly identifies the tumor for all cases, whereas the (Ra) 'leaks' into the whole dermis area for case (c). **(d,e,f)** Histograms $P(E)$ of the envelope E outside of the tumor (excluding the black air area) and fit to a Rayleigh distribution (dashed line). The KS index of the mismatch increases from (d) to (f).

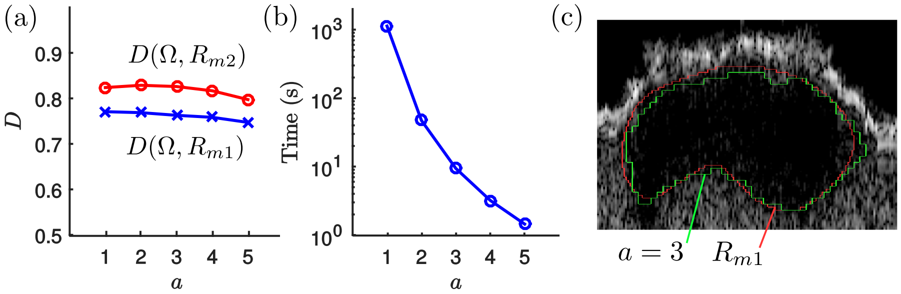

Figure 6: Properties of the segmentation on different grids with size $a = \{1, 2, 3, 4, 5\}$ on tumor 3. **(a)** Dice coefficient $D(\Omega, R_{m1})$ of the segmented volume with respect to reference contour R_{m1} (crosses) and R_{m2} (circles). **(b)** Time of segmentation in seconds versus a. **(c)** Slice of a tumor and segmented contours for a grid size $a = 3$ and the reference contour R_{m1}.

stands for the cumulative distribution function of $P_B(I)$ ($P_B^R(I)$). Roughly speaking, in the present context the fit to a Rayleigh distribution is satisfactory for $KS \leq 0.05$.

The KS fit of the background region to a Rayleigh distribution is measured for all 13 synthetic images, and plotted in Fig. 7(a), the data is found in Table 2. The image Fig. 5(a) is the first $i = 1$ and the image in Fig. 5(c) is the last $i = 13$, all other synthetic images are intermediate cases.

In Fig. 7(b), the Dice coefficient for (AdLL) method is shown, demonstrating the accuracy of the algorithm on all synthetic images. In Fig. 7(c-d), the Dice coefficient for the (Ra) and (Na) segmentation is shown. For very low KS, the (Ra) and (Na) segmentation can be slightly more accurate than the (AdLL) segmentation. This can be attributed to the fact that a parametric law has more predictive power than a non-parametric estimate, provided the data actually follows the assumed law. Yet, the (Ra) and (Na) segmentation sometimes diverge for too large KS. This instability is caused by a region of high echogenicity, as the epidermis region in clinical images, or regions of low echogenicity as the subcutaneous tissues.

To summarize the results, we find that the log-likelihood non-parametric method is able to discriminate two tissues in a more robust way than parametric methods. The parametric methods may fail whenever the tissue distribution is too different from the assumed parametric form, as expected. On the other hand, the non-parametric method typically has a larger computational cost than the simple parametric methods. Notice however that using more sophisticated parametric laws, such as statistical mixtures, would make the parametric laws slower than the non-parametric method. Finally, the non-parametric method is applied on the log-compressed image, naturally defined in a limited range of values such as $[0, 255]$. On the contrary, the parametric methods considered here are based on the envelope signal, which spans orders of magnitudes - typically $[10^{-2}, 10^5]$ in our images, making the whole processing chain more sensitive to outliers.

Results on the clinical dataset

The 9 lesions are segmented using the (AdLL) method, the log-likelihood method (LL) with a non-adaptive β, the Rayleigh method (Ra) and the Nakagami method (Na). For all methods, the parameter α, tuning the smoothness of the contour, is sampled over the range $[0.2, 0.8]$ and the best segmentation is kept. For the non-adaptive methods (LL) and (Ra), the parameter β is also optimized in the range of values $\beta \in [0.3, 1.5]$. We also perform a segmentation with a fixed set of parameters on the whole dataset. The parameters are chosen to maximize the average (over lesions) Dice coefficient.

The segmented volumes and original data for several lesions are depicted in Fig. 8. Figure 8(a) shows the result of the (Ra) segmentation. This is an example of a diverging (Ra) segmentation: the whole dermis region is segmented instead of the tumor, which is correctly delimited using (AdLL). The result is qualitatively similar to the divergence found on synthetic images in Fig. 5(c). In Fig. 8(b)

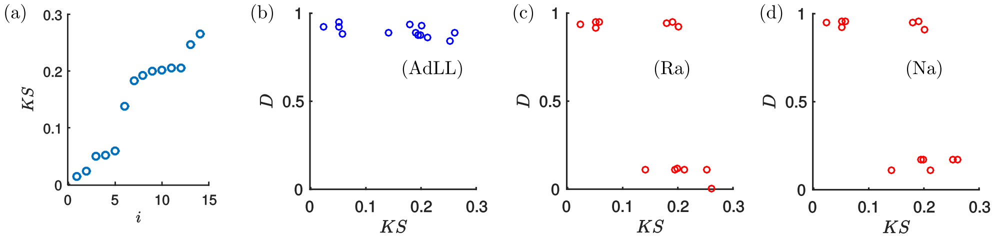

Figure 7: **(a)** KS divergence of the fit of the background region to a Rayleigh distribution, for all synthetic images $i = 1, \dots, 13$. **(b)** Dice coefficient D for the (AdLL) segmentation as a function of the KS fit to a Rayleigh (plot **(a)**), giving good agreement on all 13 synthetic images. **(c)** Dice coefficient D for the Rayleigh segmentation (Ra), versus KS. **(d)** Dice coefficient D for the Nakagami segmentation (Na), versus KS. (Ra) and (Na) segmentation fail for some synthetic images with a large KS, giving a very low Dice coefficient.

i	1	2	3	4	5	6	7	8	9	10	11	12	13
KS	0.023	0.052	0.053	0.059	0.142	0.181	0.191	0.195	0.198	0.200	0.212	0.253	0.261
D(AdLL)	0.859	0.858	0.903	0.792	0.805	0.880	0.799	0.779	0.780	0.863	0.763	0.730	0.796
D(Ra)	0.880	0.840	0.902	0.907	0.060	0.888	0.906	0.060	0.061	0.859	0.059	0.060	0.000
D(Na)	0.901	0.858	0.917	0.917	0.060	0.908	0.919	0.095	0.095	0.838	0.059	0.094	0.094

Table 2: Data for Figure 7. Except for cases $i = 4, 7$, the AdLL segmentation is as good as the Rayleigh or the Nakagami segmentation, and yields correct results (non-divergent) on a broader range of cases.

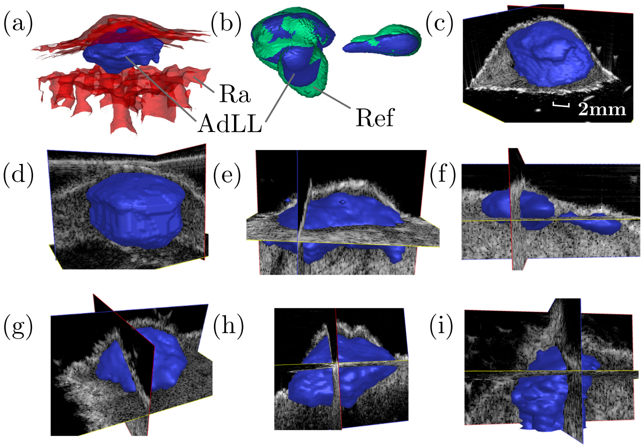

Figure 8: **(a)** (Ra) and (AdLL) segmented volumes on tumor 1. The (Ra) volume approximately covers the dermis area, which is not the intended behavior. **(b)** (AdLL) segmented volume and manual expert reference volume on tumor 4, made of two disconnected volumes. **(c-i)** (AdLL) segmented volumes on the whole clinical dataset.

the (AdLL) segmented volume is shown with one reference contour on tumor 4, which has two disconnected parts. Figure 8(c-i) depicts the segmented volume using (AdLL).

The Dice coefficient $D(\Omega, R_{m1})$ is computed for each segmented volume Ω with respect to the first expert segmentation R_{m1}. The Dice coefficients are shown in Fig. 9. For the non-adaptive methods, (LL), (Ra) and (Na), Fig. 9 shows the best Dice coefficent $\max_\beta D$ and average Dice coefficent $\langle D\rangle_\beta$ for all values of β sampled. Table 3 contains the average over all tumors of the Dice coefficients from Fig. 9, for the optimized case and for a fixed set of parameters, as well as the computational time. Table 4 contains the same quantities, computed with the MAD criterion. Finally, the mutual Dice coefficient $D(R_{m2}, R_{m1})$ is referred to as intra-operator variability.

Discussion

Fig. 9 shows that some lesions are easier to segment accurately than others. The accuracy tends to be lower in tumors with low contrast and/or complex shapes as in case 4. As previously observed on the synthetic data, the (AdLL) or (LL) method is much more consistent across different cases, although the (Ra) and (Na) methods may be slightly more accurate in some cases. Fig. 9 also shows the average Dice over several β parameters, which illustrates how β impacts the accuracy of the segmentation. Since in real world applications the parameters must be fixed by the user, a large variability is not desirable. In particular, the (Na) method has larger variability than the (LL) method.

According to Table 3 and 4, for optimized parameters, the proposed method (AdLL) has a Dice coefficient of $D = 0.79$ and MAD $= 167\mu$m (to be compared with the lateral resolution of 50μm) on average over all tumors, (with minimum $D = 0.71$ and maximum $D = 0.87$, not shown in Table 3), values indicating a satisfactory segmentation, with accuracy similar to the non-adaptive log-likelihood method $D = 0.81$. On the other hand, the (Ra) and (Na) segmentation have lower average Dice, $D = 0.621$ and $D = 0.683$, and higher MAD $= 175$ and 262μm. For fixed sets of parameters, which would be the realistic clinical setting, the same tendency is observed - except that the adaptive method (AdLL) has better accuracy than the log-likelihood method but slightly higher MAD. In terms of computation times, shown in Table 3, the methods rank as (Ra), (AdLL), (Na), (LL) from fastest to slowest. The (AdLL) method actually has similar speed to that of the fastest (Ra) method. Compared with (LL), the speed of (AdLL) can be attributed to the adaptive algorithm, which increases the rate of convergence of the gradient descent.

Summarizing the results shown above, the adaptive algorithm (AdLL) reaches optimal results (with respect to the other algorithms tested) with one free parameter less.

The segmented volumes and the clinical images can be visualized online using our web demonstrator at: `http://www.atyscrea.insa-lyon.fr/en/demo_echo3D`. It is based on the *desk* library (Jacinto et al. (2012)).

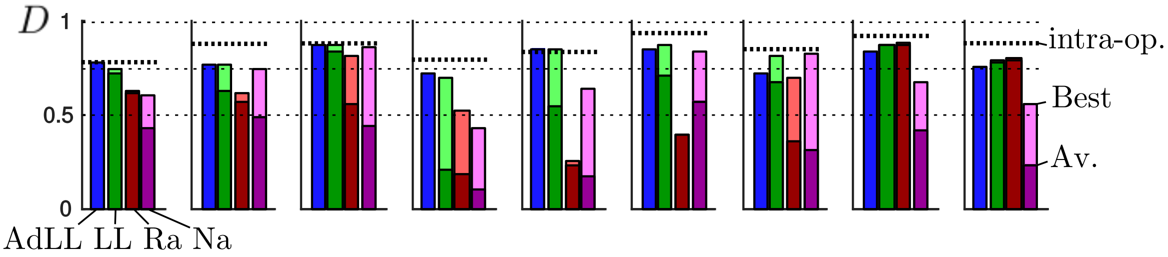

Figure 9: Dice coefficient D for the optimized α, β parameter on the 9 clinical images. The segmentation methods are (AdLL) in blue, (LL) in green, (Ra) in red, (Na) in purple. For the (LL), (Ra) and (Na) methods, the average Dice coefficient over all tested β is shown, as well as the overall best result. The thick dashed line indicates $D(R_{m1}, R_{m2})$, giving an indication of the intra-operator variability.

	(AdLL)	(LL)	(Ra)	(Na)	Intra-op.
Opt. Dice	0.793	**0.808**	0.621	0.683	0.861
Fixed Dice	**0.756**	0.732	0.586	0.535	0.861
Time (s)	13.0	18.5	12.1	14.7	300

Table 3: Average over all cases of the Dice index $D(\Omega, R_{m1})$, for all methods. The row "Opt. Dice" is a segmentation with all parameters α, β (when applicable) optimized on each single case. The row "Fixed. Dice" is a segmentation with a fixed set of parameters applied to all cases (different depending on the method). The average time of a single segmentation is indicated in the last row. The intra-operator variability $D(R_{m2}, R_{m1})$ is indicated in the last column "Intra-op.", and the corresponding time refers to the average time for a manual segmentation.

	(AdLL)	(LL)	(Ra)	(Na)	Intra-op.
Opt. MAD	167	**113**	175	262	82
Fixed MAD	184	**177**	203	347	82

Table 4: Same as Table 3 with the Mean Absolute Distance (unit: μm). The intra-operator variability $\mathrm{MAD}(R_{m2}, R_{m1})$ is indicated in the last column "Intra-op.".

Conclusions

We propose a flexible, robust and fast segmentation algorithm for 3D ultrasound skin tumor images based on nonparametric estimates of the intensity distributions in tumors. We have shown that this algorithm performs well on a clinical dataset of 9 patients, with an average Dice coefficient of $D = 0.79$. The method outperforms two parametric methods with Rayleigh or Nakagami distributions, while having similar computational time. Using a coarse grid, the speed of segmentation is increased by a factor of 100. As a result, a segmentation takes less than a minute, making it a suitable candidate for use in clinical routine. The results also show that the intra-operator variability is high, which indicates that automated segmentation tools are of interest. Indeed, assessing the tumor shape and volumes from slices is an arduous task, whereas automated algorithms can make direct use of the 3D data and include multi-directional information. Moreover, the result of the automatic segmentation is not very sensitive to the initialization, granting good reproducibility and weak operator dependency, an issue often encountered in ultrasound. In future works, an interesting question is to assert the performance of the proposed segmentation on phantoms appropriate for high-frequency ultrasound (Jacquet et al. (2015)).

Acknowledgments

We thank P. Ceccato for her contribution in early stages of this work. This work was funded by the ANR-14-LAB3-0006-01 LabCom AtysCrea and was supported by the LABEX CeLyA (ANR-10-LABX-0060) of Université de Lyon, within the "Investissements d'Avenir" program (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR).

Appendix A. Derivation of the gradient flow

The total energy (1) is minimized by gradient descent on the discrete variables $\phi_{\mathbf{x}}$. To derive Eq. (A.7), we rewrite (5) with the Heaviside distribution θ as:

$$\hat{P}_A(I) = \int \mathrm{d}\mathbf{x}\, \theta(\phi_{\mathbf{x}}) K_\lambda(I(\mathbf{x}) - I)/(\int \mathrm{d}\mathbf{x}\, \theta(\phi_{\mathbf{x}})) \tag{A.1}$$

$$\hat{P}_B(I) = \int \mathrm{d}\mathbf{x}\, \theta(-\phi_{\mathbf{x}}) K_\lambda(I(\mathbf{x}) - I)/(\int \mathrm{d}\mathbf{x}\, \theta(-\phi_{\mathbf{x}})) \tag{A.2}$$

The derivative of LL in Eq. (4) yields:

$$\frac{\delta\left(|\Omega_A| \sum_I \hat{P}_A(I) \log \hat{P}_A(I)\right)}{\delta\phi_{\mathbf{x}}} = \delta(\phi_{\mathbf{x}}) \sum_I \hat{P}_A(I) \log \hat{P}_A(I) \ldots \tag{A.3}$$

$$+ |\Omega_A| \sum_I \delta(\phi_{\mathbf{x}}) \frac{\left(K_\lambda(I - I_{\mathbf{x}}) - \hat{P}_A(I)\right)}{|\Omega_A|} \times \left[\log \hat{P}_A(I) + 1\right] \tag{A.4}$$

In the above, δ is the Dirac delta function.
We now use the approximation:
$K_\lambda(I - I_x) \to \delta(I - I_x)$ in the sum over I to obtain:

$$\frac{\delta\left(|\Omega_A| \sum_I \hat{P}_A(I) \log \hat{P}_A(I)\right)}{\delta\phi_{\mathbf{x}}} \simeq \delta(\phi_{\mathbf{x}}) \log \hat{P}_A(I_{\mathbf{x}}) \tag{A.5}$$

In contrast to (Kim et al. (2005)), we have used the approximation $K_\lambda(I - I_{\mathbf{x}}) \to \delta(I - I_{\mathbf{x}})$ in the sum over intensity appearing in the computation of (A.7). The computational cost of the algorithm is reduced, with minor numerical difference, since the Gaussian kernel $K_\lambda(I - I_{\mathbf{x}})$ is narrow. A similar calculation holds for the second region Ω_B. Following (Vese and Chan (2002)) for the regularization term, the gradient descent reads $\partial_t \phi_{\mathbf{x}} = -\delta E/\delta\phi_{\mathbf{x}}$, with:

$$\delta E/\delta\phi_{\mathbf{x}} = -\delta(\phi_{\mathbf{x}}) f_{\mathbf{x}} \tag{A.6}$$

$$f_{\mathbf{x}} = \kappa_{\mathbf{x}} + \beta \log\left(\frac{\hat{P}_A(I_{\mathbf{x}})}{\hat{P}_B(I_{\mathbf{x}})}\right) \tag{A.7}$$

This equation is then recast onto (A.7), changing the Dirac term $\delta(\phi_{\mathbf{x}})$ into a $|\nabla\phi_{\mathbf{x}}|$ following (Vese and Chan (2002)).

Appendix B. Grid solver

Let us now derive the equations of evolution of the contour on a grid. Let $\overline{\mathbf{x}}$ be the coordinates of a grid cell. We define the cell volume $\Omega_{\overline{\mathbf{x}}} = \{\mathbf{x} \in \overline{\mathbf{x}}\}$ and $\hat{P}_{\overline{\mathbf{x}}}(I)$ the Parzen estimate of the distribution on the volume $\Omega_{\overline{\mathbf{x}}}$, defined in Eq. (5) with $\Omega_{\overline{\mathbf{x}}}$ instead of Ω_A. For the derivation we assume that the curvature $\kappa_{\mathbf{x}}$ and $\phi_{\mathbf{x}}$ vary weakly in the whole cell $\overline{\mathbf{x}}$. Essentially, this means that the underlying contour can be represented with a reasonable precision on the coarse grid. The gradient descent in discrete time k reads:

$$\phi_{\overline{\mathbf{x}}}^{k+1} = \phi_{\overline{\mathbf{x}}}^{k} - \Delta t \; \delta E / \delta \phi_{\overline{\mathbf{x}}} \tag{B.1}$$

$$\delta E / \delta \phi_{\overline{\mathbf{x}}} = \sum_{\mathbf{x} \in \overline{\mathbf{x}}} \delta E / \delta \phi_{\mathbf{x}} = -\delta(\phi_{\overline{\mathbf{x}}}) |\Omega_{\overline{\mathbf{x}}}| f_{\overline{\mathbf{x}}} \tag{B.2}$$

$$f_{\overline{\mathbf{x}}} = \alpha \kappa_{\overline{\mathbf{x}}} + \beta \sum_I \hat{P}_{\overline{\mathbf{x}}}(I) \; \log\left(\frac{\hat{P}_A(I)}{\hat{P}_B(I)}\right) \tag{B.3}$$

$$= \alpha f_{\overline{\mathbf{x}}}^{\mathrm{reg}} + \beta f_{\overline{\mathbf{x}}}^{\mathrm{DA}} \tag{B.4}$$

To derive Eq.(B.3), one must change the sum over $\mathbf{x}$ into a sum over intensity I:

$$\sum_{\mathbf{x} \in \overline{\mathbf{x}}} \log\left(\frac{\hat{P}_A(I_{\mathbf{x}})}{\hat{P}_B(I_{\mathbf{x}})}\right) \simeq \sum_I |\Omega_{\overline{\mathbf{x}}}| \hat{P}_{\overline{\mathbf{x}}}(I) \log\left(\frac{\hat{P}_A(I)}{\hat{P}_B(I)}\right) \tag{B.5}$$

Using the fact that $|\Omega_{\overline{\mathbf{x}}}| \hat{P}_{\overline{\mathbf{x}}}(I)$ is (an estimate of) the number of points $\mathbf{x} \in \overline{\mathbf{x}}$ where the intensity is equal to I.

One can see that the map $f_{\overline{\mathbf{x}}}$ is similar (with a multiplicative factor $|\Omega_{\overline{\mathbf{x}}}|$) to the map on the finer scale (A.7) *averaged* over the cell $\overline{\mathbf{x}}$. The evolution contains two parts f^{reg} and f^{DA} to describe the discretized solver. It is important to note that the full information from the original image in region Ω_A, Ω_B and in the cell $\Omega_{\overline{\mathbf{x}}}$ is contained into $\hat{P}_A(I)$, $\hat{P}_B(I)$ and $\hat{P}_{\overline{\mathbf{x}}}(I)$. As a consequence, this grid solver allows the contour to move on the grid with large steps, while retaining all the information from the original image. The improvement in computational time is high for a minimal loss in precision, as shown in the section on the accuracy of the grid solver.

Adalsteinsson D, Sethian JA. A fast level set method for propagating interfaces. Journal of Computational Physics, 1995;118:269 – 277.

Anquez J, Angelini E, Grange G, Bloch I. Automatic segmentation of antenatal 3-d ultrasound images. Biomedical Engineering, IEEE Transactions on, 2013;60:1388–1400.

Cardinal MH, Meunier J, Soulez G, Maurice R, Therasse E, Cloutier G. Intravascular ultrasound image segmentation: a three-dimensional fast-marching method based on gray level distributions. Medical Imaging, IEEE Transactions on, 2006;25:590–601.

Chan T, Shen J. Image Processing And Analysis: Variational, Pde, Wavelet, And Stochastic Methods. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2005.

Dahdouh S, Angelini ED, Grangé G, Bloch I. Segmentation of embryonic and fetal 3d ultrasound images based on pixel intensity distributions and shape priors. Medical Image Analysis, 2015;24:255 – 268.

Destrempes F, Cloutier G. A critical review and uniformized representation of statistical distributions modeling the ultrasound echo envelope. Ultrasound in Medicine and Biology, 2010;36:1037 – 1051.

Gao Y, Tannenbaum A, Chen H, Torres M, Yoshida E, Yang X, Wang Y, Curran W, Liu T. Automated skin segmentation in ultrasonic evaluation of skin toxicity in breast cancer radiotherapy. Ultrasound in Medicine & Biology, 2013;39:2166 – 2175.

Jacinto H, Kéchichian R, Desvignes M, Prost R, Valette S. A web interface for 3D visualization and interactive segmentation of medical images. In: Proceedings of the 17th International Conference on 3D Web Technology, 2012. pp. 51–58.

Jacquet JR, Levassort F, Ossant F, Grégoire JM. 3d printed phantom for high frequency ultrasound imaging. In: Ultrasonics Symposium (IUS), 2015 IEEE International. IEEE, 2015. pp. 1–4.

Jensen JA. Field: A program for simulating ultrasound systems. In: 10TH Nordic Baltic conference on biomedical imaging, vol. 4, supplement 1, part 1:351–353, 1996. pp. 351–353.

Kim J, Fisher J, Yezzi A, Cetin M, Willsky A. A nonparametric statistical method for image segmentation using information theory and curve evolution. Image Processing, IEEE Transactions on, 2005;14:1486–1502.

Lecellier F, Fadili J, Jehan-Besson S, Aubert G, Revenu M, Saloux E. Region-based active contours with exponential family observations. Journal of Mathematical Imaging and Vision, 2010;36:28–45.

Machet L, Samimi L, Georgesco G, Mourtada Y, Naouri M, Grégoire JM, Ossant F, Patat F, Vaillant L. High resolution ultrasound imaging of melanocytic and other pigmented lesions of the skin. In: Tanabe M (Ed.), Ultrasound Imaging. InTech, 2010.

Maev RG. Advances in acoustic microscopy and high resolution ultrasonic imaging: from principles to new applications. Proc. SPIE, 2014;9040:904007–904007–6.

Michailovich O, Rathi Y, Tannenbaum A. Image segmentation using active contours driven by the bhattacharyya gradient flow. Image Processing, IEEE Transactions on, 2007;16:2787–2801.

Olivier J, Paulhac L. 3d ultrasound image segmentation: Interactive texture-based approaches. In: Medical Imaging. InTech, 2011. pp. 44–65.

Pereyra M, Batatia H, McLaughlin S. Exploiting information geometry to improve the convergence of nonparametric active contours. Image Processing, IEEE Transactions on, 2015;24:836–845.

Pereyra M, Dobigeon N, Batatia H, Tourneret J. Segmentation of skin lesions in 2-d and 3-d ultrasound images using a spatially coherent generalized rayleigh mixture model. Medical Imaging, IEEE Transactions on, 2012;31:1509–1520.

Pons G, Martí J, Martí R, Noble J. Simultaneous lesion segmentation and bias correction in breast ultrasound images. In: Vitrià J, Sanches J, Hernández M (Eds.), Pattern Recognition and Image Analysis. Vol. 6669 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2011. pp. 692–699.

Qiu W, Yuan J, Kishimoto J, McLeod J, Chen Y, de Ribaupierre S, Fenster A. User-guided segmentation of preterm neonate ventricular system from 3-d ultrasound images using convex optimization. Ultrasound in Medicine & Biology, 2015;41:542 – 556.

Rahmati P, Adler A, Hamarneh G. Mammography segmentation with maximum likelihood active contours. Medical Image Analysis, 2012;16:1167 – 1186.

Sarti A, Corsi C, Mazzini E, Lamberti C. Maximum likelihood segmentation of ultrasound images with rayleigh distribution. Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on, 2005;52:947–960.

Schmid-Wendtner M, Dill-Müller D. Ultrasound technology in dermatology. Semin. Cutan. Med. Surg., 2008;27:44–51.

Sciolla B, Ceccato P, Cowell L, Dambry T, Guibert B, Delachartre P. Segmentation of inhomogeneous skin tissues in high-frequency 3d ultrasound images, the advantage of non-parametric log-likelihood methods. In: International Congress on Ultrasonics. Metz, France, 2015a.

Sciolla B, Ceccato P, Dambry T, Guibert B, Delachartre P. A comparison of non-parametric segmentation methods. In: Colloque GRETSI. Lyon, France, 2015b.
URL `https://hal.archives-ouvertes.fr/hal-01307318`

Sethian JA. A fast marching level set method for monotonically advancing fronts. Proceedings of the National Academy of Sciences, 1996;93:1591–1595.

Slabaugh G, Unal G, Fang T, Wels M. Ultrasound-specific segmentation via decorrelation and statistical region-based active contours. In: Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on. Vol. 1, 2006. pp. 45–53.

Vese LA, Chan TF. A multiphase level set framework for image segmentation using the mumford and shah model. Int. J. Comput. Vision, 2002;50:271–293.

Vogt M, Ermert H. High-Frequency Ultrasonic Systems for High-Resolution Ranging and Imaging. Wiley-VCH Verlag GmbH & Co. KGaA, 2013. pp. 93–123.

Weickert J, Romeny B, Viergever M. Efficient and reliable schemes for nonlinear diffusion filtering. Image Processing, IEEE Transactions on, 1998;7:398–410.

Xu H, Mandal M. Epidermis segmentation in skin histopathological images based on thickness measurement and k-means algorithm. EURASIP Journal on Image and Video Processing, 2015;2015.

Zhang Y, Matuszewski B, Shark L, Moore C. Medical image segmentation using new hybrid level-set method. In: BioMedical Visualization, 2008. MEDIVIS '08. Fifth International Conference, 2008. pp. 71–76.